

Skew-symmetric approximations of posterior distributions

Francesco Pozza¹ , Daniele Durante^{1,2}  and Botond Szabo^{1,2} 

¹Bocconi Institute for Data Science and Analytics, Bocconi University, Via Roentgen 1, 20136 Milan, Italy

²Department of Decision Sciences, Bocconi University, Via Roentgen 1, 20136 Milan, Italy

Address for correspondence: Francesco Pozza, Bocconi Institute for Data Science and Analytics, Bocconi University, Via Roentgen 1, 20136 Milan, Italy. Email: francesco.pozza2@unibocconi.it

Abstract

Popular deterministic approximations of posterior distributions from, e.g. the Laplace method, variational Bayes and expectation-propagation, generally rely on symmetric families, often taken to be Gaussian. This choice facilitates optimization and inference, but typically affects the quality of the approximation. In fact, even in basic parametric models, posterior distributions often display asymmetries that yield bias and reduced accuracy when considering symmetric approximations. Recent research has moved towards more flexible approximations incorporating skewness. However, current solutions are often model specific, lack general supporting theory, increase the computational complexity of the optimization problem, and do not provide broadly applicable solutions to incorporate skewness in any symmetric approximation. This article addresses such a gap by introducing a general and provably optimal strategy to perturb any symmetric approximation of a generic posterior distribution. The proposed solution is derived without additional optimization steps, and yields a similarly tractable approximation within the class of skew-symmetric densities that enhances the finite sample accuracy of the original symmetric counterpart. Furthermore, under suitable assumptions, it improves the convergence rate to the exact posterior by at least a \sqrt{n} factor, in asymptotic regimes. These advancements are illustrated in numerical studies focusing on skewed perturbations of state-of-the-art Gaussian approximations.

Keywords: expectation-propagation, Laplace approximation, skew-symmetric distribution, total variation distance, variational Bayes, α -divergence

1 Introduction

Deterministic approximations of intractable posterior distributions provide a popular alternative to general Markov chain Monte Carlo methods within Bayesian inference (e.g. [Blei et al., 2017](#); [Minka, 2001](#); [Rue et al., 2009](#); [Walker, 1969](#)). Noteworthy examples in such a class are the approximations obtained under the classical Laplace method and its generalizations ([Rossell et al., 2021](#); [Rue et al., 2009](#); [Tierney & Kadane, 1986](#); [Walker, 1969](#)), variational Bayes (VB) ([Blei et al., 2017](#); [Opper & Archambeau, 2009](#)), including its widely applicable stochastic black-box extensions ([Kucukelbir et al., 2017](#); [Ranganath et al., 2014](#)), and expectation-propagation (EP) ([Minka, 2001](#); [Vehtari et al., 2020](#)); see also [Bishop \(2006, Chapters 4.4 and 10\)](#) and [Chopin and Ridgway \(2017\)](#) for a general overview. Albeit derived under different arguments and optimization strategies, all these solutions share a common trade-off between the attempt to facilitate optimization and posterior inference via a sufficiently tractable approximation, and the need to avoid an overly simplified characterization of the intractable posterior distribution, which may undermine inference accuracy. The overarching focus to date has been on addressing the first objective at the expense of the second. This has been accomplished by searching for tractable approximating densities, often within the Gaussian family. Under the classical Laplace

Received: September 23, 2024. Revised: October 21, 2025. Accepted: October 25, 2025

© The Royal Statistical Society 2026.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

method this choice is implicitly encoded in the second-order Taylor expansion of the log-posterior density (e.g. Bishop, 2006, Chapter 4.4), while for EP (e.g. Minka, 2001) and VB (e.g. Challis & Barber, 2013; Opper & Archambeau, 2009) it is enforced as a constraint within a suitable optimization problem seeking the closest approximating density to the target posterior under a selected divergence. Even when replacing such a constraint in VB with mean-field factorization assumptions (e.g. Blei et al., 2017), Gaussian approximations can still arise either as a solution of the optimization problem under routinely implemented Bayesian models, such as logit (Durante & Rigon, 2019) and probit (Consonni & Marin, 2007) regression, or as a building block to ease implementation in more general settings (e.g. Carbonetto & Stephens, 2012; Kucukelbir et al., 2017; Ray & Szabó, 2022; Rezende & Mohamed, 2015; C. Wang & Blei, 2013).

The above approximations yield state-of-the-art solutions which are implemented in routinely employed \mathbb{R} libraries (e.g. `rstan`). In addition, these solutions also find theoretical justification in the asymptotic normality of the posterior distribution ensured by Bernstein-von Mises type results (e.g. Van der Vaart, 2000) that have been extensively studied in the context of classical Laplace (e.g. Kasprzak et al., 2025; Panov & Spokoiny, 2015; Spokoiny, 2025; Spokoiny & Panov, 2025; Walker, 1969), VB (e.g. Katsevich & Rigollet, 2024; Y. Wang & Blei, 2019) and EP (e.g. Dehaene & Barthelmé, 2018) approximations. While such results yield theoretical support within asymptotic regimes, in practice the accuracy of Gaussian approximations ultimately depends on the actual shape of the posterior distribution in finite samples, which often displays non-negligible skewness that might affect inference accuracy when relying on Gaussian approximations. In fact, although the design of tractable and computationally efficient skewed approximations of posterior distributions has attracted less interest to date, the available contributions along these lines have proven effective in refining the accuracy of the Gaussian counterparts (see, e.g. Anceschi et al., 2023; Challis & Barber, 2012; Durante et al., 2024; Fasano & Durante, 2022; Fasano et al., 2022; Katsevich, 2024; Ormerod, 2011; Rue et al., 2009; Ruli et al., 2016; Salomone et al., 2024; Smith et al., 2020; Tan & Chen, 2024; Zhou et al., 2024). While such gains should stimulate an increasing adoption of skewed approximations, there is still a gap between the promising methodological avenues opened by this perspective and its routine use in practice. In fact, when compared to classical approximations based on Gaussians, or on other symmetric densities, there are at least three major barriers, which still undermine the routine adoption of these skewed extensions.

- **(a) Methods:** The current solutions are developed under specific models and, even when targeting more general posterior distributions, the focus remains on improving a single strategy, often VB. Moreover, the overarching interest is on including skewness in Gaussian approximating densities via extensions of skew-normal families (e.g. Arellano-Valle & Azzalini, 2006; González-Farías et al., 2004), rather than in more general symmetric ones.
- **(b) Computation:** The schemes to derive the above approximations require further optimization steps that often yield substantial increments in computational costs. In some cases these schemes necessitate even the calculation of the exact moments for the intractable posterior, a requirement in contrast with the motivations for considering deterministic approximations.
- **(c) Theory:** Unlike for the classical versions of the Laplace method, VB and EP, most skewed extensions lack general supporting theory on the accuracy in approximating the whole posterior distribution and its functionals, both in finite samples and also in asymptotic regimes.

Two exceptions to point (c) can be found in the recent contributions by Fasano et al. (2022) and Durante et al. (2024). The former demonstrates that, in Bayesian probit regression, a partially-factorized generalization of mean-field VB leads to a unified skew-normal (SUN) (Arellano-Valle & Azzalini, 2006) approximation, which provably matches the exact posterior for fixed sample size n and dimension $d \rightarrow \infty$. The latter derives, instead, a refined version of the Bernstein-von Mises theorem (see, e.g. Van der Vaart, 2000) which replaces the classical Gaussian limiting law with a valid skewed perturbation (arising from a third-order version of the Laplace method) to obtain an improvement in the convergence rates to the target posterior by a \sqrt{n} factor, under the total variation (TV) distance; see also Katsevich (2024) for alternative higher-order skewed expansions, which, however, do not guarantee valid approximating densities. Albeit being derived in

specific models (Fasano et al., 2022), or with a focus to a single approximation strategy (Durante et al., 2024), these results suggest that the inclusion of skewness in popular symmetric approximations has the potential to substantially improve inference accuracy. However, it is still not clear if a broadly applicable, computationally tractable and theoretically supported solution addressing points (a)–(c) above actually exists and can be effectively derived.

In this article we provide a positive and innovative answer to the above question by deriving a broadly applicable strategy to perturb, at no additional optimization costs, any off-the-shelf symmetric approximation of a generic posterior distribution. Crucially, such a perturbation yields a provably more accurate skewed approximation of the target posterior, which belongs to the tractable class of skew-symmetric distributions reviewed in Section 2.2 (Azzalini & Capitanio, 2003; Genton & Loperfido, 2005; Ma & Genton, 2004; J. Wang et al., 2004). Our proposed solution follows from two fundamental results, which we derive in Section 2.3. The first (see Proposition 2), shows that the density of any generic posterior distribution admits a skew-symmetric representation decomposing this density as the product of (i) a suitably symmetrized version of the target posterior and (ii) a tractable skewing factor, which is available in closed form and does not depend on additional unknown parameters. This novel characterization, whose scope extends beyond our contribution, relates to an important existence result of skew-symmetric densities (J. Wang et al., 2004) which has been never explored within the context of Bayesian inference and approximations, despite its unique potential. In fact, our first result in Proposition 2 opens the avenues to prove a second important property (see Lemma 1), which ensures that any symmetric approximation of a posterior distribution, including the popular ones reviewed in Section 2.1, is more accurate in approximating the symmetrized version of such a posterior in its skew-symmetric representation, rather than the original posterior itself. This holds under the TV distance, Kullback–Leibler (KL) divergence (Kullback & Leibler, 1951), reverse-KL and generic α -divergences.

The aforementioned results yield, therefore, a natural strategy to derive skewed approximations of a generic posterior distribution via direct perturbation of any symmetric approximation for such a posterior. In particular, these can be obtained, at no additional optimization costs, by simply replacing the symmetrized version of the posterior density in its skew-symmetric representation with the available symmetric approximation to be improved. As proved in Section 3.1, this yields a more accurate asymmetric alternative that matches exactly the closed-form skewing factor of the target posterior, while preserving the tractability of its unperturbed symmetric counterpart. This tractability follows from the fact that the density of the resulting skewed approximation falls in the skew-symmetric family (see, e.g. Azzalini & Capitanio, 2003; Ma & Genton, 2004; J. Wang et al., 2004). Such a family admits a closed-form normalizing constant along with straightforward i.i.d. sampling schemes, which enable inference via Monte Carlo evaluation of any functional of the approximate posterior. These schemes only require simulation from the original symmetric approximation and evaluation of the skewing factor. As shown in Section 2.3, this factor crucially depends only on the ratios of posterior densities, thereby allowing cancellation of the intractable normalizing constant, and hence, direct computation without additional optimizations.

Our theoretical results within Section 3.1 show that the strategy we propose is not only tractable, but also yields skew-symmetric approximations that are provably more accurate than the original symmetric counterpart, under any sample size n and several popular divergences (i.e. the TV distance and any α -divergence, including its limiting KL and reverse-KL forms (see, e.g. Cichocki & Amari, 2010; Margossian et al., 2025)). Even more, the derived skewing factor is shown to be the optimal among those yielding approximations within the skew-symmetric class. This allows to formally interpret the proposed solution as the optimum of a minimization problem that seeks the closest approximation to the target posterior in the entire class of skew-symmetric densities generated by perturbations of the known symmetric approximation to be improved. The accuracy improvements formalized in Section 3.1 are quantified asymptotically in Section 3.2 with a focus on perturbations of specific symmetric approximations which can be shown to improve by at least a \sqrt{n} factor the convergence rate to the target posterior relative to the original unperturbed counterpart, under all the divergences considered within Section 3.1. As illustrated in the simulation studies in Section 4, the theoretical results presented in Sections 3.1 and 3.2 find empirical evidence also in practice, including in challenging high-dimensional contexts beyond the regimes considered by the asymptotic theory within Section 3.2. In Section 5, these accuracy gains are further confirmed in real-data applications under realistic Bayesian hierarchical formulations showcasing both

the relative and absolute improvements of the proposed skew-symmetric perturbation of several state-of-the-art Gaussian approximations from the Laplace method, black-box VB, and EP. In these real-data applications we also illustrate the remarkable gains in the effective sample size (ESS) when replacing such Gaussian approximations with the corresponding skew-symmetric counterparts as proposals in importance sampling (e.g. [Chopin & Papaspiliopoulos, 2020](#), Chapter 8). These improvements clarify the broad impact and applicability of the ideas presented in this article. Section 6 concludes by further stressing this point through the discussion of promising research directions motivated by the general scope of the skew-symmetric representation of posterior densities behind our contribution. Proofs can be found in the [online supplementary material](#). Code is available at https://github.com/Francesco16p/SkewApr_Post.

2 Skew-symmetric approximations of posterior distributions

Let $L(\boldsymbol{\theta}; \mathbf{y}_{1:n}) = \prod_{i=1}^n p(y_i | \boldsymbol{\theta})$ denote the likelihood that is induced by a generic statistical model characterizing the probabilistic generative mechanism of the data $\mathbf{y}_{1:n} = (y_1, \dots, y_n)$ through a family of distributions indexed by the parameter $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^d$. To perform inference on $\boldsymbol{\theta}$, Bayesian statistics defines a prior $\pi(\boldsymbol{\theta})$ for $\boldsymbol{\theta}$, which, combined with $L(\boldsymbol{\theta}; \mathbf{y}_{1:n})$, yields the corresponding posterior density $\pi_n(\boldsymbol{\theta}) := \pi(\boldsymbol{\theta} | \mathbf{y}_{1:n})$ via Bayes rule, i.e. $\pi_n(\boldsymbol{\theta}) = \pi(\boldsymbol{\theta})L(\boldsymbol{\theta}; \mathbf{y}_{1:n})/c(\mathbf{y}_{1:n})$, with $c(\mathbf{y}_{1:n}) = \int_{\Theta} \pi(\boldsymbol{\theta})L(\boldsymbol{\theta}; \mathbf{y}_{1:n}) d\boldsymbol{\theta}$. Inference under such a posterior is then performed through evaluation of functionals of interest. The feasibility of this task inherently depends on the tractability of $\pi_n(\boldsymbol{\theta})$. In fact, $\pi_n(\boldsymbol{\theta})$ is even not available analytically in most Bayesian models, due, in general, to the lack of a closed-form expression for $c(\mathbf{y}_{1:n})$. To address such an issue, standard practice relies either on Markov chain Monte Carlo schemes or on tractable deterministic approximations of $\pi_n(\boldsymbol{\theta})$. Here we consider the latter class of solutions and focus on improving those state-of-the-art symmetric approximations that routinely arise in practice.

More specifically, let $\bar{q}_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta})$ be an already-derived symmetric approximation of the posterior density $\pi_n(\boldsymbol{\theta})$, with $\boldsymbol{\theta}^*$ denoting the known symmetry point, i.e. $\bar{q}_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta}) = \bar{q}_{n,\boldsymbol{\theta}^*}(2\boldsymbol{\theta}^* - \boldsymbol{\theta})$ for all $\boldsymbol{\theta} \in \Theta$. Our main objective is to obtain a more accurate approximation $q_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta})$ from the perturbation of $\bar{q}_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta})$ via a skewing factor $w_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta})$, which does not require optimization of additional parameters, while preserving the inference tractability of the original symmetric approximation. Although the novel strategy we derive within Section 2.3 is guaranteed to improve the accuracy of any symmetric approximating density $\bar{q}_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta})$, in practice it is natural to focus on perturbing the outputs of routinely implemented Gaussian approximations from, e.g. Laplace, VB and EP; see Section 2.1 for an overview of these methods and refer to [Bishop \(2006, Chapters 4.4 and 10\)](#), [Challis and Barber \(2013\)](#), [Kucukelbir et al. \(2017\)](#), [Oppen and Archambeau \(2009\)](#) and [Vehtari et al. \(2020\)](#) for a detailed treatment. As clarified in these contributions, such symmetric approximations provide popular solutions which are effectively coded in standard platforms (e.g. STAN) and softwares (e.g. R, python and julia). Since the proposed skew-symmetric approximation arises as the direct perturbation of an already-available symmetric one, without further optimization, it is useful to place particular emphasis on those Gaussian approximations provided by state-of-the-art softwares so as to stimulate direct adoption of the proposed perturbation. Note also that standard symmetric approximations often depend on additional parameters beyond the symmetry point $\boldsymbol{\theta}^*$. For instance, Gaussian approximations also require estimation of the covariance matrix. As clarified in Section 2.3, the skewing factor we derive only depends on $\boldsymbol{\theta}^*$ and is applied to an already-available symmetric approximation whose parameters have been previously estimated under standard algorithms. Hence, to ease notation, we index $\bar{q}_{n,\boldsymbol{\theta}^*}(\boldsymbol{\theta})$ only by $\boldsymbol{\theta}^*$, and leave the remaining, already-estimated, parameters implicit. Theoretical guarantees for the finite sample and asymptotic accuracy of the proposed approximation are derived in Section 3.

2.1 An overview of symmetric approximations of posterior distributions

When seeking a tractable approximation of a generic posterior distribution for the parameter $\boldsymbol{\theta}$, a simple option is to consider the classical Gaussian approximation that arises from the Laplace method (see, e.g. [Bishop, 2006, Chapter 4.4](#)). Such a solution follows directly from a second-order Taylor expansion for the un-normalized log-posterior $\log[\pi(\boldsymbol{\theta})L(\boldsymbol{\theta}; \mathbf{y}_{1:n})]$ at the maximum a posteriori (MAP). This yields a Gaussian approximating density, which is centered at the MAP and has

covariance matrix given by the inverse of the negative Hessian for $\log[\pi(\theta)L(\theta; y_{1:n})]$, again evaluated at the MAP (e.g. [Gelman et al., 2013](#), Chapter 13). The simplicity of this solution has stimulated broad applicability and several extensions, including, among others, integrated nested Laplace approximation (INLA) ([Rue et al., 2009](#)) and approximate Laplace approximation (ALA) ([Rossell et al., 2021](#)). INLA combines efficient numerical integration and analytical approximations to derive accurate characterizations of posterior marginals for parameters of interest, under latent Gaussian models. The scheme is inspired by [Tierney and Kadane \(1986\)](#) Laplace approximation of the marginal posterior density for a subset of parameters of interest, which expresses this marginal as proportional to the ratio between the joint posterior and the full conditional density of the remaining parameters, and then applies the Gaussian approximation from the Laplace method to the latter density. This strategy can yield asymmetric approximations of marginal posterior densities, but still relies on nested Gaussian approximations. As such, the skew-symmetric solution derived in Section 2.3 has the potential to further improve also the accuracy of INLA. Our proposal can be also applied directly to improve ALA. This scheme does not differ from the classical Laplace method in the shape of the approximating density. Rather, it provides a scalable strategy which avoids MAP estimation.

While Laplace type schemes yield simple approximation strategies, the resulting solution may fail to incorporate global characteristics of the posterior beyond the local behaviour at the MAP ([Bishop, 2006](#), Chapter 4.4). This issue has motivated interest in alternative schemes, with a main focus on VB (e.g. [Blei et al., 2017](#)) and EP (e.g. [Vehtari et al., 2020](#)). VB specifies a tractable family \mathcal{F} of approximating densities and then identifies, within this family, the one that is closest to the intractable posterior in KL divergence ([Kullback & Leibler, 1951](#)). The common practice in specifying \mathcal{F} relies either on parametric, often Gaussian, families (see, e.g. [Challis & Barber, 2013](#); [Kucukelbir et al., 2017](#); [Opper & Archambeau, 2009](#); [Tan & Nott, 2018](#)), or on mean-field assumptions (e.g. [Blei et al., 2017](#)) enforcing the joint approximating density for θ to factorize as the product of marginals for suitably selected non-overlapping subsets of parameters. In the first case the final output of the optimization problem is, by definition, a symmetric density which can be readily improved under our proposed skew-symmetric approximation. Conversely, the second can yield skewed solutions matching the shape of the actual, not necessarily symmetric, full-conditional density for the subset of parameters in θ comprising the mean-field factors (e.g. [Blei et al., 2017](#)). Nonetheless, several routine use Bayesian models, such as logit ([Durante & Rigon, 2019](#)) and probit ([Consonni & Marin, 2007](#)) regression, admit Gaussian full-conditionals for the coefficients. Therefore, the proposed perturbation might be useful even within mean-field VB to improve the accuracy of the symmetric density factors, and hence, of the overall approximation for the entire posterior. Finally, it is important to emphasize that these Gaussian approximations appear either as the final output or as a building block also within Laplace variational inference, delta-method variational inference ([C. Wang & Blei, 2013](#)) and stochastic black-box extensions of VB ([Ranganath et al., 2014](#)), including automatic differentiation variational inference (ADVI) ([Kucukelbir et al., 2017](#)) and solutions relying on normalizing flows (e.g. [Papamakarios et al., 2021](#); [Rezende & Mohamed, 2015](#)). Such methods are crucial to facilitate the derivation of generally applicable and more flexible VB approximations that can be automatically implemented even in settings lacking closed-form expressions for the quantities involved in the optimization problem. Therefore, the proposed skew-symmetric perturbation can be either applied directly to improve the approximation accuracy of the Gaussian output provided by these methods (see Sections 4–5), or may be even included within the stochastic optimization scheme to enlarge the family (and hence the quality) of the approximating densities, without arguably affecting the tractability of the routine. This broad applicability of the simple skew-symmetric perturbation we propose is a unique feature, which generally lacks in more sophisticated solutions yielding accurate, possibly asymmetric, approximations. Successful ones are those provided by the VB strategies leveraging normalizing flows (e.g. [Papamakarios et al., 2021](#); [Rezende & Mohamed, 2015](#)). Interestingly, as discussed in Section 6, such strategies often use suitable Gaussians as initial densities, and hence, our proposal could be possibly leveraged to provide improved (i.e. closer to the target posterior), yet tractable, skew-symmetric perturbations of these initial Gaussian densities.

Although VB is arguably the most widely studied and implemented deterministic approximation strategy, commonly used variational approximations often suffer from an underestimation of posterior uncertainty (see, e.g. [Blei et al., 2017](#); [Giordano et al., 2018](#)). This is implicit in the expression of the KL minimized under VB that penalizes densities placing mass to areas of low posterior probability, but it does not enforce a similarly strong penalty in the opposite direction. As a

consequence, minimizing such a KL yields more concentrated approximations avoiding those regions where the target posterior does not place substantial mass. While improved estimates of variances and covariances have been proposed in the context of VB (Giordano et al., 2018), another natural solution is to consider EP (e.g. Vehtari et al., 2020), which addresses such an issue by minimizing a reverse form of the KL considered under VB. This implies penalizations in the opposite direction than those of VB, and hence, a tendency to favour more global *mass-covering* approximations that match more closely the variability encoded in the target posterior. To obtain these approximations, EP postulates that the target posterior density itself can be expressed as a product of factors, and then iteratively approximates each of these factors with an element of a tractable parametric family, almost always Gaussian. This results in a computational scheme updating each factor at-a-time via moment matching between the global approximating density and a hybrid one, more tractable than the target posterior, where the other factors are kept fixed at the most recent approximation (e.g. Vehtari et al., 2020). Being often Gaussian, these EP approximating densities can be readily perturbed under our proposed strategy to further improve the quality of the approximation. In fact, although several empirical studies have highlighted the remarkable accuracy of Gaussian EP (Anceschi et al., 2023 ; Chopin & Ridgway, 2017 ; Vehtari et al., 2020) relative to Laplace and VB, there is still a lack of tractable solutions which are capable of including skewness within these Gaussian approximations to further improve the quality of EP. As illustrated in the empirical studies in Sections 4 and 5, accounting for these skewed behaviours yields further improvements over such an already-successful strategy.

It is important to emphasize that our strategy applies directly to any symmetric approximation, not necessarily Gaussian. This can be useful to perturb further extensions of the aforementioned methods aimed at capturing higher-order properties (e.g. tails). For example, generalizations from Gaussian approximating densities to Student- t ones have been explored in the context of Laplace, VB and EP (e.g. Daudel et al., 2023; Ding et al., 2011; Futami et al., 2017; Gelman et al., 2013; Liang et al., 2022), but there are no simple strategies, to date, for including skewness within such extensions. Our proposed solution applies also to these symmetric approximations. As proved in Section 3.2, the perturbation of these increasingly accurate symmetric densities opens also the avenues to further improve the asymptotic convergence rates to the target posterior.

2.2 An overview of skew-symmetric distributions

As anticipated in Section 1, the novel strategy we propose to perturb an already-available symmetric approximation of a generic target posterior leads to an improved asymmetric counterpart, which crucially belongs to the skew-symmetric family (e.g. Azzalini & Capitanio, 2003; Genton & Loperfido, 2005; Ma & Genton, 2004; J. Wang et al., 2004). This broad family has been originally proposed to increase the flexibility of symmetric densities via the inclusion of a suitably designed skewing factor that preserves the tractability of the unperturbed counterparts. Definition 1 (e.g. Azzalini & Capitanio, 2013; J. Wang et al., 2004), clarifies that this objective can be accomplished through a direct perturbation scheme that meets simple conditions.

Definition 1 (Skew-symmetric densities). Let $\bar{q}_{\theta^*}(\theta)$ be a density for $\theta \in \Theta \subseteq \mathbb{R}^d$, which is symmetric at θ^* , i.e. $\bar{q}_{\theta^*}(\theta) = \bar{q}_{\theta^*}(2\theta^* - \theta)$ for all $\theta \in \Theta$, and denote with $w_{\theta^*}(\theta) \in [0, 1]$ a skewing factor such that $w_{\theta^*}(\theta) = 1 - w_{\theta^*}(2\theta^* - \theta)$ for any $\theta \in \Theta$. Then

$$q_{\theta^*}(\theta) = 2\bar{q}_{\theta^*}(\theta)w_{\theta^*}(\theta), \quad (1)$$

is a skew-symmetric density.

Definition 1 provides a general skewness-inducing mechanism that leverages a suitably designed perturbation factor $w_{\theta^*}(\theta)$ to redistribute the density between each point θ and its polar opposite $2\theta^* - \theta$, with respect to a notion of central symmetry. While general, such a definition crucially encompasses several special cases of routine use in practice. A remarkable example is provided by the sub-class of generalized skew-elliptical distributions which can be obtained by replacing $\bar{q}_{\theta^*}(\theta)$ in (1) with an elliptical density, including, for instance, those of multivariate Gaussians, Student- t , and Cauchy, among others (e.g. Genton & Loperfido, 2005). Perturbing these latter

densities through a specific skewing factor $w_{\theta^*}(\theta)$ set equal to a suitably defined cumulative distribution function induced by the family of densities $\bar{q}_{\theta^*}(\theta)$ further yields classical skew-normal, skew- t and skew-Cauchy densities, among others (Azzalini & Capitanio, 1999, 2003; Azzalini & Dalla Valle, 1996; Genton & Loperfido, 2005).

Albeit different, all the above alternatives share a common property, which characterizes the general skew-symmetric family in Definition 1. In particular, as is clear from (1), the tractability of $q_{\theta^*}(\theta)$ inherently depends on the one of its symmetric counterpart $\bar{q}_{\theta^*}(\theta)$ and on the choice of the skewing factor $w_{\theta^*}(\theta)$. When both quantities are tractable, the density in (1) can be computed analytically, and further admits a straightforward i.i.d. sampling scheme that facilitates direct Monte Carlo evaluation of any functional of $q_{\theta^*}(\theta)$. This follows from the stochastic representation of skew-symmetric distributions in Proposition 1 (e.g. J. Wang et al., 2004).

Proposition 1 (Stochastic representation). Let θ be a random variable having density $q_{\theta^*}(\theta)$ as in (1), and denote with $\mathbb{1}[\cdot]$ the indicator function, then

$$\theta \stackrel{d}{=} \mathbb{1}[U \leq w_{\theta^*}(\bar{\theta})]\bar{\theta} + (1 - \mathbb{1}[U \leq w_{\theta^*}(\bar{\theta})])(2\theta^* - \bar{\theta}),$$

where $\bar{\theta}$ has symmetric density $\bar{q}_{\theta^*}(\bar{\theta})$, and U is a uniform in $[0, 1]$.

As a consequence of Proposition 1, i.i.d. sampling from skew-symmetric distributions reduces to drawing values from the symmetric component and then retaining the sampled value or its symmetric with respect to θ^* depending on whether the skewing factor evaluated at such a sampled value is below or above the realization from a uniform distribution in $[0, 1]$. This strategy is crucial to facilitate tractable Monte Carlo inference under the proposed skew-symmetric approximation $q_{n,\theta^*}(\theta)$ of posterior densities $\pi_n(\theta)$. As clarified in Section 2.3, this approximation arises from the perturbation of tractable symmetric counterparts $\bar{q}_{n,\theta^*}(\theta)$ (often Gaussians) via a skewing factor $w_{n,\theta^*}(\theta)$ that involves only prior $\pi(\theta)$ and likelihood $L(\theta; \mathbf{y}_{1:n})$ evaluation.

2.3 Skew-symmetric perturbation of symmetric approximations

Recalling Section 1, the class of skewed approximations we propose is motivated by a fundamental, but yet overlooked, skew-symmetric representation of posterior densities stated in Proposition 2.

Proposition 2 (Skew-symmetric representation of posterior densities). Consider the generic posterior density $\pi_n(\theta) = \pi(\theta)L(\theta; \mathbf{y}_{1:n})/c(\mathbf{y}_{1:n})$ for the parameter $\theta \in \Theta$, where $\pi(\theta)$ denotes the prior, $L(\theta; \mathbf{y}_{1:n})$ corresponds to the likelihood, while $c(\mathbf{y}_{1:n})$ is the normalizing constant. Moreover, denote with

$$\bar{\pi}_{n,\theta^*}(\theta) = \frac{\pi_n(\theta) + \pi_n(2\theta^* - \theta)}{2}, \tag{2}$$

the symmetrized form of this posterior density about a point $\theta^* \in \Theta$, and let

$$w_{n,\theta^*}(\theta) = \frac{\pi(\theta)L(\theta; \mathbf{y}_{1:n})}{\pi(\theta)L(\theta; \mathbf{y}_{1:n}) + \pi(2\theta^* - \theta)L(2\theta^* - \theta; \mathbf{y}_{1:n})}, \tag{3}$$

with the convention that $w_{n,\theta^*}(\theta) = 0.5$ if both $\pi(\theta)L(\theta; \mathbf{y}_{1:n}) = 0$ and $\pi(2\theta^* - \theta)L(2\theta^* - \theta; \mathbf{y}_{1:n}) = 0$ (see Remark 3). Then, the posterior density $\pi_n(\theta)$ can be equivalently re-expressed in skew-symmetric form as

$$\pi_n(\theta) = 2\bar{\pi}_{n,\theta^*}(\theta)w_{n,\theta^*}(\theta), \tag{4}$$

for any θ^* and n , with $w_{n,\theta^*}(\theta) \in [0, 1]$ and $w_{n,\theta^*}(\theta) = 1 - w_{n,\theta^*}(2\theta^* - \theta)$.

Remark 1 While other forms of symmetrization can be considered, the one defined in (2) is arguably the most natural, direct and simple, in that it redistributes equally to

every pair of symmetric points the total posterior density assigned to this pair. Although it has been generally overlooked and, to our knowledge, never explored in the context of Bayesian approximations, such a symmetrization has been successfully employed in classical frequentist literature to improve standard estimators of empirical distribution functions associated with underlying symmetric densities (e.g. [Hinkley, 1976](#); [Lo, 1985](#); [Schuster, 1975, 1987](#)). Our contribution leverages such a symmetrization with a substantially different focus. Namely that of designing principled and tractable skew-symmetric approximations of generic posterior densities from the direct perturbation of already-derived symmetric ones.

The proof of Proposition 2 simply requires to notice that $\pi_n(\theta)$ can be equivalently re-written as $2\bar{\pi}_{n,\theta^*}(\theta)[\pi_n(\theta)/(2\bar{\pi}_{n,\theta^*}(\theta))]$. Therefore, replacing in the ratio $\pi_n(\theta)/(2\bar{\pi}_{n,\theta^*}(\theta))$ the quantity $\bar{\pi}_{n,\theta^*}(\theta)$ with its expression in (2), and $\pi_n(\theta)$ with $\pi(\theta)L(\theta; \mathbf{y}_{1:n})/c(\mathbf{y}_{1:n})$, yields the skewness-inducing factor $w_{n,\theta^*}(\theta)$ in (3), after noticing that the intractable normalizing constant $c(\mathbf{y}_{1:n})$ cancels out within the ratio between the target posterior density and its symmetrized form. The $[0, 1]$ range for $w_{n,\theta^*}(\theta)$ and the property $w_{n,\theta^*}(\theta) = 1 - w_{n,\theta^*}(2\theta^* - \theta)$ follow from its definition in (3). As clarified later, such properties for $w_{n,\theta^*}(\theta)$ are fundamental to ensure that the proposed perturbation for $\bar{q}_{n,\theta^*}(\theta)$ belongs to the skew-symmetric family (e.g. [Azzalini & Capitanio, 2003](#); [Ma & Genton, 2004](#); [J. Wang et al., 2004](#)).

Proposition 2 relates to a core result in [J. Wang et al. \(2004\)](#), which establishes the existence and uniqueness of skew-symmetric representations for generic densities. This parallel ensures that the equivalent expression for $\pi_n(\theta)$ in Proposition 2 is the one of a skew-symmetric density. According to (4), such a density is equal to the product between a tractable and analytically available skewing factor $w_{n,\theta^*}(\theta)$ (where the troublesome normalizing constant $c(\mathbf{y}_{1:n})$ cancels out) and a symmetrized posterior $\bar{\pi}_{n,\theta^*}(\theta)$ that is often as intractable as the target $\pi_n(\theta)$. Such a latter issue prevents (4) from having a direct impact in facilitating posterior inference and, for this reason, the representation in Proposition 2 has never appeared within Bayesian statistics. One possibility to address the above issue, would be to replace $\bar{\pi}_{n,\theta^*}(\theta)$ in (4) with a tractable symmetric approximation. Lemma 1 suggests that a natural candidate is the original, already-derived, symmetric approximation $\bar{q}_{n,\theta^*}(\theta)$ of the target posterior density $\pi_n(\theta)$, which we aim to improve through a skewed perturbation. In particular, Lemma 1 proves that $\bar{q}_{n,\theta^*}(\theta)$ approximates $\bar{\pi}_{n,\theta^*}(\theta)$ more accurately than $\pi_n(\theta)$, under the TV distance $\mathcal{D}_{\text{TV}}[p \parallel q] = (1/2)\int |p(\theta) - q(\theta)|d\theta$ and any generic α -divergence $\mathcal{D}_\alpha[p \parallel q] = [1/(\alpha(1-\alpha))](1 - \int p(\theta)^\alpha q(\theta)^{1-\alpha} d\theta)$ for $\alpha \in \mathbb{R} \setminus \{0, 1\}$ (see, e.g. [Cichocki & Amari, 2010](#); [Margossian et al., 2025](#); [Póczos & Schneider, 2011](#)), including also its limiting KL ($\text{KL}[q \parallel p] = \int q(\theta) \log [q(\theta)/p(\theta)] d\theta$) and reverse-KL ($\text{KL}[p \parallel q] = \int p(\theta) \log [p(\theta)/q(\theta)] d\theta$) forms, which can be obtained by letting $\alpha \rightarrow 0$ and $\alpha \rightarrow 1$, respectively, in \mathcal{D}_α (e.g. [Cichocki & Amari, 2010](#)). Hence, Lemma 1, and the subsequent theoretical results we derive, hold under the most widely used divergences in the context of deterministic approximations. Recalling Section 2.1, the divergences $\text{KL}[q \parallel p]$ and $\text{KL}[p \parallel q]$ enter the optimization problems in VB and EP (e.g. [Blei et al., 2017](#); [Vehtari et al., 2020](#)), while TV is a reference distance in the study of the asymptotic accuracy of standard approximations. Recent literature has also explored variational approximations based on generic α -divergences ([Hernandez-Lobato et al., 2016](#); [Yang et al., 2020](#)). Hence, our methods and theory apply also to the symmetric approximations from these strategies.

Lemma 1 Let $\pi_n(\theta)$ be a generic posterior density for the parameter $\theta \in \Theta$, and denote with $\bar{q}_{n,\theta^*}(\theta)$ an already-derived approximation of $\pi_n(\theta)$ which is symmetric about the point $\theta^* \in \Theta$. Moreover, consider the symmetrized posterior density $\bar{\pi}_{n,\theta^*}(\theta)$ about θ^* defined as in (2). Then, for any $\theta^* \in \Theta$ and n , it holds

$$\mathcal{D}[\bar{\pi}_{n,\theta^*} \parallel \bar{q}_{n,\theta^*}] \leq \mathcal{D}[\pi_n \parallel \bar{q}_{n,\theta^*}],$$

where \mathcal{D} is either the TV distance, KL, reverse-KL or a generic α -divergence.

As clarified in Definition 2, combining Proposition 2 with Lemma 1 suggests a natural strategy to obtain the newly proposed skew-symmetric approximation. In particular, by Lemma 1, $\bar{q}_{n,\theta^*}(\theta)$ approximates $\bar{\pi}_{n,\theta^*}(\theta)$ more accurately than $\pi_n(\theta)$, where $\bar{\pi}_{n,\theta^*}(\theta)$ can be in turn expressed via $\bar{\pi}_{n,\theta^*}(\theta) = \pi_n(\theta) / [2w_{n,\theta^*}(\theta)]$ as a direct consequence of (4). Combining these two results yields $\bar{q}_{n,\theta^*}(\theta) \approx \bar{\pi}_{n,\theta^*}(\theta) = \pi_n(\theta) / [2w_{n,\theta^*}(\theta)]$, which implies $2\bar{q}_{n,\theta^*}(\theta)w_{n,\theta^*}(\theta) \approx \pi_n(\theta)$. Hence, setting $q_{n,\theta^*}(\theta) = 2\bar{q}_{n,\theta^*}(\theta)w_{n,\theta^*}(\theta)$ gives the improved skew-symmetric approximation in Definition 2.

Definition 2 (Skew-symmetric approximation of posterior densities). Consider the generic posterior density $\pi_n(\theta)$ for $\theta \in \Theta$, and let $\bar{q}_{n,\theta^*}(\theta)$ denote an already-derived approximation for $\pi_n(\theta)$, which is symmetric about the point $\theta^* \in \Theta$. Moreover, let $w_{n,\theta^*}(\theta) \in [0, 1]$ be the skewing factor in (3). Then, the skew-symmetric approximation of $\pi_n(\theta)$ arising from the perturbation of $\bar{q}_{n,\theta^*}(\theta)$ is defined as

$$\begin{aligned} q_{n,\theta^*}(\theta) &= 2\bar{q}_{n,\theta^*}(\theta)w_{n,\theta^*}(\theta) \\ &= 2\bar{q}_{n,\theta^*}(\theta) \frac{\pi(\theta)L(\theta; \mathbf{y}_{1:n})}{\pi(\theta)L(\theta; \mathbf{y}_{1:n}) + \pi(2\theta^* - \theta)L(2\theta^* - \theta; \mathbf{y}_{1:n})}, \end{aligned} \tag{5}$$

for every symmetry point $\theta^* \in \Theta$ and sample size n , where $\pi(\theta)$ and $L(\theta; \mathbf{y}_{1:n})$ are, respectively, the prior and likelihood inducing the posterior $\pi_n(\theta)$.

Remark 2 Notice that $w_{n,\theta^*}(\theta)$ in (5) admits a natural interpretation, in that it coincides with the relative proportion of the posterior density at $\theta \in \Theta$ with respect to the total assigned to such a θ and its symmetric counterpart $2\theta^* - \theta$. This yields a skewing factor which quantifies differences in the posterior density at the symmetric points $\{\theta, 2\theta^* - \theta\}$ for any $\theta \in \Theta$. Therefore, if the posterior density is actually symmetric about θ^* , then $w_{n,\theta^*}(\theta) = 0.5$ for all $\theta \in \Theta$, and hence, $q_{n,\theta^*}(\theta)$ reduces to $\bar{q}_{n,\theta^*}(\theta)$, as expected. Conversely, whenever there are asymmetries within $\pi_n(\theta)$, the original symmetric approximation $\bar{q}_{n,\theta^*}(\theta)$ is re-weighted by $w_{n,\theta^*}(\theta)$ in order to properly redistribute the total density at each pair $\{\theta, 2\theta^* - \theta\}$ according to the one assigned by the actual posterior to θ and $2\theta^* - \theta$. This yields an improved approximation $q_{n,\theta^*}(\theta)$ incorporating the skewness of $\pi_n(\theta)$ with respect to the known symmetry point θ^* .

Remark 3 In the skew-symmetric representation in Proposition 2, we rely on the convention that $w_{n,\theta^*}(\theta) = 0.5$ when $\pi(\theta)L(\theta; \mathbf{y}_{1:n}) = \pi(2\theta^* - \theta)L(2\theta^* - \theta; \mathbf{y}_{1:n}) = 0$, where the latter equality also implies that $\pi_n(\theta) = \pi_n(2\theta^* - \theta) = 0$. In this setting, which arises, for example, when the support of θ is bounded, both the numerator and the denominator of the skewing factor in (3) are zero, and hence, $w_{n,\theta^*}(\theta)$ is undefined, thus requiring an alternative specification. Note that this alternative specification is not necessary to guarantee the validity of the skew-symmetric representation for $\pi_n(\theta)$, provided that also the symmetrized posterior $\bar{\pi}_{n,\theta^*}(\theta)$ is zero whenever $\pi_n(\theta) = \pi_n(2\theta^* - \theta) = 0$. Hence, in (4) different values of $w_{n,\theta^*}(\theta)$ are irrelevant at these points with a zero density. However, $\pi_n(\theta) = \pi_n(2\theta^* - \theta) = 0$ does not imply that $\bar{q}_{n,\theta^*}(\theta) = 0$. This situation arises, for example, when Gaussian approximations are considered for parameters θ having bounded support. Therefore, when $\pi_n(\theta) = \pi_n(2\theta^* - \theta) = 0$ and $\bar{q}_{n,\theta^*}(\theta) > 0$, the skewing factor must be specified alternatively to guarantee that (5) yields a valid skew-symmetric density. From Definition 1, a central condition for this result to hold is that $w_{n,\theta^*}(\theta) = 1 - w_{n,\theta^*}(2\theta^* - \theta)$ with $w_{n,\theta^*}(\theta) \in [0, 1]$, for any point θ satisfying $\bar{q}_{n,\theta^*}(\theta) > 0$. To this end, setting $w_{n,\theta^*}(\theta) = 0.5$ whenever $\pi_n(\theta) = \pi_n(2\theta^* - \theta) = 0$ and $\bar{q}_{n,\theta^*}(\theta) > 0$ is arguably

the most direct strategy to meet such a key condition. Recalling Remark 2, $w_{n,\theta^*}(\theta) = 0.5$ implies that $\pi_n(\theta) = \pi_n(2\theta^* - \theta)$, and hence, the convention that $w_{n,\theta^*}(\theta) = 0.5$ if $\pi_n(\theta) = \pi_n(2\theta^* - \theta) = 0$ is also the most natural one.

As clarified in Definition 2, the proposed approximation $q_{n,\theta^*}(\theta)$ results from the re-weighting of the known $\bar{q}_{n,\theta^*}(\theta)$ by a skewing factor $w_{n,\theta^*}(\theta)$ via a strategy which does not necessitate additional optimization costs relative to those required for deriving $\bar{q}_{n,\theta^*}(\theta)$. In fact, the expression for $w_{n,\theta^*}(\theta)$ in (3) does not depend on additional unknown parameters beyond θ^* , which is in turn available as the output of the already-solved optimization problem that targeted the posterior $\pi_n(\theta)$ via the symmetric density $\bar{q}_{n,\theta^*}(\theta)$ to be perturbed. Proposition 3 below guarantees that, albeit more flexible than $\bar{q}_{n,\theta^*}(\theta)$, the deterministic approximation $q_{n,\theta^*}(\theta)$ in Definition 2 preserves similar tractability in inference, in that it belongs to the known class of skew-symmetric densities presented in Section 2.2.

Proposition 3 The expression for $q_{n,\theta^*}(\theta)$ given in (5) coincides with the density of a skew-symmetric distribution having symmetric component $\bar{q}_{n,\theta^*}(\theta)$ and skewing factor $w_{n,\theta^*}(\theta)$.

Proposition 3 follows directly from Definition 1, after noticing that, by construction, $\bar{q}_{n,\theta^*}(\theta)$ is symmetric at θ^* and, in view of Proposition 2, $w_{n,\theta^*}(\theta)$ has support within $[0, 1]$ and satisfies $w_{n,\theta^*}(\theta) = 1 - w_{n,\theta^*}(2\theta^* - \theta)$. As anticipated in Section 2.2, the connection with skew-symmetric distributions established in Proposition 3 is crucial in facilitating inference also under $q_{n,\theta^*}(\theta)$. More concretely, the stochastic representation in Proposition 1 yields a simple rejection-free i.i.d. sampling scheme from any density within the skew-symmetric family, including $q_{n,\theta^*}(\theta)$, thereby allowing tractable and effective Monte Carlo evaluation of any functional of interest under the improved skewed approximation. This sampling scheme from $q_{n,\theta^*}(\theta)$ is outlined in Algorithm 1.

Notice that Algorithm 1 only requires simulation from the symmetric approximation $\bar{q}_{n,\theta^*}(\theta)$ and computation of the skewing factor $w_{n,\theta^*}(\theta)$, which is analytically-available in Definition 2, and it does not depend on intractable quantities. The first task is straightforward whenever the unperturbed density $\bar{q}_{n,\theta^*}(\theta)$ arises from one of the routinely implemented approximation schemes discussed in Section 2.1. Moreover, when the interest lies in more complex functionals, sampling is often needed also for inference under $\bar{q}_{n,\theta^*}(\theta)$. The second task requires instead the evaluation of the likelihood, which enters (together with the prior) the definition of $w_{n,\theta^*}(\theta)$. Although this yields some increments in sampling costs relative to inference under $\bar{q}_{n,\theta^*}(\theta)$, it is important to emphasize that multiple internal likelihood evaluations are standard in popular algorithms for Bayesian inference. Important examples include Metropolis–Hastings, Hamiltonian Monte Carlo, sequential Monte Carlo, EP, coordinate ascent variational inference, and its black-box extensions (e.g. Blei et al., 2017; Chopin & Papaspiliopoulos, 2020; Chopin & Ridgway, 2017; Kucukelbir et al., 2017; Ranganath et al., 2014; Tan & Chen, 2024; Vehtari et al., 2020). However, unlike for Algorithm 1, these schemes introduce additional complexities, which often imply further tuning and non-negligible increments in the overall runtimes. Hence, the cost of inference under $q_{n,\theta^*}(\theta)$ is often dominated by the one required to obtain its less accurate symmetric counterpart $\bar{q}_{n,\theta^*}(\theta)$ via state-of-the-art algorithms, and it substantially reduces the runtimes of other routine use computational methods. These important gains are illustrated in Section 5.2 for a challenging

Algorithm 1 i.i.d. sampling from the skew-symmetric approximation in (5)

```

for  $s = 1, \dots, N_{\text{SAMPLE}}$  do
  1. Sample  $\bar{\theta}^{(s)}$  from the distribution with symmetric density  $\bar{q}_{n,\theta^*}(\bar{\theta})$ .
  2. Sample  $u^{(s)} \sim \text{Unif}[0, 1]$ .
  3. If  $u^{(s)} \leq w_{n,\theta^*}(\bar{\theta}^{(s)})$  set  $\theta^{(s)} = \bar{\theta}^{(s)}$ , otherwise set  $\theta^{(s)} = 2\theta^* - \bar{\theta}^{(s)}$ .
end for
output: i.i.d. samples  $\theta^{(1)}, \dots, \theta^{(N_{\text{SAMPLE}})}$  from the skew-symmetric approximation in (5).

```

real-data application with $n = 30,524$ and $d = 62$. In this case, sampling 10,000 values from the skew-symmetric approximation under Algorithm 1 requires only 8 seconds. This runtime is orders of magnitude lower than, for example, the ≈ 1 hour required to obtain a sample of the same size from Hamiltonian Monte Carlo implemented in `rstan`.

Note also that, unlike for most of the aforementioned computational methods, Algorithm 1 is inherently amenable to parallel implementations both across samples and likelihood factors. Thus, the runtime of Algorithm 1 can be easily reduced to fractions of seconds. Moreover, as clarified within Section 2.4 below, in several routinely implemented models such as, e.g. generalized linear models, it is possible to devise computationally efficient strategies that further optimize the cost of the likelihood evaluations required to compute the skewing factor.

2.4 Efficient evaluation of the skewing factor

Recalling the above discussion, inference under the newly proposed skew-symmetric approximation $q_{n,\theta^*}(\theta)$ requires the evaluation of the two un-normalized posterior densities $\pi(\theta)L(\theta; \mathbf{y}_{1:n})$ and $\pi(2\theta^* - \theta)L(2\theta^* - \theta; \mathbf{y}_{1:n})$ in the skewing factor $w_{n,\theta^*}(\theta)$. Hence, the increments in cost relative to performing inference under $\bar{q}_{n,\theta^*}(\theta)$ mainly depend on those of evaluating the likelihoods $L(\theta; \mathbf{y}_{1:n})$ and $L(2\theta^* - \theta; \mathbf{y}_{1:n})$. As mentioned above, multiple likelihoods evaluations are standard in Bayesian computation and can be performed with negligible cost for most models of practical interest. Nevertheless, designing strategies that further optimize this cost could lead to additional reductions in the overall runtimes. This is the case, for example, in high-dimensional (i.e. large d) models, with the parameters entering the likelihood function through a linear predictor, as in routinely implemented generalized linear models.

Under these models, Algorithm 2 provides an efficient approach for computing $w_{n,\theta^*}(\theta)$ at a cost which is essentially that of a single likelihood evaluation instead of the two required in the expression of $w_{n,\theta^*}(\theta)$. More specifically, Algorithm 2 can be applied to the broad class of models whose log-likelihood $\ell(\theta; \mathbf{y}_{1:n}) = \log L(\theta; \mathbf{y}_{1:n})$ can be expressed as $\ell(\theta; \mathbf{y}_{1:n}) = \sum_{i=1}^n g_i(\mathbf{y}_i, \eta_{\theta,i})$, where $g_i(\cdot)$, $i = 1, \dots, n$, denote functions that can be evaluated with a limited computational effort, while $\eta_{\theta,i} = \mathbf{x}_i^T \theta$ for $i = 1, \dots, n$, are linear predictors depending on θ and on a d -dimensional vector of explanatory variables \mathbf{x}_i . Within this context, for (at least moderately) large d , the most expensive operation is the evaluation of the linear predictor. Algorithm 2 reduces the cost required for computing $w_{n,\theta^*}(\theta)$ by exploiting the fact that

$$\ell(\theta; \mathbf{y}_{1:n}) = \sum_{i=1}^n g_i(\mathbf{y}_i, \eta_{\theta^*,i} + \eta_{\theta - \theta^*,i}), \quad \text{and} \quad \ell(2\theta^* - \theta; \mathbf{y}_{1:n}) = \sum_{i=1}^n g_i(\mathbf{y}_i, \eta_{\theta^*,i} - \eta_{\theta - \theta^*,i}),$$

where $\eta_{\theta^*,i} = \mathbf{x}_i^T \theta^*$ and $\eta_{\theta - \theta^*,i} = \mathbf{x}_i^T (\theta - \theta^*)$, $i = 1, \dots, n$. Crucially, $\boldsymbol{\eta}_{\theta^*} = (\eta_{\theta^*,1}, \dots, \eta_{\theta^*,n})^T$ can be pre-computed, meaning that only a single evaluation of $\boldsymbol{\eta}_{\theta - \theta^*} = (\eta_{\theta - \theta^*,1}, \dots, \eta_{\theta - \theta^*,n})^T$ is required in order to compute the skewing factor. As a result, since the cost of $\mathbf{x}_i^T \theta$ and $g_i(\cdot)$ is $O(d)$ and

Algorithm 2 Efficient evaluation of $w_{n,\theta^*}(\theta)$ in models with linear predictors

require: $\boldsymbol{\eta}_{\theta^*}$, $\theta - \theta^*$ and the design matrix \mathbf{X} with rows \mathbf{x}_i^T .
do
1. Evaluate $\boldsymbol{\eta}_{\theta - \theta^*,i} = \mathbf{x}_i^T (\theta - \theta^*)$, for $i = 1, \dots, n$.
2. Compute $\ell(\theta; \mathbf{y}_{1:n}) = \sum_{i=1}^n g_i(\mathbf{y}_i, \eta_{\theta^*,i} + \eta_{\theta - \theta^*,i})$.
3. Compute $\ell(2\theta^* - \theta; \mathbf{y}_{1:n}) = \sum_{i=1}^n g_i(\mathbf{y}_i, \eta_{\theta^*,i} - \eta_{\theta - \theta^*,i})$.

output:

$$w_{n,\theta^*}(\theta) = \frac{\pi(\theta) \exp(\ell(\theta; \mathbf{y}_{1:n}))}{\pi(\theta) \exp(\ell(\theta; \mathbf{y}_{1:n})) + \pi(2\theta^* - \theta) \exp(\ell(2\theta^* - \theta; \mathbf{y}_{1:n}))}.$$

$O(1)$, respectively, [Algorithm 2](#) allows to evaluate $w_{n,\theta^*}(\theta)$ in a number of operations which, for d large, is essentially reduced by a factor of two. This yields a sampling strategy whose computational complexity interestingly matches, or even improves, those of other available skewed approximations (see, e.g. [Anceschi et al., 2023](#); [Durante et al., 2024](#); [Fasano et al., 2022](#)). Notice that, unlike for the general perturbation strategy we propose, these alternatives are developed in the context of specific deterministic approximations (e.g. Laplace and VB) and/or models (e.g. probit regression and its extensions). Hence, matching the costs of these ad-hoc implementations is arguably a remarkable result.

3 Theoretical properties of skew-symmetric approximations

Sections 3.1–3.2 clarify that the skew-symmetric approximation $q_{n,\theta^*}(\theta)$ derived in Section 2 (see also Definition 2) is not only computationally tractable, but also yields a provably more accurate characterization of the exact posterior $\pi_n(\theta)$. See the [online supplementary material](#) for proofs.

3.1 Finite sample properties and optimality

The original motivation behind the skew-symmetric approximation $q_{n,\theta^*}(\theta)$ in (5) is to improve the accuracy of the unperturbed $\bar{q}_{n,\theta^*}(\theta)$. Theorem 1 provides finite sample theoretical support to such an accuracy gain and clarifies that the quality of $q_{n,\theta^*}(\theta)$ only depends on how accurate is $\bar{q}_{n,\theta^*}(\theta)$ in approximating the symmetrized posterior $\bar{\pi}_{n,\theta^*}(\theta)$ in (2). These results are deepened in Theorem 2, which proves that the skewing factor $w_{n,\theta^*}(\theta)$ in Definition 2 is optimal among all those yielding a skew-symmetric approximation for $\pi_n(\theta)$, with $\bar{q}_{n,\theta^*}(\theta)$ as symmetric component.

Theorem 1 (Finite sample accuracy). Consider the generic posterior density $\pi_n(\theta)$ for the parameter $\theta \in \Theta$, and let $\bar{q}_{n,\theta^*}(\theta)$ correspond to an already-derived approximation for $\pi_n(\theta)$ which is symmetric about the point $\theta^* \in \Theta$. In addition, let $q_{n,\theta^*}(\theta) = 2\bar{q}_{n,\theta^*}(\theta)w_{n,\theta^*}(\theta)$, where $w_{n,\theta^*}(\theta)$ is defined as in (3). Then, for any symmetry point $\theta^* \in \Theta$ and sample size n , it holds

$$\mathcal{D}[\pi_n \parallel q_{n,\theta^*}] = \mathcal{D}[\bar{\pi}_{n,\theta^*} \parallel \bar{q}_{n,\theta^*}], \quad (6)$$

where \mathcal{D} is either the TV distance, KL, reverse-KL or a generic α -divergence, while $\bar{\pi}_{n,\theta^*}(\theta)$ corresponds to the symmetrized posterior density defined in (2). In view of Lemma 1, the result in (6) implies also

$$\mathcal{D}[\pi_n \parallel q_{n,\theta^*}] \leq \mathcal{D}[\pi_n \parallel \bar{q}_{n,\theta^*}], \quad (7)$$

for any $\theta^* \in \Theta$ and sample size n .

Theorem 1 states two important results. First, as clarified in (6), the overall quality of $q_{n,\theta^*}(\theta)$ coincides with the one achieved by the unperturbed $\bar{q}_{n,\theta^*}(\theta)$ in approximating the symmetrized posterior $\bar{\pi}_{n,\theta^*}(\theta)$ defined in (2). Second, according to (7), $q_{n,\theta^*}(\theta)$ is, provably, never less accurate than the original $\bar{q}_{n,\theta^*}(\theta)$ in approximating the target posterior density $\pi_n(\theta)$, irrespectively of the chosen $\bar{q}_{n,\theta^*}(\theta)$, its symmetry point θ^* , and the sample size n . Notice also that (6) is not only an intermediate step to obtain (7), but it is of direct practical interest. Indeed, among the approximating densities $\bar{q}_{n,\theta^*}(\theta)$ with the same symmetry point θ^* , it suggests to prioritize those that give a more accurate approximation of the symmetrized posterior $\bar{\pi}_{n,\theta^*}(\theta)$ in (2), rather than the original $\pi_n(\theta)$. Therefore, although this objective goes beyond our original scope of perturbing an already-available $\bar{q}_{n,\theta^*}(\theta)$, equations (6) and (7) stimulate the development of novel symmetric approximations explicitly targeting $\bar{\pi}_{n,\theta^*}(\theta)$, rather than the original posterior $\pi_n(\theta)$. In fact, as a consequence of (6)–(7), the skewed perturbation of these approximations under the proposed strategy can yield an increasingly accurate characterization of $\pi_n(\theta)$.

The results in Theorem 1 follow from the specific form of the skewing factor $w_{n,\theta^*}(\theta)$ defined in (3). Theorem 2 proves the optimality of such a factor.

Theorem 2 (Optimality of the skewing factor). Let $\pi_n(\theta)$, $\bar{q}_{n,\theta^*}(\theta)$ and $q_{n,\theta^*}(\theta)$ be defined as in Theorem 1. Moreover, consider the alternative skew-symmetric perturbation $\tilde{q}_{n,\theta^*}(\theta) = 2\bar{q}_{n,\theta^*}(\theta)\tilde{w}_{\theta^*}(\theta)$ of $\bar{q}_{n,\theta^*}(\theta)$, where $\tilde{w}_{\theta^*}(\theta)$ is a generic skewing factor such that $\tilde{w}_{\theta^*}(\theta) \in [0, 1]$ and $\tilde{w}_{\theta^*}(\theta) = 1 - \tilde{w}_{\theta^*}(2\theta^* - \theta)$. Then, for any $\theta^* \in \Theta$, sample size n , and skewing factor $\tilde{w}_{\theta^*}(\theta)$, it holds

$$\mathcal{D}[\pi_n \parallel q_{n,\theta^*}] \leq \mathcal{D}[\pi_n \parallel \tilde{q}_{n,\theta^*}],$$

where \mathcal{D} is either the TV distance, KL, reverse-KL or a generic α -divergence.

According to Theorem 2, the skewing factor $w_{n,\theta^*}(\theta)$ in (3) is guaranteed to provide a perturbed version $q_{n,\theta^*}(\theta)$ of $\bar{q}_{n,\theta^*}(\theta)$ that is never less accurate in approximating the target posterior $\pi_n(\theta)$ when compared to any other skew-symmetric density $\tilde{q}_{n,\theta^*}(\theta)$ having symmetric component $\bar{q}_{n,\theta^*}(\theta)$ and generic skewing factor $\tilde{w}_{\theta^*}(\theta)$. Notice that to ensure $\tilde{q}_{n,\theta^*}(\theta) = 2\bar{q}_{n,\theta^*}(\theta)\tilde{w}_{\theta^*}(\theta)$ is a skew-symmetric density as in Definition 1, it suffices that the skewing factor satisfies $\tilde{w}_{\theta^*}(\theta) \in [0, 1]$ and $\tilde{w}_{\theta^*}(\theta) = 1 - \tilde{w}_{\theta^*}(2\theta^* - \theta)$. Hence, in principle, there are infinitely many options to perturb the original symmetric approximation so that the resulting density falls within the skew-symmetric class. As discussed in Section 2.2, some interesting examples of skewing factors have been derived in Azzalini and Capitanio (2003), Genton and Loperfido (2005) and Ma and Genton (2004) with a focus on generalizations of skew-normal and skew-elliptical densities, which belong to the skew-symmetric family. According to Theorem 2, all these options are sub-optimal relative to $w_{n,\theta^*}(\theta)$. This is because, unlike for other alternatives, $w_{n,\theta^*}(\theta)$ exactly matches the skewing factor of the target posterior density, when expressed in skew-symmetric form as in Proposition 2.

Besides proving the optimality of the skewing factor $w_{n,\theta^*}(\theta)$, Theorem 2 allows also to formalize the proposed skew-symmetric approximation in Definition 2 as the solution of a well defined optimization problem. This result is stated in Corollary 1, and is useful to establish direct connections with state-of-the-art approximation strategies that arise from the optimization of specific divergences, such as, for example, VB (e.g. Blei et al., 2017) and EP (e.g. Vehtari et al., 2020). Moreover, as discussed in Remark 4, it provides the premises to further expand the scope of the novel perspective considered in this article.

Corollary 1 Let $\pi_n(\theta)$, $\bar{q}_{n,\theta^*}(\theta)$ and $q_{n,\theta^*}(\theta)$ be defined as in Theorem 1. Moreover, let

$$\mathcal{Q} = \{\tilde{q}_{n,\theta^*}(\theta) : \tilde{q}_{n,\theta^*}(\theta) = 2\bar{q}_{n,\theta^*}(\theta)\tilde{w}_{\theta^*}(\theta)\},$$

denote the general family of skew-symmetric densities, which arise from the perturbation of \bar{q}_{n,θ^*} through a generic skewing factor $\tilde{w}_{\theta^*}(\theta) \in [0, 1]$ satisfying $\tilde{w}_{\theta^*}(\theta) = 1 - \tilde{w}_{\theta^*}(2\theta^* - \theta)$. Then, for any symmetry point $\theta^* \in \Theta$ and sample size n , it holds

$$\min_{\tilde{q}_{n,\theta^*} \in \mathcal{Q}} \mathcal{D}[\pi_n \parallel \tilde{q}_{n,\theta^*}] = \mathcal{D}[\pi_n \parallel q_{n,\theta^*}],$$

where \mathcal{D} is either the TV distance, KL, reverse-KL or a generic α -divergence, and, recalling (6), $\mathcal{D}[\pi_n \parallel q_{n,\theta^*}] = \mathcal{D}[\bar{\pi}_{n,\theta^*} \parallel \bar{q}_{n,\theta^*}]$.

Although the proposed skew-symmetric approximation has not been derived in Section 2 as the solution of an optimization problem, Corollary 1 clarifies that, in fact, $q_{n,\theta^*}(\theta)$ can be formalized also under such a perspective. In particular, the skew-symmetric density $q_{n,\theta^*}(\theta)$ in Definition 2 actually coincides with the solution of the constrained minimization for a suitable divergence \mathcal{D} between the target posterior $\pi_n(\theta)$ and a given approximating density $\tilde{q}_{n,\theta^*}(\theta)$ within the family \mathcal{Q} . Such a family comprises all the skew-symmetric densities having symmetric component fixed at the already-available approximating density $\bar{q}_{n,\theta^*}(\theta)$.

The above interpretation allows us to establish direct connections with the optimization-based perspectives of VB (e.g. Blei et al., 2017) and EP (e.g. Vehtari et al., 2020) solutions. However,

unlike for these strategies, Corollary 1 holds under a broader class of divergences, rather than a specific one, and, when compared to routinely implemented VB and EP schemes yielding symmetric approximations, it considers an expanded family which ensures improvements in accuracy. Note that, consistent with the focus of this article, $\bar{q}_{n,\theta^*}(\theta)$ and, as a consequence, θ^* are known and fixed in Corollary 1. Therefore, the only quantity to be derived is the skewing factor. Crucially, as clarified in Theorem 2, the solution $w_{n,\theta^*}(\theta)$ of this minimization with respect to $\tilde{w}_{\theta^*}(\theta)$ does not require optimization of additional parameters beyond the already-available θ^* . Although extending the optimization problem in Corollary 1 to the case in which also $\bar{q}_{n,\theta^*}(\theta)$ is unknown goes beyond our scope, as clarified in Remark 4, this direction can be of substantial interest to further improve the accuracy of $q_{n,\theta^*}(\theta)$, and our results open several avenues to stimulate future advancements along these lines.

Remark 4 Recalling Sections 1 and 2, the overarching focus of this article is to improve the accuracy of state-of-the-art symmetric approximations of posterior distributions via a broadly applicable perturbation scheme, which can be derived at no additional optimization costs and applied directly to the output $\bar{q}_{n,\theta^*}(\theta)$ of standard implementations. To this end, $\bar{q}_{n,\theta^*}(\theta)$ is kept fixed and known within our derivations. However, although the optimization of such a symmetric component goes beyond the scope of our contribution, combining the results in Theorem 1 and Corollary 1 with the skew-symmetric representation of posterior densities in Proposition 2, opens promising directions to further improve the approximation accuracy via the additional optimization of the symmetric component. In particular, as clarified in Corollary 1, minimizing $\mathcal{D}[\pi_n \parallel \bar{q}_{n,\theta^*}]$ also with respect to the symmetric component in $\bar{q}_{n,\theta^*}(\theta)$, simply requires to find the closest symmetric density to the symmetrized posterior in (2), and then perturb such a density with the already-derived optimal skewing factor $w_{n,\theta^*}(\theta)$. This is expected to further improve accuracy relative to perturbations of currently implemented symmetric approximations that target the actual posterior instead of its symmetrized version. In fact, to our knowledge, such a different target has never been considered before and, hence, our results can open unexplored avenues to derive improved classes of tractable deterministic approximations, along with novel computational methods to obtain these approximations. When specializing \mathcal{D} to the KL minimized under VB a promising direction could be to solve such an optimization problem via automatic differentiation variational inference schemes (Kucukelbir et al., 2017).

Section 3.2 quantifies the aforementioned accuracy gains of $q_{n,\theta^*}(\theta)$ in asymptotic settings.

3.2 Asymptotic properties

Theorems 1 and 2 provide theoretical guarantees in finite samples for the improved approximation accuracy of the newly developed skew-symmetric solution $q_{n,\theta^*}(\theta)$, compared to its symmetric counterpart, i.e. $\bar{q}_{n,\theta^*}(\theta)$. However, these results do not quantify the magnitude of such improvements. Below, we address this important point from an asymptotic perspective for $n \rightarrow \infty$. This focus further clarifies how the overall quality of the proposed skew-symmetric approximation crucially depends on the one achieved by its symmetric counterpart $\bar{q}_{n,\theta^*}(\theta)$ in approximating the symmetrized posterior $\bar{\pi}_{n,\theta^*}(\theta)$.

To address the above objectives, we study the skewed perturbations of two particular symmetric approximations, both centered at the maximum a posteriori $\theta_{\text{MAP}} = \arg \max_{\theta \in \Theta} \pi(\theta)L(\theta; \mathbf{y}_{1:n})$ (i.e. $\theta^* = \theta_{\text{MAP}}$), and approximating the target posterior density $\pi_n(\theta)$ with the same rate. Since the symmetry points coincide, the perturbations of these two symmetric approximations share the same skewing factor $w_{n,\theta_{\text{MAP}}}(\theta)$, as a direct consequence of (3). Nonetheless, as clarified in the following, the resulting skew-symmetric perturbations achieve different rates in approximating $\pi_n(\theta)$. These rates coincide with those obtained by the symmetric counterparts in approximating the symmetrized posterior $\bar{\pi}_{n,\theta_{\text{MAP}}}(\theta)$, which can be substantially different from those achieved when the target is the actual posterior $\pi_n(\theta)$. Note that, unlike for the finite sample properties in Section 3.1, which

hold broadly for any skew-symmetric approximation, the asymptotic analysis we consider below focuses on the sub-class of skewed perturbations of Gaussians or higher-order extensions of these Gaussians, both centered at the posterior mode, i.e. θ_{MAP} . Although generalizations of such an asymptotic theory to the general class of skew-symmetric approximations can be envisioned, the focus on this sub-class facilitates theoretical derivations and allows to leverage available results on the rates of Gaussian approximations. In addition, these Gaussian approximations and the corresponding higher-order versions are directly related to ubiquitous Laplace type methods, and hence, it is natural to deepen the asymptotic analysis of the skew-symmetric approximations arising from the perturbation of such symmetric densities. Note that, although we consider the MAP as symmetry point, our theory relies on assumptions which allow to replace θ_{MAP} with any efficient estimator of θ .

To derive the aforementioned asymptotic results, let $\ell_n(\theta) = \ell(\theta; \mathbf{y}_{1:n}) = \log L(\theta; \mathbf{y}_{1:n})$ and $\log \pi(\theta)$ be, respectively, the log-likelihood function and log-prior density, both evaluated at $\theta \in \Theta \subseteq \mathbb{R}^d$. Moreover, denote with

$$\ell_{n,\theta}^{(k)} = (\partial^{\otimes k} / \partial \theta^{\otimes k} \ell_n(\theta)), \quad \text{and} \quad \log \pi_{\theta}^{(k)} = (\partial^{\otimes k} / \partial \theta^{\otimes k} \log \pi(\theta)),$$

the d^k -dimensional arrays containing the associated k th order partial derivatives at θ . Under these settings, the first symmetric density whose skewed perturbation is studied in asymptotic regimes is the Gaussian approximation from the Laplace method. Namely,

$$\bar{q}_{n,\theta_{\text{MAP}},(1)}(\theta) = \phi_d(\theta; \theta_{\text{MAP}}, \mathbf{J}_{\theta_{\text{MAP}}}^{-1}), \tag{8}$$

where $\theta_{\text{MAP}} = \arg \max_{\theta \in \Theta} \pi(\theta)L(\theta; \mathbf{y}_{1:n})$ and $\mathbf{J}_{\theta_{\text{MAP}}} = -(\ell_{n,\theta_{\text{MAP}}}^{(2)} + \log \pi_{\theta_{\text{MAP}}}^{(2)})$.

The second is, instead, a novel higher-order extension of the Gaussian in (8), which relates to the family of semi-nonparametric distributions (Gallant & Nychka, 1987) and has density

$$\bar{q}_{n,\theta_{\text{MAP}},(2)}(\theta) = \frac{\phi_d(\theta; \theta_{\text{MAP}}, \mathbf{J}_{\theta_{\text{MAP}}}^{-1})f(\theta - \theta_{\text{MAP}})}{\int \phi_d(\theta; \theta_{\text{MAP}}, \mathbf{J}_{\theta_{\text{MAP}}}^{-1})f(\theta - \theta_{\text{MAP}}) d\theta}, \tag{9}$$

where $f(\theta - \theta_{\text{MAP}})$ is a non-negative polynomial obtained from a fourth-order expansion of the symmetrized log-posterior at θ_{MAP} that yields

$$f(\theta - \theta_{\text{MAP}}) = 1 + \frac{1}{24} \langle \ell_{n,\theta_{\text{MAP}}}^{(4)}, (\theta - \theta_{\text{MAP}})^{\otimes 4} \rangle + \frac{1}{2} \left[\frac{1}{24} \langle \ell_{n,\theta_{\text{MAP}}}^{(4)}, (\theta - \theta_{\text{MAP}})^{\otimes 4} \rangle \right]^2 + \frac{1}{2} \left[\frac{1}{6} \langle \ell_{n,\theta_{\text{MAP}}}^{(3)}, (\theta - \theta_{\text{MAP}})^{\otimes 3} \rangle \right]^2.$$

In the above expression, the generic quantity $\langle \ell_{n,\theta_{\text{MAP}}}^{(k)}, (\theta - \theta_{\text{MAP}})^{\otimes k} \rangle$ denotes the polynomial term associated to the k th element of the Taylor expansion for $\ell_n(\theta)$ at θ_{MAP} . Notice that, in view of Lemma B.2 within Appendix B of the online supplementary material, $f(\theta - \theta_{\text{MAP}})$ is always non-negative, and hence, (9) is a proper density function symmetric about θ_{MAP} . Although (9) is of less direct applicability than (8), such a higher-order approximating density can capture the behaviour of the symmetrized posterior more flexibly than the Gaussian in (8). Therefore, it provides an interesting theoretical alternative to (8) for quantifying the gains of the proposed skew-symmetric family of approximations when applied to symmetric densities achieving different accuracies in capturing the behaviour of the symmetrized posterior. Since (8) and (9) share the same symmetry point θ_{MAP} , the resulting skew-symmetric perturbations

$$q_{n,\theta_{\text{MAP}},(1)}(\theta) = 2\bar{q}_{n,\theta_{\text{MAP}},(1)}(\theta)w_{n,\theta_{\text{MAP}}}(\theta), \quad \text{and} \quad q_{n,\theta_{\text{MAP}},(2)}(\theta) = 2\bar{q}_{n,\theta_{\text{MAP}},(2)}(\theta)w_{n,\theta_{\text{MAP}}}(\theta),$$

have the same skewing factor $w_{n,\theta_{\text{MAP}}}(\theta)$.

Theorem 3 below formalizes the above discussion. When the dimension d is fixed, this theorem states that several key divergences (including the TV distance, KL, reverse-KL and α -divergences) between the target posterior density and the skew-symmetric perturbation $q_{n,\theta_{\text{MAP}},(1)}(\boldsymbol{\theta})$ of (8) converge to zero in probability with a rate $1/n$, up to a poly-log term, which substantially improves the $1/\sqrt{n}$ rate of the unperturbed Gaussian approximation in (8). This accuracy gain is further refined by $q_{n,\theta_{\text{MAP}},(2)}(\boldsymbol{\theta})$ which achieves a $1/n^2$ rate, again up to a poly-log term. Such a latter result is even more remarkable, provided that (9) has the same rate of (8) in approximating the target posterior $\pi_n(\boldsymbol{\theta})$. This is due to the fact that both (8) and (9) are not able to capture the skewness of the posterior distribution. However, when the target is the symmetrized posterior $\bar{\pi}_{n,\theta_{\text{MAP}}}(\boldsymbol{\theta})$ (i.e. the skewness is removed from the target), the higher-order extension (9) of (8) improves the rates of such a latter Gaussian approximation. Combining this result with Theorem 1, the same level of accuracy is maintained by $q_{n,\theta_{\text{MAP}},(2)}(\boldsymbol{\theta})$ in approximating the target posterior density $\pi_n(\boldsymbol{\theta})$. This explains the improved rates of $q_{n,\theta_{\text{MAP}},(2)}(\boldsymbol{\theta})$ relative to $q_{n,\theta_{\text{MAP}},(1)}(\boldsymbol{\theta})$, and provides a result with important methodological consequences. In particular, it suggests to target the symmetrized posterior rather than the original one with state-of-the-art symmetric approximations, and then perturb these approximations with the proposed skewing factor.

The proof of Theorem 3 is provided in [Appendix A of the online supplementary material](#) and relies on a number of regularity conditions listed below. These conditions are similar to those recently considered by [Durante et al. \(2024\)](#) for deriving a skewed extension of the Bernstein-von Mises theorem (e.g. [Van der Vaart, 2000](#)). Recalling related discussions in [Durante et al. \(2024\)](#), we shall emphasize that Assumptions 1–6 provide natural extensions of standard conditions in asymptotic studies of this type, and hold under several, commonly adopted, regular statistical models, such as generalized linear models.

Assumption 1 For every sample size $n \in \mathbb{N}$, the data $\mathbf{y}_{1:n}$ are realizations from a sequence of random variables $\mathbf{Y}_{1:n}$ having true underlying distribution P_0^n which may not be necessarily included within the parametric family $\{P_\theta^n, \theta \in \Theta\}$ defining the assumed model with log-likelihood $\ell_n(\boldsymbol{\theta})$. Furthermore, the KL projection $P_{\theta_0}^n$ of P_0^n into $\{P_\theta^n, \theta \in \Theta\}$, is unique.

Assumption 2 The log-prior density $\log \pi(\boldsymbol{\theta})$ is four times continuously differentiable inside a neighbourhood of $\boldsymbol{\theta}_0$, and $0 < \pi(\boldsymbol{\theta}_0) < \infty$.

Assumption 3 For every $M_n \rightarrow \infty$ there exists a positive constant c_1 such that

$$\lim_{n \rightarrow \infty} P_0^n \left\{ \sup_{\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\| > M_n \sqrt{d}/\sqrt{n}} \{\ell_n(\boldsymbol{\theta}) - \ell_n(\boldsymbol{\theta}_0)\}/n < -c_1 M_n^2 d/n \right\} = 1,$$

where $\ell_n(\boldsymbol{\theta}_0)$ is the log-likelihood evaluated at $\boldsymbol{\theta}_0$.

Assumption 4 The maximum a posteriori estimator $\boldsymbol{\theta}_{\text{MAP}} = \arg \max_{\boldsymbol{\theta} \in \Theta} \{\ell_n(\boldsymbol{\theta}) + \log \pi(\boldsymbol{\theta})\}$, satisfies $\mathbb{E}_0^n \|\boldsymbol{\theta}_{\text{MAP}} - \boldsymbol{\theta}_0\|^2 = O(d/n)$.

Assumption 5 There exist two positive constants $\bar{\eta}_1$ and $\bar{\eta}_2$ such that the event

$$\tilde{A}_{n,0} = \{\lambda_{\text{MIN}}(\mathbf{J}_{\boldsymbol{\theta}_{\text{MAP}}}/n) > \bar{\eta}_1\} \cap \{\lambda_{\text{MAX}}(\mathbf{J}_{\boldsymbol{\theta}_{\text{MAP}}}/n) < \bar{\eta}_2\}$$

holds with probability $P_0^n(\tilde{A}_{n,0}) = 1 - o(1)$. Moreover, there exist two positive constants $\delta > 0$ and $L > 0$ such that the inequalities

$$\|\ell_n^{(3)}(\boldsymbol{\theta})/n\| < L, \quad \|\ell_n^{(4)}(\boldsymbol{\theta})/n\| < L, \quad \|\log \pi^{(2)}(\boldsymbol{\theta})\| < L,$$

hold uniformly in $B_\delta(\boldsymbol{\theta}_{\text{MAP}}) = \{\boldsymbol{\theta} \in \Theta : \|\boldsymbol{\theta}_{\text{MAP}} - \boldsymbol{\theta}\| < \delta\}$, with P_0^n -probability tending to one, where $\|\cdot\|$ denotes the spectral norm. When $q_{n,\boldsymbol{\theta}_{\text{MAP}},(2)}(\boldsymbol{\theta})$ is considered, for the same $\delta > 0$ and $L > 0$ as above, also the inequalities

$$\|\ell_n^{(5)}(\boldsymbol{\theta})/n\| < L, \quad \|\ell_n^{(6)}(\boldsymbol{\theta})/n\| < L, \quad \|\log \pi^{(3)}(\boldsymbol{\theta})\| < L, \quad \|\log \pi^{(4)}(\boldsymbol{\theta})\| < L,$$

hold uniformly in $B_\delta(\boldsymbol{\theta}_{\text{MAP}})$, with P_0^n -probability tending to one.

Assumption 6 Let $\mathbf{h}_{\text{MAP}} = \sqrt{n}(\boldsymbol{\theta} - \boldsymbol{\theta}_{\text{MAP}})$. Then, there exist positive constants c and M , such that $|\log \bar{\pi}_{n,\boldsymbol{\theta}_{\text{MAP}}}(\mathbf{h}_{\text{MAP}})| \leq c\|\mathbf{h}_{\text{MAP}}\|^M \vee 1$.

Assumptions 1–3 are needed to guarantee that the assumed model is sufficiently regular and that the induced posterior distribution $\pi_n(\boldsymbol{\theta})$ concentrates asymptotically fast enough around $\boldsymbol{\theta}_0$. Similarly, for d fixed, Assumption 4 provides the parametric $1/\sqrt{n}$ concentration rate of the posterior mode $\boldsymbol{\theta}_{\text{MAP}}$ around $\boldsymbol{\theta}_0$, required to control the likelihood around the random location $\boldsymbol{\theta}_{\text{MAP}}$. Assumption 5 presents, instead, requirements on the regularity of the prior and the likelihood, depending on the approximation procedure considered. The second, more complex, approximation $q_{n,\boldsymbol{\theta}_{\text{MAP}},(2)}(\boldsymbol{\theta})$ requires a stronger control, but the corresponding rates are also faster. These conditions allow us to control the accuracy of the Taylor expansion for the log-posterior, up to the required order, and for its approximation induced by either $q_{n,\boldsymbol{\theta}_{\text{MAP}},(1)}(\boldsymbol{\theta})$ or $q_{n,\boldsymbol{\theta}_{\text{MAP}},(2)}(\boldsymbol{\theta})$. Finally, Assumption 6 controls the tail behaviour of the symmetrized posterior to ensure finite cross entropy between the skewed perturbations of (8) (or (9)) and $\pi_n(\boldsymbol{\theta})$. While this condition could be possibly relaxed, we conjecture that explicit control of the tails is necessary to obtain the results for the KL and reverse-KL divergences in Theorem 3.

The combination of Assumptions 1–6 described above allows us to state the following theorem regarding the asymptotic accuracy of $q_{n,\boldsymbol{\theta}_{\text{MAP}},(1)}(\boldsymbol{\theta})$ and $q_{n,\boldsymbol{\theta}_{\text{MAP}},(2)}(\boldsymbol{\theta})$ in approximating the target posterior $\pi_n(\boldsymbol{\theta})$.

Theorem 3 (Asymptotic accuracy). Under Assumptions 1–5, it holds

$$\mathcal{D}[\pi_n \parallel q_{n,\boldsymbol{\theta}_{\text{MAP}},(1)}] = O_{P_0^n}(M_n^{c_3} d^3/n), \tag{10}$$

and

$$\mathcal{D}[\pi_n \parallel q_{n,\boldsymbol{\theta}_{\text{MAP}},(2)}] = O_{P_0^n}(M_n^{c_4} d^6/n^2), \tag{11}$$

where $M_n = \sqrt{c_0 \log n}$, the quantities $c_0, c_3, c_4 > 0$ are fixed positive constants not depending on n and d , and \mathcal{D} denotes either the TV distance or any generic α -divergence with $\alpha \in (0, 1)$. If Assumption 6 holds (in addition to 1–5), then (10)–(11) are valid also when \mathcal{D} is the KL or reverse-KL divergence.

Remark 5 Under the same assumptions, it can be also shown that the convergence rate of the symmetric components (8) and (9) to the target posterior density $\pi_n(\boldsymbol{\theta})$ is $O_{P_0^n}(M_n^{c_5} d^{3/2}/\sqrt{n})$ for some $c_5 > 0$. This result suggests that the difference in (10)–(11) are due to how the symmetric component approximates $\bar{\pi}_{n,\boldsymbol{\theta}_{\text{MAP}}}(\boldsymbol{\theta})$.

Remark 6 All the rates presented in Theorem 3 and Remark 5 rely on bounds which are guaranteed to vanish also when the dimension d grows with the sample size n , as long as $d = o(n^{1/3})$, up to a poly-log term. This condition is common in high-dimensional studies of Gaussian approximations (Panov & Spokoiny, 2015; Spokoiny, 2025; Spokoiny & Panov, 2025). However, unlike these asymptotic studies, the rates presented in Theorem 3 vanish with n (or even n^2), up to a

poly-log term, instead of \sqrt{n} , for any fixed d . Finally, let us emphasize that the constants c_0 , c_3 and c_4 in Theorem 3 arise from specific technicalities in the proof that can be found in the [online supplementary material](#). For what concerns the interpretation of the rates in Theorem 3, the only key requirement is that c_0 , c_3 and c_4 are fixed positive constants not depending on n and d .

Note that $q_{n,\theta_{\text{MAP}}(1)}$ is closely related to the skew-modal approximation recently introduced in [Durante et al. \(2024\)](#). These two approximations share a similar symmetric component and the same order of convergence to the target posterior, but have different skewing factor. In [Durante et al. \(2024\)](#) such a quantity is derived using asymptotic arguments that involve the evaluation of third-order log-likelihood derivatives at θ_{MAP} . Conversely, the skewing factor $w_{n,\theta_{\text{MAP}}}(\theta)$ introduced in this article depends only on the un-normalized posterior density, and hence, is simpler to implement. Moreover, while the focus of [Durante et al. \(2024\)](#) is on perturbing Gaussian approximations from the Laplace method to refine the classical Bernstein-von Mises theorem, the proposed skew-symmetric approximations broadly apply to any symmetric density, beyond Gaussians centered at the MAP. Besides its practical benefits illustrated in Sections 4 and 5 below, this generality yields also theoretical gains which are evident in (11). This latter result clarifies that the theory in [Durante et al. \(2024\)](#) can be further expanded to higher-order approximations achieving faster convergence rates, while remaining within the skew-symmetric family.

4 Simulation studies

Before assessing the skew-symmetric approximations in real-data applications, let us first consider two simulation studies quantifying to what extent the accuracy gains and improved rates derived theoretically in Sections 3.1 and 3.2 find empirical evidence also in practice, even beyond the settings explored by our theory. To this end, we compare the accuracies achieved by three popular Gaussian approximations from the Laplace method ([Gelman et al., 2013](#), Chapter 13), black-box VB ([Kucukelbir et al., 2017](#); [Ranganath et al., 2014](#)) and EP ([Minka, 2001](#); [Vehtari et al., 2020](#)), with those obtained under the corresponding skew-symmetric counterparts. These gains are quantified through different accuracy measures in two simulations studies. The first (see Section 4.1) aims at providing empirical evidence for the asymptotic rates presented in Theorem 3 (see also Remark 5), under a parametric setting that meets the conditions underlying the theory within Section 3.2. In this first study, we consider a simple one-dimensional scenario, which allows us to evaluate with precision the different divergences in Theorem 3 via numerical integration methods. This provides a reliable quantification of the rates, which is not affected by Monte Carlo error. Large d settings are instead addressed in the second simulation study (see Section 4.2) whose aim is to explore the accuracy improvements of the skew-symmetric approximations beyond the settings and quantities studied theoretically in Sections 3.1 and 3.2. In particular, the focus of Section 4.2 is on assessing the accuracy in approximating the posterior marginals under high-dimensional regimes where d grows with n at a rate that does not meet the condition $d = o(n^{1/3})$ discussed in Remark 6.

4.1 One-dimensional Poisson model

The upper bound (10) presented in Theorem 3 suggests that, for a sufficiently large n and fixed d , the log-divergence between the target posterior and the proposed skew-symmetric perturbation of suitable Gaussian approximations should decrease as a linear function of $\log n$ with slope that is -1 or lower. According to the discussion in Remark 5, this slope should be, instead, -0.5 , or lower, for the unperturbed Gaussian counterpart.

To assess whether the above results find empirical evidence in practice, we simulate i.i.d. data $y_{1:n}$, for growing sample size n from $n = 15$ to $n = 145$ with step size 10, from a Poisson distribution having rate $\exp(\theta_0) = 1$. Leveraging these samples of growing size we perform Bayesian inference by assuming a one-dimensional Poisson model with rate $\exp(\theta)$, where $\theta \in \mathbb{R}$ is assigned a Student- t prior with one degree of freedom. Such a prior choice coincides with a hierarchical Bayesian formulation assuming a zero mean Gaussian prior for θ combined with a χ_1^2 hyperprior for its precision parameter, thereby providing a realistic elicitation which is aligned with the common practice of considering higher-level priors for the hyperparameters to mitigate sensitivity

Table 1. Empirical comparison of the rates achieved by state-of-the-art Gaussian approximations from the Laplace method, black-box VB and EP, with those obtained under the corresponding skew-symmetric perturbations

	TV	KL	reverse-KL
Laplace (Gaussian)	-0.48 (0.01)	-0.93 (0.02)	-0.97 (0.02)
Laplace (Skew-symmetric)	-1.04 (0.02)	-1.80 (0.08)	-3.11 (0.26)
black-box VB (Gaussian)	-0.48 (0.01)	-0.95 (0.02)	-0.98 (0.02)
black-box VB (Skew-symmetric)	-1.05 (0.02)	-1.73 (0.12)	-3.18 (0.29)
EP (Gaussian)	-0.47 (0.01)	-0.93 (0.02)	-0.99 (0.02)
EP (Skew-symmetric)	-0.99 (0.02)	-1.76 (0.11)	-3.47 (0.26)

Note. For three routinely employed divergences \mathcal{D} (i.e. TV, KL and reverse-KL) the table displays the slope of the fitted linear regression between $\log \mathcal{D}(\pi_n \parallel q_n)$ and $\log n$, where q_n is either \bar{q}_{n,θ^*} (i.e. Gaussian) or q_{n,θ^*} (i.e. Skew-symmetric), and n is defined on a grid from 15 to 145 with step size 10. The results are averaged across 50 replicated experiments (standard error within brackets). Bold values indicate the best performance for each pair of Gaussian and skew-symmetric approximation.

issues. These settings yield a sequence of posteriors $\pi_n(\theta)$, $n = 15, 25, 35, \dots, 135, 145$, for θ , which we approximate under the three Gaussian densities resulting from the Laplace method (Gelman et al., 2013, Chapter 13), black-box VB (Kucukelbir et al., 2017; Ranganath et al., 2014) and EP (Minka, 2001; Vehtari et al., 2020), respectively, along with the corresponding skew-symmetric perturbations presented in Section 2. As anticipated in Sections 1 and 2, all these approximations can be readily derived leveraging standard softwares. In particular, the three Gaussians alternatives can be obtained as a direct output of `rstan` or simple R coding, whereas the skew-symmetric counterparts only require the additional calculation of the closed-form skewing factor, without further optimization.

Table 1 quantifies the rates achieved by the aforementioned approximations under routinely employed divergences among those considered in Theorem 3, i.e. TV, KL and reverse-KL. Consistent with the motivation underlying this simulation study, we display, in particular, the slope of the fitted linear regression between $\log \mathcal{D}(\pi_n \parallel q_n)$ and $\log n$, for $n = 15, 25, 35, \dots, 135, 145$, where \mathcal{D} is the TV, KL or reverse-KL, and q_n is either one of the Gaussian approximations under analysis (i.e. $q_n = \bar{q}_{n,\theta^*}$) or its skew-symmetric counterpart (i.e. $q_n = q_{n,\theta^*}$). In Table 1, these slopes are averaged across 50 replicated experiments and the associated standard errors are reported within brackets. Consistent with Theorem 3 and Remark 5, the proposed skew-symmetric approximations display slopes close to -1 or below it, across all divergences and approximation methods, whereas those of the unperturbed Gaussian counterparts are approximately -0.5 or lower. Moreover, the improvements in rates achieved by the proposed perturbation strategy over its symmetric alternative are systematic and by at least a multiplicative factor of 2. These results provide empirical evidence on the fact that the finite sample accuracy gains proved in Theorem 1 and the asymptotic rates derived in Theorem 3 are visible also in practice. Interestingly, the empirical results obtained for the KL and reverse-KL further suggest that sharper upper bounds could be derived in Theorem 3 for these two divergences, thereby stimulating future research along these lines. Section 4.2 quantifies to what extent similar gains can be preserved in more challenging settings and for quantities of interest beyond those studied by our theory in Section 3.

4.2 High-dimensional Poisson regression

Let us now consider a high-dimensional setting where $d = 6 \lfloor \sqrt{n/2} \rfloor$, and the focus is on assessing the accuracy improvements in approximating posterior marginals for each θ_j , $j = 1, \dots, d$. This setting is substantially more challenging than the one considered in Section 4.1 for two reasons. First, it does not meet the $d = o(n^{1/3})$ condition that is required to obtain vanishing rates under Theorem 3 (see also Remark 6). Second, the finite sample accuracy gains proved in Theorem 1 focus on the approximation of the joint posterior density $\pi_n(\theta)$, rather than its marginals $\pi_n(\theta_j)$, for $j = 1, \dots, d$. As such, it is of interest to assess whether the theoretical results in Section 3 and the

Table 2. Relative accuracy improvement (in percentage) achieved by each skew-symmetric approximation with respect to three symmetric counterparts, corresponding to the Gaussians from the Laplace method, black-box VB and EP

	$n = 200$ ($d = 60$)	$n = 400$ ($d = 84$)	$n = 600$ ($d = 102$)	$n = 800$ ($d = 120$)	$n = 1000$ ($d = 132$)	% IMPROV
Laplace (Skew-sym. vs. Gaus.)						
TV	12.9%	11.1%	10.3%	9.2%	9.5%	100%
KL	28.0%	22.7%	20.2%	18.6%	17.3%	99.8%
reverse-KL	37.1%	39.2%	40.7%	40.9%	41.2%	100%
bias	32.8%	34.3%	34.0%	34.8%	35.3%	99.8%
black-box VB (Skew-sym. vs. Gaus.)						
TV	14.1%	12.3%	12.0%	12.7%	11.6%	100%
KL	31.6%	25.7%	22.0%	20.9%	19.1%	100%
reverse-KL	42.4%	42.2%	41.0%	43.7%	41.5%	100%
bias	55.2%	48.7%	51.2%	56.1%	50.8%	95%
EP (Skew-sym. vs. Gaus.)						
TV	17.3%	16.1%	16.7%	16.7%	16.8%	100%
KL	33.1%	28.7%	26.4%	24.9%	24.1%	100%
reverse-KL	43.3%	44.6%	46.6%	46.6%	46.9%	100%
bias	24.7%	31.9%	30.6%	32.3%	36.8%	86%

Note. The improvement is quantified, for varying sample sizes $n = 200, 400, 600, 800, 1000$, under four different measures, namely the bias in approximating the posterior means of the standardized parameters, and the TV, KL and reverse-KL divergences between the target posterior marginals and the corresponding approximation, for each $j = 1, \dots, d$ (with $d = 6\lfloor\sqrt{n/2}\rfloor$). Standardization involves dividing the parameters by the standard deviations of the associated exact posterior. As such these standard deviations are estimated via Monte Carlo using 40,000 samples from the target posterior produced by HMC via `rstan`. The relative improvement is computed under the four measures analyzed for every $j = 1, \dots, d$. The table displays the median relative improvement over j , for each $n = 200, 400, 600, 800, 1000$, averaged across 10 replicated experiments. The quantity % IMPROV corresponds, instead, to the percentage of marginals for which the proposed skew-symmetric approximation improves the Gaussian counterpart with respect to the associated accuracy measure. This percentage is averaged across n .

empirical gains in Section 4.1 find some numerical evidence also in these more challenging settings. In fact, as discussed in Section 2, the skewness-inducing correction we propose relies on a global notion of central symmetry. Therefore, it is unclear whether re-distributing the density with respect to such a notion translates into empirical gains in characterizing local asymmetries within the marginals, which are often of interest in practice (e.g. Chopin & Ridgway, 2017; Rue et al., 2009; Tierney & Kadane, 1986).

To address the above goal, we simulate data $y_{1:n}$ for growing sample size $n = 200, 400, 600, 800, 1000$, and dimension $d = 6\lfloor\sqrt{n/2}\rfloor$, from a Poisson regression model with the generic response variable y_i having rate $\exp(\mathbf{x}_i^\top \boldsymbol{\theta}_0)$, $i = 1, \dots, n$. In this simulation, each \mathbf{x}_i denotes a d -dimensional vector with unbalanced binary entries drawn independently from Bernoulli variables having low probabilities generated by a uniform in $[0.02, 0.1]$, while $\boldsymbol{\theta}_0$ has entries simulated from independent uniforms in $[-0.5, 0.5]$, and further rescaled to control the variance of the linear predictor as d grows with n . Extending the formulation considered in Section 4.1, we model these data under a Bayesian Poisson regression with rates $\exp(\mathbf{x}_i^\top \boldsymbol{\theta})$, for $i = 1, \dots, n$, and assume independent Student- t priors with one degree of freedom for each θ_j , $j = 1, \dots, d$.

Table 2 quantifies, for each sample size n , whether the posterior marginals $\pi_n(\theta_j)$, $j = 1, \dots, d$, resulting from the above Bayesian model are more accurately approximated by those of the three Gaussians considered in Section 4.1 (i.e. $\bar{q}_{n,\theta^*}(\theta_j)$, $j = 1, \dots, d$), or by the corresponding skew-symmetric perturbations (i.e. $q_{n,\theta^*}(\theta_j)$, $j = 1, \dots, d$). This is accomplished by computing, for every n and j , the relative improvement (in percentage) achieved by the marginals of the proposed skew-symmetric approximation with respect to those of the Gaussian counterparts, under four

accuracy measures. These are the TV, KL and reverse-KL divergences between the target posterior marginals and the corresponding approximations, along with the bias in estimating the posterior expectations of the standardized parameters (see the caption of [Table 2](#) for details on the standardization). This bias is useful to assess if the accurate characterization of the marginal densities translates also into practical improvements in approximating functionals of key interest for inference. Note that, unlike in [Section 4.1](#) (where the focus is on one-dimensional settings), all the performance measures discussed above are estimated here via Monte Carlo leveraging 10,000 i.i.d. draws from each of the approximations studied, while considering 4 chains of size 10,000 produced by HMC (via `rstan`) as benchmark samples from the target posterior $\pi_n(\theta)$.

[Table 2](#) displays the medians of the above relative accuracy gains across j , further averaged over 10 replicated experiments. The results show that, albeit tested in a more challenging setting, the proposed skew-symmetric approximation still showcases systematic improvements over the Gaussian counterparts. Such gains are non-negligible for all the measures, and particularly remarkable for the reverse-KL and bias. Interestingly, the ability to reduce both the KL and the reverse-KL divergences clarifies that the proposed perturbation is able to improve the original Gaussians in both mode seeking and mass covering behaviours, leading also to reductions in bias. According to %IMPROV, these improvements are again systematic across marginals, sample sizes and approximation methods. Notice that, the 86% associated with the bias of EP should be interpreted as a remarkable result. In fact, EP aims, by design, at providing Gaussian approximations centered at the actual posterior mean. As such, its skewed perturbation may not yield, in principle, noticeable improvements when the focus is specifically on the bias. [Table 2](#) shows that, in practice, the skew-symmetric approximation provides clear gains also under this measure, likely due to a positive effect in removing possible biases that arise from EP optimization schemes.

[Section 5](#) clarifies that the improvements achieved by the skew-symmetric approximations in simulations are found also in real-data applications, both in relative and absolute terms.

5 Real-data applications

We conclude by showcasing the practical gains of the proposed skew-symmetric approximation in two real-data applications requiring realistic hierarchical Bayesian implementations. The first (see [Section 5.1](#)) shows that relative improvements similar to those presented in the simulation study in [Section 4.2](#) can be obtained also under a more sophisticated zero-inflated negative binomial regression aimed at modelling school attendance behaviours leveraging the dataset `Attendance` in the R library `mixpoissonreg`. The second (see [Section 5.2](#)) considers, instead, a semi-parametric hierarchical logistic regression model to infer the determinants of contraceptive usage behaviours in 33 Indian states based on data retrieved from the HDS II survey (see, e.g. [Rigon et al., 2019](#)). In this latter application $n = 30,524$ and $d = 62$, thereby providing an interesting setting to study also the runtimes of inference under the proposed skew-symmetric approximation, in addition to its gains with respect to the Gaussian counterparts in absolute terms.

Consistent with the simulation studies in [Section 4](#), the above gains are again quantified with a focus on the skew-symmetric perturbations for the Gaussian approximations from the Laplace method ([Gelman et al., 2013](#), Chapter 13), black-box VB ([Kucukelbir et al., 2017](#); [Ranganath et al., 2014](#)) and EP ([Minka, 2001](#); [Vehtari et al., 2020](#)), leveraging the four accuracy measures in [Table 2](#) (i.e. TV, KL, reverse-KL and bias). As for the analyses within [Section 4.2](#), also in the two applications these measures are estimated via Monte Carlo based on 10,000 samples drawn from each of the approximations studied. Inference under the target posterior, which provides the benchmark to assess the accuracy of the approximations analyzed, leverages instead 4 chains of length 10,000 under HMC from the R library `rstan`. This library is also employed for obtaining the Laplace and black-box VB approximations, via the functions `optimizing` and `vb`, respectively. The EP solution is instead implemented via a custom-built R code. Once these three symmetric approximations are available, the corresponding skew-symmetric counterparts can be obtained as in [Definition 2](#), and i.i.d. sampling from these skewed corrections proceeds via [Algorithm 1](#).

In interpreting the results in [Sections 5.1–5.2](#) (along with those in [Sections 4.1–4.2](#)) it is important to emphasize that the proposed skewness-inducing mechanism is constrained to redistribute the density of the symmetric approximation only among pairs $\{\theta, 2\theta^* - \theta\}$, $\theta \in \Theta$. Thus, the systematic

Table 3. Relative accuracy improvement (in percentage) achieved by each skew-symmetric approximation with respect to three symmetric counterparts, corresponding to the Gaussians from the Laplace method, black-box VB and EP

	TV	KL	reverse-KL	bias	% IMPROV
Laplace (Skew-sym. vs. Gaus.)	22.0%	39.5%	47.5%	29.5%	100%
black-box VB (Skew-sym. vs. Gaus.)	22.0%	46.1%	33.4%	26.8%	96%
EP (Skew-sym. vs. Gaus.)	21.2%	36.3%	34.2%	27.0%	90%

Note. The improvement is quantified under four different measures, namely the bias in approximating the posterior means of the standardized parameters, and the TV, KL and reverse-KL divergences between the target posterior marginals and the corresponding approximation. Standardization proceeds as discussed in the caption of Table 2. The relative gain is computed under these four measures for every $j = 1, \dots, 9$. The table displays the median relative improvement over j . The quantity % IMPROV corresponds, instead, to the percentage of marginals for which the skew-symmetric approximation improves the Gaussian counterpart, averaged across the accuracy measures.

improvements achieved by the skew-symmetric approximations in Tables 1–3 and Figure 1 are arguably remarkable when accounting for the simplicity of the proposed perturbation strategy. These results suggest future research to devise even more sophisticated skewness-inducing mechanisms redistributing the density at more flexible configurations than $\{\theta, 2\theta^* - \theta\}$, $\theta \in \Theta$. Such extensions should, however, account for the tractability of the resulting approximation. Recalling Section 2.3 and Remark 1, this is a key property of the skewed perturbation we propose, which is fundamental to facilitate its broad applicability. To extend this applicability even further, we illustrate in the two real-data applications the use of skew-symmetric approximations as improved proposals in importance sampling (see, e.g. Chopin & Papaspiliopoulos, 2020, Chapter 8). This perspective yields remarkable gains in the effective sample sizes (ESS) relative to those obtained when leveraging the Gaussian counterparts as proposals.

5.1 Zero-inflated negative binomial regression

As a first application let us consider a zero-inflated negative binomial regression (e.g. Lawal, 2012) applied to a study on school attendance of $n = 314$ students. The dataset is available within the R library `mixpoissonreg` and comprises, for each student i , information on the number of days of absence at school y_i , along with two categorical explanatory variables x_{i1} and x_{i2} corresponding, respectively, to the gender of the student (`female/male`) and the instructional programme in which such a student is enrolled (`General/Academic/Vocational`). To learn how the school attendance behaviours relate to these two variables, while accounting for the excess of zeros and overdispersion in the observed responses y_1, \dots, y_n , we consider a zero-inflated negative binomial regression (e.g. Lawal, 2012). This model incorporates gender and instructional programme effects both in the zero inflation probabilities ψ_i , and also in the negative binomial expected counts μ_i . In particular, denoting with $\mathbb{1}(\cdot)$ the indicator function, we let

$$\begin{aligned} \text{logit}(\psi_i) &= \alpha_0 + \alpha_1 \mathbb{1}(x_{i1} = \text{male}) + \alpha_2 \mathbb{1}(x_{i2} = \text{Academic}) + \alpha_3 \mathbb{1}(x_{i2} = \text{Vocational}), \\ \log \mu_i &= \beta_0 + \beta_1 \mathbb{1}(x_{i1} = \text{male}) + \beta_2 \mathbb{1}(x_{i2} = \text{Academic}) + \beta_3 \mathbb{1}(x_{i2} = \text{Vocational}), \end{aligned}$$

for every unit $i = 1, \dots, n$, and assume a constant overdispersion $\exp(\gamma)$, with $\gamma \in \mathbb{R}$. In order to perform Bayesian inference on the parameters $\theta = (\gamma, \alpha_0, \alpha_1, \alpha_2, \alpha_3, \beta_0, \beta_1, \beta_2, \beta_3) \in \mathbb{R}^9$ we rely on independent Gaussian priors having mean 0 and variance 2, and then assess in Table 3 whether the proposed skewed perturbation improves the Gaussian alternatives from the Laplace method (e.g. Gelman et al., 2013, Chapter 13), black-box VB (Kucukelbir et al., 2017; Ranganath et al., 2014) and EP (Minka, 2001; Vehtari et al., 2020) in approximating the resulting posterior. Note that, unlike for the simulations in Sections 4.1–4.2, we consider here Gaussian priors rather than Student- t ones. This is motivated by the aim to assess our skew-symmetric approximations under several elicitation of routine use, including those relying on Gaussian priors.

Consistent with the simulations in Section 4, Table 3 demonstrates that the proposed skew-symmetric approximations systematically outperform the corresponding symmetric counterparts

for all the measures considered also in this realistic real-data application. Recalling the remarks before Section 5.1, these relative gains are particularly remarkable, especially in the light of the simplicity of the perturbation strategy employed, and can be further strengthened by studying its use as an improved proposal within importance sampling targeting the joint posterior density $\pi_n(\theta)$ (e.g. [Chopin & Papaspiliopoulos, 2020](#), Chapter 8). Such an additional assessment is not only useful to clarify the broad applicability of the proposed skew-symmetric approximation in Bayesian computation, but also complements the analyses within [Table 3](#). In fact, the magnitude of the gains in the ESS (e.g. [Chopin & Papaspiliopoulos, 2020](#), Chapter 8.6) relative to those obtained under the unperturbed Gaussian proposals are useful for quantifying the improvements of the proposed skew-symmetric approximation in characterizing the shape of the entire target posterior density $\pi_n(\theta)$, beyond its marginals. The percentage gains in the ESS, estimated in 100 replicated studies, are 90.95% (sd: 10.10%), 87.25% (sd: 7.37%) and 89.91% (sd: 6.57%) for the skew-symmetric perturbations of the Gaussian proposals from the Laplace method, black-box VB and EP, respectively. This means that replacing such Gaussian proposals with the corresponding skewed counterparts essentially doubles the ESS, at negligible computational costs. In fact, in importance sampling, the multiple likelihood evaluations we require to compute the skewing factor are also necessary under the Gaussian proposals to evaluate the importance weights. These findings strengthen the impact of our skew-symmetric approximation, and motivate further research on its use within importance sampling and its extensions.

5.2 Hierarchical semi-parametric logistic regression

Let us conclude by illustrating the accuracy gains of skew-symmetric approximations in a demographic application with $n = 30,524$ and $d = 62$. The data are retrieved from the HDS II survey whose focus, among others, is to study women contraceptive choices within 33 indian states. Recently, this dataset has been analyzed by [Rigon et al. \(2019\)](#) via a Bayesian semi-parametric sequential logistic regression to investigate how different demographic factors affect the probability of four different contraceptive choices (including no contraceptive). In the following, we focus instead on the binary decision on whether to use or not contraceptives as a function of state of residence, age, education (factor with four levels), religion (factor with four levels) and area (urban or rural).

Let \mathbf{x}_{ik}^\top be the 8-dimensional vector comprising the binary encoding of the variables education, religion and area for the i th woman within state $k = 1, \dots, K$. Consistent with the hierarchical structure of the data, we assume that the binary decision $y_{ik} \in \{0, 1\}$ to use or not contraceptives is modeled, for each woman $i = 1, \dots, n_k$ in state $k = 1, \dots, K$, through a Bernoulli variable with probability ψ_{ik} defined as

$$\text{logit}(\psi_{ik}) = \alpha_k + g(\text{age}_{ik}) + \mathbf{x}_{ik}^\top \boldsymbol{\beta}, \quad \text{for every } i = 1, \dots, n_k, \quad \text{and } k = 1, \dots, K,$$

where n_k denotes the number of observations in state k , α_k is a state-specific intercept, $\boldsymbol{\beta} \in \mathbb{R}^8$ denotes a vector encoding the fixed effects, while $g(\cdot)$ is a flexible functional effect for the variable age. Following [Rigon et al. \(2019\)](#), $g(\cdot)$ is defined via a linear combination of B-spline functions $B_m(\cdot)$, $m = 1, \dots, M$, with coefficients vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_M)^\top \in \mathbb{R}^M$. Bayesian inference under this model proceeds via suitably defined priors for the $d = 62$ parameters in $\theta = (\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\gamma})$. In particular, we assume improper uniform priors for the entries in $\boldsymbol{\beta}$, and independent Student- t with 3 degrees of freedom for the K intercepts within $\boldsymbol{\alpha}$ and the M entries of $\boldsymbol{\Sigma}^{-1/2} \boldsymbol{\gamma}$, where $\boldsymbol{\Sigma}$ is specified as in [Rigon et al. \(2019\)](#) to enforce similar coefficient values for contiguous splines. As discussed in [Section 4.1](#), the choice of Student- t coincides with the realistic assumption of Gaussian priors for the parameters combined with χ^2_3 hyperpriors for the corresponding precisions.

[Figure 1](#) extends the comparisons in [Section 5.1](#) by studying the accuracy in both relative and absolute terms of the skew-symmetric approximations over the unperturbed Gaussian counterparts from the Laplace method, black-box VB and EP, with focus on the marginals of the posterior induced by the above model. More specifically, the first three panels within [Figure 1](#) display the boxplots of the estimated $\mathcal{D}(\pi_{j,n} \parallel q_{j,n})$, $j = 1, \dots, 62$, where $q_{j,n}$ is the j th marginal

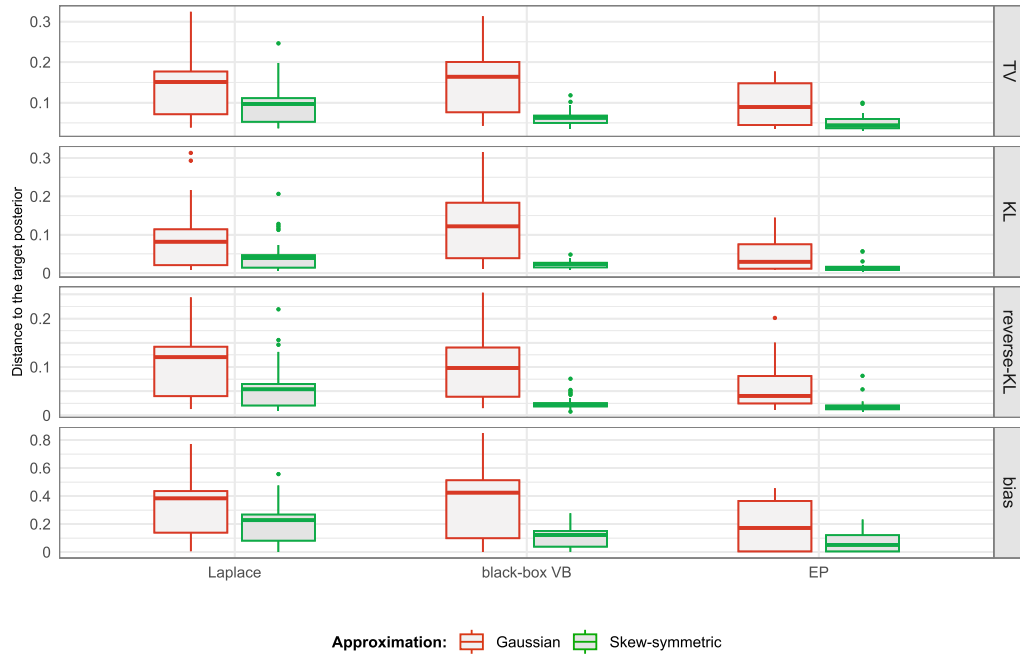


Figure 1. Empirical comparison of the accuracy achieved by three state-of-the-art Gaussian approximations from the Laplace method, black-box VB and EP, versus the corresponding skew-symmetric perturbations. For three routinely-employed divergences \mathcal{D} (TV, KL, reverse-KL) the first three panels display the boxplots of $\mathcal{D}(\pi_{j,n} \parallel q_{j,n})$, $j = 1, \dots, 62$, where $q_{j,n}$ is the j th marginal of either \bar{q}_{n,θ^*} (Gaussian) or q_{n,θ^*} (Skew-symmetric). The fourth panel shows instead the boxplot of the absolute differences between the approximated and actual posterior means of the $d = 62$ standardized parameters (standardization proceeds as discussed in the caption of Table 2).

of either \bar{q}_{n,θ^*} (Gaussian) or q_{n,θ^*} (Skew-symmetric), and \mathcal{D} corresponds to the TV, KL and reverse-KL divergences, respectively. The fourth panel shows instead the boxplot of the absolute differences between the approximated and actual posterior means of the $d = 62$ standardized parameters. Similarly to the empirical results in Sections 4.1–4.2 and 5.1, correcting for asymmetry yields also in this case approximations that systematically outperform the symmetric counterparts on all accuracy measures. Comparing the medians of the boxplots in Figure 1, these improvements are, generally, by a multiplicative factor of 2 or more, and provide evident gains also in absolute terms. For instance, since the bias is computed on standardized parameters, achieving systematic reductions of 0.2 or more via simple perturbations is a remarkable result in absolute terms.

Notice also that, among the deterministic approximations in Figure 1, the skewed perturbation of EP provides the most competitive strategy in terms of overall accuracy across the different summaries. Recalling, e.g. Vehtari et al. (2020), EP already provides highly-accurate deterministic approximations in its standard Gaussian form. As a result, perturbing these symmetric densities through the proposed skewing factor yields even more accurate skewed approximations. This point is also confirmed by the analysis of the relative improvements in the ESS when leveraging the skew-symmetric approximations instead of the Gaussian counterparts as proposals in importance sampling. In this case, the percentage gains in the ESS, averaged across 100 replicated studies, are 76.88% (sd: 9.25%), 59.88% (sd: 6.90%) and 87.59% (sd: 7.17%) for the skewed perturbations of the Gaussian proposals from the Laplace method, black-box VB and EP, respectively.

Notably, the above improvements are obtained at negligible computational costs. In this application with $n = 30,524$ and $d = 62$, the runtime of Monte Carlo inference via 10,000 samples from Algorithm 1 is approximately 8 seconds on a standard laptop, without considering any parallel implementations.

6 Discussion

This article introduces a novel and general strategy to perturb any given symmetric approximation of a generic posterior density for obtaining an improved, yet similarly tractable, skew-symmetric counterpart. Such an approximation is shown to improve, both theoretically and practically, the accuracy of the symmetric density which is perturbed. Unlike recently developed deterministic approximations based on generalizations of skew-normal distributions, the proposed solution (i) applies to generic posterior densities and to any symmetric approximation of such densities, from e.g. Laplace, VB and EP, (ii) does not imply additional optimization costs relative to those required for obtaining the original symmetric approximation to be improved, (iii) is substantially simpler and can be applied directly to any output of state-of-the-art softwares yielding symmetric approximations of posterior densities, (iv) has strong theoretical guarantees in terms of accuracy, and (v) achieves systematic empirical gains over the unperturbed symmetric counterparts in both simulation studies and real-data applications. Such advantages, combined with the limited computational burden of inference under these skew-symmetric approximations (see Sections 2.3–2.4 and 5.2), motivate extensive use of this novel solution, especially in situations where the posterior density displays non-negligible skewness and the symmetric approximation to be perturbed provides an accurate characterization of the symmetrized posterior.

Recalling Section 2.3, our proposal arises from a yet unexplored skew-symmetric representation of posterior densities in Proposition 2 that is of independent interest and stimulates several directions of future research. A promising one is discussed in Remark 4 and refers to the case in which also the symmetric approximating density $\bar{q}_{n,\theta^*}(\theta)$ is unknown and part of the optimization problem formalized in Corollary 1. This extension, combined with our results in Section 2, supports a change of perspective that suggests to focus on symmetrized posterior densities defined in (2), rather than on the original posterior, as the target of symmetric approximations. To the best of our knowledge, such a perspective has not been considered so far, thus stimulating active research motivated by the novel questions associated with this task. For instance, this would require extracting a suitable symmetric component from the target posterior density, which is both as close as possible to such a posterior and can be also accurately approximated by a tractable symmetric density whose perturbation yields the final skew-symmetric approximation of $\pi_n(\theta)$. A related problem can be found in classical, yet overlooked, literature aimed at improving estimators of empirical distribution functions for symmetric densities (Hinkley, 1976; Lo, 1985; Schuster, 1975, 1987). Inheriting these results within our framework for deterministic approximations provides a promising direction to address the objectives discussed in Remark 4.

The above perspective is also useful to further extend the theory in Section 3.2 on the asymptotic accuracy of skew-symmetric approximations. In particular, as a consequence of the proof of Theorem 3, a natural direction is to show that these solutions can be made arbitrarily accurate via higher order approximations only for the symmetrized component. From a more practical perspective, another key direction is to illustrate the empirical performance of the proposed skew-symmetric perturbation when applied to other symmetric approximations and in more complex hierarchical formulations beyond those considered in Sections 4–5. Notice that our focus in Sections 4–5 on Gaussian approximations arising from commonly adopted Laplace, black-box VB and EP methods is motivated by the attempt to stimulate implementation of the proposed skew-symmetric approximation in routinely used softwares. Nonetheless, showcasing the empirical gains of our proposal within the broader context of symmetric approximations discussed in Section 2.1 (e.g. ALA, INLA, t -exponential approximating families and delta-method variational inference), is a natural direction of future research which could further boost the impact of our contribution. Recalling Section 2.1, it is also of interest to explore the benefits of the proposed skew-symmetric approximation not only as a direct perturbation applied to the symmetric outputs of such methods, but also as a building-block within the corresponding optimization routines. For instance, as shown by a recent contribution (Kock et al., 2025) appeared after our article and leveraging our results, the skew-symmetric approximation we derived could be incorporated directly within the algorithms of stochastic black-box extensions of VB (Ranganath et al., 2014), including ADVI (Kucukelbir et al., 2017), without affecting the tractability of the stochastic gradient step, while enlarging (and hence improving) the class of approximating distributions. This is because the density of the skew-symmetric approximation is available in closed form, and i.i.d.

samples from it can be obtained via a direct perturbation of those from the symmetric counterparts routinely used in black-box implementations of VB. Such a tractability could also motivate the use of the proposed skewed perturbation to improve the initial Gaussian densities often employed in variational inference with normalizing flows (e.g. [Rezende & Mohamed, 2015](#)), while stimulating research to assess the benefits of this idea. Finally, the remarkable improvements in the effective sample size achieved by the skew-symmetric solution in Section 5, motivate its future use beyond the context of deterministic approximations, with a particular focus on importance sampling and its extensions ([Chopin & Papaspiliopoulos, 2020](#)).

Recalling the discussions and results in Sections 2.3, 2.4, and 5.2, the computational burden associated with the proposed skew-symmetric approximation is not a major source of concern in practice, and can be further reduced through parallel implementations both across the i.i.d. samples and likelihood factors. Nonetheless, future research exploring ideas to further reduce the cost of the likelihood evaluations (both in the sample size n and dimension d) is of interest. Data subsampling could be a direction, but naive implementations of this idea would not be compatible with the skew-symmetric representation of the posterior induced by the entire sample, and hence, developing a principled solution requires a careful treatment.

Acknowledgments

The authors are grateful to the Editor, the Associate Editor, and the referees for the constructive suggestions, which helped in improving the initial version of this article. In addition, the authors would like to thank Giacomo Zanella for the useful feedbacks on the ideas and methods developed in this article.

Conflicts of interest: None declared.

Funding

F.P. is funded by the European Union (ERC, PrSc-HDBayLe, n. 101076564). B.S. is co-funded by the European Union (ERC, BigBayesUQ, n. 101041064). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. D.D. acknowledges the support by the MUR-PRIN 2022 project ‘CARONTE’ (2022KBTEBN, funded by the European Union – Next Generation EU, Mission 4, CUP: J53D23009400001).

Data availability

Data and codes are available at https://github.com/Francesco16p/SkewAppr_Post.

Supplementary material

Supplementary material is available online at [Journal of the Royal Statistical Society: Series B](#).

References

- Anceschi N., Fasano A., Durante D., & Zanella G. (2023). Bayesian conjugacy in probit, tobit, multinomial probit and extensions: A review and new results. *Journal of the American Statistical Association*, 118(542), 1451–1469. <https://doi.org/10.1080/01621459.2023.2169150>
- Arellano-Valle R., & Azzalini A. (2006). On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics*, 33(3), 561–574. <https://doi.org/10.1111/j.1467-9469.2006.00503.x>
- Azzalini A., & Capitanio A. (1999). Statistical applications of the multivariate skew normal distribution. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, 61(3), 579–602. <https://doi.org/10.1111/1467-9868.00194>
- Azzalini A., & Capitanio A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, 65(2), 367–389. <https://doi.org/10.1111/1467-9868.00391>
- Azzalini A., & Capitanio A. (2013). *The skew-normal and related families*. Cambridge University Press.
- Azzalini A., & Dalla Valle A. (1996). The multivariate skew-normal distribution. *Biometrika*, 83(4), 715–726. <https://doi.org/10.1093/biomet/83.4.715>

- Bishop C. (2006). *Pattern recognition and machine learning*. Springer.
- Blei D. M., Kucukelbir A., & McAuliffe J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518), 859–877. <https://doi.org/10.1080/01621459.2017.1285773>
- Carbonetto P., & Stephens M. (2012). Scalable variational inference for Bayesian variable selection in regression, and its accuracy in genetic association studies. *Bayesian Analysis*, 7(1), 73–108. <https://doi.org/10.1214/12-BA703>
- Challis E., & Barber D. (2012). Affine independent variational inference. *Advances in Neural Information Processing Systems*, 25, 1–9.
- Challis E., & Barber D. (2013). Gaussian Kullback–Leibler approximate inference. *Journal of Machine Learning Research*, 14(68), 2239–2286.
- Chopin N., & Papaspiliopoulos O. (2020). *An introduction to sequential Monte Carlo*. Springer.
- Chopin N., & Ridgway J. (2017). Leave Pima Indians alone: Binary regression as a benchmark for Bayesian computation. *Statistical Science*, 32(1), 64–87. <https://doi.org/10.1214/16-STS581>
- Cichocki A., & Amari S.-I. (2010). Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6), 1532–1568. <https://doi.org/10.3390/e12061532>
- Consonni G., & Marin J.-M. (2007). Mean-field variational approximate Bayesian inference for latent variable models. *Computational Statistics & Data Analysis*, 52(2), 790–798. <https://doi.org/10.1016/j.csda.2006.10.028>
- Daudel K., Douc R., & Roueff F. (2023). Monotonic alpha-divergence minimisation for variational inference. *Journal of Machine Learning Research*, 24(64), 1–76.
- Dehaene G., & Barthelmé S. (2018). Expectation propagation in the large data limit. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, 80(1), 199–217. <https://doi.org/10.1111/rssb.12241>
- Ding N., Qi Y., & Vishwanathan S. (2011). t-divergence based approximate inference. *Advances in Neural Information Processing Systems*, 24, 1–9.
- Durante D., Pozza F., & Szabo B. (2024). Skewed Bernstein-von Mises theorem and skew-modal approximations. *Annals of Statistics*, 52(6), 2714–2737. <https://doi.org/10.1214/24-AOS2429>
- Durante D., & Rigon T. (2019). Conditionally conjugate mean-field variational Bayes for logistic models. *Statistical Science*, 34(3), 472–485. <https://doi.org/10.1214/19-STS712>
- Fasano A., & Durante D. (2022). A class of conjugate priors for multinomial probit models which includes the multivariate normal one. *Journal of Machine Learning Research*, 23(30), 1–26.
- Fasano A., Durante D., & Zanella G. (2022). Scalable and accurate variational Bayes for high-dimensional binary regression models. *Biometrika*, 109(4), 901–919. <https://doi.org/10.1093/biomet/asac026>
- Futami F., Sato I., & Sugiyama M. (2017). Expectation propagation for t-exponential family using q-algebra. *Advances in Neural Information Processing Systems*, 30, 1–10.
- Gallant A. R., & Nychka D. W. (1987). Semi-nonparametric maximum likelihood estimation. *Econometrica: Journal of the Econometric Society*, 55(2), 363–390. <https://doi.org/10.2307/1913241>
- Gelman A., Carlin J. B., Stern H. S., Dunson D. B., Vehtari A., & Rubin D. B. (2013). *Bayesian data analysis* (3rd ed.). Chapman and Hall/CRC.
- Genton M. G., & Loperfido N. M. (2005). Generalized skew-elliptical distributions and their quadratic forms. *Annals of the Institute of Statistical Mathematics*, 57(2), 389–401. <https://doi.org/10.1007/BF02507031>
- Giordano R., Broderick T., & Jordan M. I. (2018). Covariances, robustness and variational Bayes. *Journal of Machine Learning Research*, 19(51), 1–49.
- González-Farías G., Dominguez-Molina A., & Gupta A. K. (2004). Additive properties of skew normal random vectors. *Journal of Statistical Planning and Inference*, 126(2), 521–534. <https://doi.org/10.1016/j.jspi.2003.09.008>
- Hernandez-Lobato J., Li Y., Rowland M., Bui T., Hernández-Lobato D., & Turner R. (2016). Black-box α -divergence minimization. In *Proceedings of the 33rd international conference on machine learning* (Vol. 48, pp. 1511–1520). PMLR.
- Hinkley D. (1976). On estimating a symmetric distribution. *Biometrika*, 63(3), 680–681. <https://doi.org/10.1093/biomet/63.3.680>
- Kasprzak M. J., Giordano R., & Broderick T. (2025). How good is your Gaussian approximation of the posterior? Finite-sample computable error bounds for a variety of useful divergences. *Journal of Machine Learning Research*, 26(87), 1–81.
- Katsevich A. (2024). ‘The Laplace approximation accuracy in high dimensions: a refined analysis and new skew adjustment’, arXiv, arXiv:2306.07262, preprint: not peer reviewed. <https://doi.org/10.48550/arXiv.2306.07262>.
- Katsevich A., & Rigollet P. (2024). On the approximation accuracy of Gaussian variational inference. *Annals of Statistics*, 52(4), 1384–1409. <https://doi.org/10.1214/24-AOS2393>
- Kock L., Tan L. S., Bansal P., & Nott D. J. (2025). ‘Variational inference for hierarchical models with conditional scale and skewness corrections’, arXiv, arXiv:2503.18075, preprint: not peer reviewed. <https://doi.org/10.48550/arXiv.2503.18075>

- Kucukelbir A., Tran D., Ranganath R., Gelman A., & Blei D. M. (2017). Automatic differentiation variational inference. *Journal of Machine Learning Research*, 18(14), 1–45.
- Kullback S., & Leibler R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86. <https://doi.org/10.1214/aoms/1177729694>
- Lawal B. H. (2012). Zero-inflated count regression models with applications to some examples. *Quality & Quantity*, 46(1), 19–38. <https://doi.org/10.1007/s11135-010-9324-x>
- Liang F., Mahoney M., & Hodgkinson L. (2022). Fat-tailed variational inference with anisotropic tail adaptive flows. In *Proceedings of the 39th international conference on machine learning* (Vol. 162, pp. 13257–13270). PMLR.
- Lo S.-H. (1985). Estimation of a symmetric distribution. *Annals of Statistics*, 13(3), 1097–1113. <https://doi.org/10.1214/aos/1176349658>
- Ma Y., & Genton M. G. (2004). Flexible class of skew-symmetric distributions. *Scandinavian Journal of Statistics*, 31(3), 459–468. <https://doi.org/10.1111/sjos.2004.31.issue-3>
- Margossian C. C., Pillaud-Vivien L., & Saul L. K. (2025). Variational inference for uncertainty quantification: An analysis of trade-offs. *Journal of Machine Learning Research*, 26(202), 1–41.
- Minka T. P. (2001). Expectation propagation for approximate Bayesian inference. In *Proceedings of the 17th conference on uncertainty in artificial intelligence* (Vol. 17, pp. 362–369). Morgan Kaufmann Publishers.
- Opper M., & Archambeau C. (2009). The variational Gaussian approximation revisited. *Neural Computation*, 21(3), 786–792. <https://doi.org/10.1162/neco.2008.08-07-592>
- Ormerod J. T. (2011). *Skew-normal variational approximations for Bayesian inference* (Technical Report). School of Mathematics and Statistics, University of Sydney.
- Panov M., & Spokoiny V. (2015). Finite sample Bernstein–von Mises theorem for semiparametric problems. *Bayesian Analysis*, 10(3), 665–710. <https://doi.org/10.1214/14-BA926>
- Papamakarios G., Nalisnick E., Rezende D. J., Mohamed S., & Lakshminarayanan B. (2021). Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57), 1–64.
- Póczos B., & Schneider J. (2011). On the estimation of α -divergences. In *Proceedings of the 14th international conference on artificial intelligence and statistics* (Vol. 15, pp. 609–617). PMLR.
- Ranganath R., Gerrish S., & Blei D. (2014). Black box variational inference. In *Proceedings of the 17th international conference on artificial intelligence and statistics* (Vol. 33, pp. 814–822). PMLR.
- Ray K., & Szabó B. (2022). Variational Bayes for high-dimensional linear regression with sparse priors. *Journal of the American Statistical Association*, 117(539), 1270–1281. <https://doi.org/10.1080/01621459.2020.1847121>
- Rezende D., & Mohamed S. (2015). Variational inference with normalizing flows. In *Proceedings of the 32nd international conference on machine learning* (Vol. 37, pp. 1530–1538). PMLR.
- Rigon T., Durante D., & Torelli N. (2019). Bayesian semiparametric modelling of contraceptive behaviour in India via sequential logistic regressions. *Journal of the Royal Statistical Society: Series A, Statistics in Society*, 182(1), 225–247. <https://doi.org/10.1111/rssa.12361>
- Rossell D., Abril O., & Bhattacharya A. (2021). Approximate Laplace approximations for scalable model selection. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, 83(4), 853–879. <https://doi.org/10.1111/rssb.12466>
- Rue H., Martino S., & Chopin N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, 71(2), 319–392. <https://doi.org/10.1111/j.1467-9868.2008.00700.x>
- Ruli E., Sartori N., & Ventura L. (2016). Improved Laplace approximation for marginal likelihoods. *Electronic Journal of Statistics*, 10(2), 3986–4009. <https://doi.org/10.1214/16-EJS1218>
- Salomone R., Yu X., Nott D. J., & Kohn R. (2024). Structured variational approximations with skew normal decomposable graphical models. *Journal of Computational and Graphical Statistics*, 33(4), 1329–1338. <https://doi.org/10.1080/10618600.2024.2319159>
- Schuster E. F. (1975). Estimating the distribution function of a symmetric distribution. *Biometrika*, 62(3), 631–635. <https://doi.org/10.1093/biomet/62.3.631>
- Schuster E. F. (1987). Identifying the closest symmetric distribution or density function. *Annals of Statistics*, 15(2), 865–874. <https://doi.org/10.1214/aos/1176350380>
- Smith M. S., Loaiza-Maya R., & Nott D. J. (2020). High-dimensional copula variational approximation through transformation. *Journal of Computational and Graphical Statistics*, 29(4), 729–743. <https://doi.org/10.1080/10618600.2020.1740097>
- Spokoiny V. (2025). Inexact Laplace approximation and the use of posterior mean in Bayesian inference. *Bayesian Analysis*, 20(1), 1303–1330. <https://doi.org/10.1214/23-BA1391>
- Spokoiny V., & Panov M. (2025). Accuracy of Gaussian approximation for high-dimensional posterior distributions. *Bernoulli*, 31(2), 843–867. <https://doi.org/10.3150/21-BEJ1412>
- Tan L. S., & Chen A. (2024). Variational inference based on a subclass of closed skew normals. *Journal of Computational and Graphical Statistics*, 34(2), 422–436. <https://doi.org/10.1080/10618600.2024.2402278>

- Tan L. S., & Nott D. J. (2018). Gaussian variational approximation with sparse precision matrices. *Statistics and Computing*, 28(2), 259–275. <https://doi.org/10.1007/s11222-017-9729-7>
- Tierney L., & Kadane J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393), 82–86. <https://doi.org/10.1080/01621459.1986.10478240>
- Van der Vaart A. W. (2000). *Asymptotic statistics*. Cambridge University Press.
- Vehtari A., Gelman A., Sivula T., Jylänki P., Tran D., Sahai S., Blomstedt P., Cunningham J. P., Schiminovich D., & Robert C. P. (2020). Expectation propagation as a way of life: A framework for Bayesian inference on partitioned data. *Journal of Machine Learning Research*, 21(17), 1–53.
- Walker A. M. (1969). On the asymptotic behaviour of posterior distributions. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, 31(1), 80–88. <https://doi.org/10.1111/j.2517-6161.1969.tb00767.x>
- Wang C., & Blei D. M. (2013). Variational inference in nonconjugate models. *Journal of Machine Learning Research*, 14(32), 1005–1031.
- Wang J., Boyer J., & Genton M. G. (2004). A skew-symmetric representation of multivariate distributions. *Statistica Sinica*, 14(4), 1259–1270. <https://www.jstor.org/stable/24307231>
- Wang Y., & Blei D. M. (2019). Frequentist consistency of variational Bayes. *Journal of the American Statistical Association*, 114(527), 1147–1161. <https://doi.org/10.1080/01621459.2018.1473776>
- Yang Y., Pati D., & Bhattacharya A. (2020). α -variational inference with statistical guarantees. *Annals of Statistics*, 48(2), 886–905. <https://doi.org/10.1214/19-AOS1827>
- Zhou J., Grazian C., & Ormerod J. T. (2024). Tractable skew-normal approximations via matching. *Journal of Statistical Computation and Simulation*, 94(5), 1016–1034. <https://doi.org/10.1080/00949655.2023.2277885>