

Expectation propagation for the smoothing distribution in dynamic probit

Niccolò Anceschi¹, Augusto Fasano², and Giovanni Rebaudo^{2,3}

¹ Duke University, Durham, USA. niccolo.anceschi@duke.edu

² Collegio Carlo Alberto, Turin, IT. augusto.fasano@carloalberto.org

³ University of Turin, Turin, IT. giovanni.rebaudo@unito.it

Abstract. The smoothing distribution of dynamic probit models with Gaussian state dynamics was recently proved to belong to the unified skew-normal family. Although this is computationally tractable in small-to-moderate settings, it may become computationally impractical in higher dimensions. In this work, adapting a recent more general class of expectation propagation (EP) algorithms, we derive an efficient EP routine to perform inference for such a distribution. We show that the proposed approximation leads to accuracy gains over available approximate algorithms in a financial illustration.

Keywords: Dynamic Probit Model, State-Space Model, Expectation Propagation, Unified Skew-Normal Distribution, Smoothing

1 Introduction

Dynamic binary regression is a lively field of research, with the dynamic probit model playing a central role [1,2,3,4]. In this work, we develop an expectation propagation (EP) approximation of the joint smoothing distribution in a dynamic probit model. In such a state-space model, for each time t , a known p -dimensional covariate vector \mathbf{x}_t is available, while the binary observations $y_t \in \{0, 1\}$, $t = 1, \dots, n$, compose a time series whose dependence across time is driven by a dynamic latent state $\boldsymbol{\theta}_t = (\theta_{1t}, \dots, \theta_{pt})^\top \in \mathbb{R}^p$. The latter has Markovian dynamics and affects the probability of success of the observation at the corresponding time according to the following specifications:

$$y_t \mid \boldsymbol{\theta}_t \sim \text{BERN}(\Phi(\mathbf{x}_t^\top \boldsymbol{\theta}_t)), \quad (1)$$

$$\boldsymbol{\theta}_t = \mathbf{G}_t \boldsymbol{\theta}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \stackrel{\text{ind}}{\sim} \text{N}_p(\mathbf{0}, \mathbf{W}_t), \quad t = 1, \dots, n, \quad (2)$$

with $\boldsymbol{\theta}_0 \sim \text{N}_p(\mathbf{a}_0, \mathbf{P}_0) \perp \{\boldsymbol{\varepsilon}_t\}_{t \geq 1}$. We denote with $\Phi(\cdot)$ the cumulative distribution function of the standard normal distribution, while \mathbf{G}_t , \mathbf{W}_t and \mathbf{P}_0 are known matrices. In the following, according to common practices, we set $\mathbf{a}_0 = \mathbf{0}$. Moreover, we will denote with $\Phi_m(\mathbf{z}; \boldsymbol{\Sigma})$ the cumulative distribution of an m -variate Gaussian random variable having zero mean and covariance matrix $\boldsymbol{\Sigma}$, evaluated at \mathbf{z} . Calling $\mathbf{y}_{1:n} = (y_1, \dots, y_n)^\top$ and $\boldsymbol{\theta}_{1:n} = (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_n^\top)^\top$, the smoothing distribution is given by $p(\boldsymbol{\theta}_{1:n} \mid \mathbf{y}_{1:n})$, which has been showed (see

[2]) to belong to the class of unified skew-normal (SUN) distributions [5]. This result comes however with some computational difficulties when the length of the time series is moderately large, for the reasons clarified in Section 2, motivating the development of approximate methods [3]. This can be done by resorting to expectation propagation, which constructs a Gaussian approximation of the true posterior mirroring its factorization structure [6]. To this end, it is worth noting that the model (1)-(2) can be rewritten as a probit model with parameter $\boldsymbol{\theta}_{1:n} \in \mathbb{R}^{pn}$ and fictional pn -dimensional covariates $\tilde{\mathbf{x}}_t = (\mathbf{0}_{p(t-1)}^\top, \mathbf{x}_t^\top, \mathbf{0}_{p(n-t)}^\top)^\top$, with $\mathbf{0}_m$ denoting the m -dimensional column vector of zeros. Indeed, model (1)-(2) is equivalent to

$$p(y_t | \boldsymbol{\theta}_{1:n}) = \Phi((2y_t - 1)\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_{1:n}), \quad (3)$$

$$\boldsymbol{\theta}_{1:n} \sim N_{pn}(\mathbf{0}, \boldsymbol{\Omega}), \quad (4)$$

where $\boldsymbol{\Omega}$ is composed by $(p \times p)$ -dimensional blocks whose expression follows directly from the dynamics in (2): defining $\mathbf{G}_l^t = \mathbf{G}_t \cdots \mathbf{G}_l$, $l \leq t-1$, it holds $\boldsymbol{\Omega}_{[tt]} = \text{var}(\boldsymbol{\theta}_t) = \mathbf{G}_1^t \mathbf{P}_0 \mathbf{G}_1^{t\top} + \sum_{l=2}^t \mathbf{G}_l^t \mathbf{W}_{l-1} \mathbf{G}_l^{t\top} + \mathbf{W}_t$, for $t = 1, \dots, n$, and $\boldsymbol{\Omega}_{[tl]} = \boldsymbol{\Omega}_{[tl]}^\top = \text{cov}(\boldsymbol{\theta}_t, \boldsymbol{\theta}_l) = \mathbf{G}_{l+1}^t \boldsymbol{\Omega}_{[ll]}$, for $t > l$ (see also [3]). From this,

$$p(\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}) \propto p(\boldsymbol{\theta}_{1:n}) \prod_{t=1}^n \Phi((2y_t - 1)\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_{1:n}). \quad (5)$$

Leveraging on this representation, the EP scheme constructs an approximating density of (5) by replacing each factor therein with a Gaussian density, whose parameters are optimized iteratively via moments matching between the global approximant and a so-called hybrid distribution, as detailed in Section 3. The resulting global EP approximation, which is Gaussian by construction, is shown to be accurate and tractable in Section 4. More details on the EP algorithm can be found, for instance, in Chapter 10 of [7].

2 Literature review

Exploiting (5), [2] extended the conjugacy results in the static probit regression [8] to the more challenging multivariate dynamic probit setting. In particular, [2] showed that the filtering, predictive, and smoothing densities of the state variables in this setting have unified skew-normal (SUN) [5] kernels and compute the closed form of the parameters of such distributions. To provide a brief overview, a random vector $\boldsymbol{\theta} \in \mathbb{R}^q$ is said to follow a SUN distribution, denoted as $\boldsymbol{\theta} \sim \text{SUN}_{q,h}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$, if its density function can be expressed as:

$$\phi_q(\boldsymbol{\theta} - \boldsymbol{\xi}; \boldsymbol{\Omega}) \frac{\Phi_h(\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\omega}^{-1}(\boldsymbol{\theta} - \boldsymbol{\xi}); \boldsymbol{\Gamma} - \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Delta})}{\Phi_h(\boldsymbol{\gamma}; \boldsymbol{\Gamma})},$$

where $\phi_q(\boldsymbol{\theta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$ is the density function of a Gaussian distribution with $\mathbf{0}$ mean and covariance matrix $\boldsymbol{\Omega}$ evaluated at $\boldsymbol{\theta} - \boldsymbol{\xi}$, $\boldsymbol{\omega} = (\boldsymbol{\Omega} \odot \mathbf{I}_q)^{1/2}$ is the diagonal scale matrix, with \odot denoting the element-wise Hadamard product, and $\bar{\boldsymbol{\Omega}} = \boldsymbol{\omega}^{-1} \boldsymbol{\Omega} \boldsymbol{\omega}^{-1}$ is the corresponding correlation matrix. Additional details on

SUN distributions can be found in [5]. More specifically, taking $\boldsymbol{\Omega}$ as in Section 1, under model (1)–(2) the joint smoothing distribution has the form (see [2,3])

$$(\boldsymbol{\theta}_{1:n} \mid \mathbf{y}_{1:n}) \sim \text{SUN}_{pn,n}(\mathbf{0}, \boldsymbol{\Omega}_{1:n|n}, \boldsymbol{\Delta}_{1:n|n}, \mathbf{0}, \boldsymbol{\Gamma}_{1:n|n}), \quad (6)$$

with $\boldsymbol{\Omega}_{1:n|n} = \boldsymbol{\Omega}$, $\boldsymbol{\Delta}_{1:n|n} = \bar{\boldsymbol{\Omega}}\boldsymbol{\omega}\mathbf{D}^\top\mathbf{s}^{-1}$, $\boldsymbol{\Gamma}_{1:n|n} = \mathbf{s}^{-1}(\mathbf{D}\boldsymbol{\Omega}\mathbf{D}^\top + \mathbf{I}_n)\mathbf{s}^{-1}$, where \mathbf{D} is an $n \times pn$ block-diagonal matrix having block entries $\mathbf{D}_{[tt]} = (2y_t - 1)\mathbf{x}_t^\top$, $t = 1, \dots, n$, $\mathbf{s} = [(\mathbf{D}\boldsymbol{\Omega}\mathbf{D}^\top + \mathbf{I}_n) \odot \mathbf{I}_n]^{1/2}$, \mathbf{I}_n defines the n -dimensional identity matrix. Leveraging on (6) and the additive representation of the SUN [5], one can establish a probabilistic characterization that facilitates the generation of independent and identically distributed (i.i.d.) samples from the smoothing distribution. The characterization is as follows:

$$(\boldsymbol{\theta}_{1:n} \mid \mathbf{y}_{1:n}) \stackrel{d}{=} \boldsymbol{\omega}_{1:n|n}(\mathbf{U}_{0 \ 1:n|n} + \boldsymbol{\Delta}_{1:n|n}\boldsymbol{\Gamma}_{1:n|n}^{-1}\mathbf{U}_{1 \ 1:n|n}),$$

where $\mathbf{U}_{0 \ 1:n|n} \sim N_{pn}(\mathbf{0}, \bar{\boldsymbol{\Omega}}_{1:n|n} - \boldsymbol{\Delta}_{1:n|n}\boldsymbol{\Gamma}_{1:n|n}^{-1}\boldsymbol{\Delta}_{1:n|n}^\top)$, while $\mathbf{U}_{1 \ 1:n|n}$ is distributed according to a truncated $N_n(\mathbf{0}, \boldsymbol{\Gamma}_{1:n|n})$ with lower truncation at $\mathbf{0}$. Based on this representation, one can develop an i.i.d. sampler where the most computationally intensive task is sampling from an n -variate truncated Gaussian distribution. However, despite recent advances [9] enabling efficient simulation for small-to-moderate time series (i.e., with n in the order of a few hundred), this i.i.d. sampler may become impractical for longer time series due to such computational constraints arising from the multivariate truncated Gaussian component.

To overcome this issue, [3] developed a partially-factorized variational Bayes (PFM-VB) approximation of the smoothing distribution where, adapting [10] to the dynamic setting, the multivariate truncated Gaussian component is replaced by n independent univariate truncated Gaussian terms. This leads to remarkable computational gains and great accuracy. Here, motivated by the overall improved accuracy of EP [11,12], we propose an efficient and accurate EP algorithm for the smoothing distribution of model (1)–(2). Our approach adapts to the dynamic probit setting techniques developed by [11] and further specified to the static probit in [13].

3 Expectation propagation (EP) for the dynamic probit

As clarified in Section 2, dealing with the $\text{SUN}_{pn,n}$ smoothing distribution for model (1)–(2) may become computationally intractable in scenarios where n is not small-to-moderate. As anticipated, the computational issues arise from the fact that sampling from such distribution requires sampling from an n -variate truncated normal and similar problems are faced when computing moments. Consistently, computational methods which are able to provide a good approximation of the smoothing distribution at a much lower computational time may provide a preferable alternative. In line with this goal, we derive the expectation propagation (EP) routine [7,14] for the dynamic probit model (1)–(2), by specifying more general results obtained for a wide class of models in [11] and adapting computations done for the static probit model in [13]. Section 3.1 presents

how one can leverage results on multivariate extended skew-normal distributions (SN) to obtain the equations for the EP updates. The resulting approximation is computable at a much lower cost than sampling from the exact smoothing distribution. Nevertheless, it requires matrices manipulation of dimension $pn \times pn$ within each iteration, which may be computationally inefficient. We thus show in Section 3.2 how the EP routine can be implemented avoiding such direct matrix computations, by storing and updating only lower-dimensional parameters. This results in sensible computational gains.

3.1 Implementation without $pn \times pn$ matrix inversions

The goal of expectation propagation (EP) is to approximate the smoothing distribution $p(\boldsymbol{\theta}_{1:n} \mid \mathbf{y}_{1:n})$ with an approximating distribution having density $q(\boldsymbol{\theta}_{1:n}) \propto \prod_{t=0}^n q_t(\boldsymbol{\theta}_{1:n})$, whose factorization reflects the form of the smoothing distribution (5). Differently from (5), in order to have a tractable approximating distribution, all the q_t 's, $t = 0, \dots, n$, are Gaussian densities having form $q_t(\boldsymbol{\theta}_{1:n}) \propto \exp(-\frac{1}{2}\boldsymbol{\theta}_{1:n}^\top \mathbf{Q}_t \boldsymbol{\theta}_{1:n} + \boldsymbol{\theta}_{1:n}^\top \mathbf{r}_t)$ for $t = 1, \dots, n$. As a consequence, the resulting EP approximating density will be Gaussian too, with $q(\boldsymbol{\theta}_{1:n}) = \phi_p(\boldsymbol{\theta}_{1:n} - \mathbf{Q}^{-1}\mathbf{r}, \mathbf{Q}^{-1})$, having set $\mathbf{r} = \sum_{t=0}^n \mathbf{r}_t$, $\mathbf{Q} = \sum_{t=0}^n \mathbf{Q}_t$. Recalling the analogy between the factorization of $q(\boldsymbol{\theta}_{1:n})$ and (5), the parameters for the factor, or site, 0 are set to $\mathbf{r}_0 = \mathbf{0}$ and $\mathbf{Q}_0 = \boldsymbol{\Omega}^{-1}$ and kept constant throughout the algorithm, so that $q_0(\boldsymbol{\theta}_{1:n})$ matches the prior distribution. On the other hand, the parameters \mathbf{r}_t and \mathbf{Q}_t for each site $t = 1, \dots, n$ are iteratively updated so that, keeping the parameters for the other sites fixed, the first two moments of the global approximation $q(\boldsymbol{\theta}_{1:n})$ match the ones of the hybrid distribution

$$h_t(\boldsymbol{\theta}_{1:n}) \propto p(y_t \mid \boldsymbol{\theta}_{1:n}) \prod_{j \neq t} q_j(\boldsymbol{\theta}_{1:n}) = \Phi((2y_t - 1)\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_{1:n}) \prod_{j \neq t} q_j(\boldsymbol{\theta}_{1:n}). \quad (7)$$

The moments of $q(\boldsymbol{\theta}_{1:n})$ are immediate, being $q(\boldsymbol{\theta}_{1:n})$ Gaussian, while the ones of $h_t(\boldsymbol{\theta}_{1:n})$ are also straightforward to compute, after noticing that it coincides with the kernel of an $\text{SN}_p(\boldsymbol{\xi}_t, \boldsymbol{\Omega}_t, \boldsymbol{\alpha}_t, \tau_t)$ [15] with

$$\begin{aligned} \boldsymbol{\xi}_t &= \mathbf{Q}_{-t}^{-1} \mathbf{r}_{-t}, & \boldsymbol{\Omega}_t &= \mathbf{Q}_{-t}^{-1}, & \boldsymbol{\alpha}_t &= (2y_t - 1)\boldsymbol{\omega}_t \tilde{\mathbf{x}}_t, \\ \tau_t &= (2y_t - 1)(1 + \tilde{\mathbf{x}}_t^\top \boldsymbol{\Omega}_t \tilde{\mathbf{x}}_t)^{-1/2} \tilde{\mathbf{x}}_t^\top \boldsymbol{\xi}_t, \end{aligned}$$

with $\mathbf{Q}_{-t} = \sum_{j \neq t} \mathbf{Q}_j = \mathbf{Q} - \mathbf{Q}_t$, $\mathbf{r}_{-t} = \sum_{j \neq t} \mathbf{r}_j = \mathbf{r} - \mathbf{r}_t$ and $\boldsymbol{\omega}_t = (\boldsymbol{\Omega}_t \odot \mathbf{I}_{pn})^{1/2}$. The moments of such distribution are then given by Equations (5.71) and (5.72) in [15]. Calling $\zeta_1(x) = \phi(x)/\Phi(x)$ and $\zeta_2(x) = -\zeta_1(x)^2 - x\zeta_1(x)$, they equal

$$\begin{aligned} \boldsymbol{\mu}_{h_t} &= \mathbb{E}_{h_t(\boldsymbol{\theta}_{1:n})}[\boldsymbol{\theta}_{1:n}] = \boldsymbol{\xi}_t + \zeta_1(\tau_t) s_t \boldsymbol{\Omega}_t \tilde{\mathbf{x}}_t \\ \boldsymbol{\Sigma}_{h_t} &= \text{var}_{h_t(\boldsymbol{\theta}_{1:n})}[\boldsymbol{\theta}_{1:n}] = \boldsymbol{\Omega}_t + \zeta_2(\tau_t) s_t^2 \boldsymbol{\Omega}_t \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \boldsymbol{\Omega}_t, \end{aligned}$$

with $s_t = (2y_t - 1)(1 + \tilde{\mathbf{x}}_t^\top \boldsymbol{\Omega}_t \tilde{\mathbf{x}}_t)^{-1/2}$. In EP, at each iteration the parameters for \mathbf{r}_t and \mathbf{Q}_t of each site $t = 1, \dots, n$ are updated so that the EP moment matching

Algorithm 1: EP for dynamic probit - no $pn \times pn$ matrix inversions

Initialization: $\mathbf{Q}^{-1} = \mathbf{\Omega}$; $\mathbf{r} = \mathbf{0}$; $k_t = 0$ and $m_t = 0$ for $t = 1, \dots, n$.
for s from 1 until convergence **do**
 for t from 1 to n **do**
 $\mathbf{r}_{-t} = \mathbf{r} - m_t \tilde{\mathbf{x}}_t$
 $\mathbf{\Omega}_t = \mathbf{Q}^{-1} + k_t / (1 - k_t \tilde{\mathbf{x}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{x}}_t) \mathbf{Q}^{-1} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \mathbf{Q}^{-1}$
 $s_t = (2y_t - 1)(1 + \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t)^{-1/2}$
 $\tau_t = s_t \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \mathbf{r}_{-t}$
 $k_t = -\zeta_2(\tau_t) / (1 + \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t + \zeta_2(\tau_t) \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t)$
 $m_t = \zeta_1(\tau_t) s_t + k_t \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \mathbf{r}_{-t} + k_t \zeta_1(\tau_t) s_t \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t$
 $\mathbf{r} = \mathbf{r}_{-t} + m_t \tilde{\mathbf{x}}_t$
 $\mathbf{Q}^{-1} = \mathbf{\Omega}_t + \zeta_2(\tau_t) s_t^2 \mathbf{\Omega}_t \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t$
Output: $q(\boldsymbol{\theta}_{1:n}) = \phi_p(\boldsymbol{\theta}_{1:n} - \mathbf{Q}^{-1} \mathbf{r}; \mathbf{Q}^{-1})$

conditions

$$\begin{cases} (\mathbf{Q}_{-t} + \mathbf{Q}_t^{\text{NEW}})^{-1} (\mathbf{r}_{-t} + \mathbf{r}_t^{\text{NEW}}) = \boldsymbol{\mu}_{h_t} \\ (\mathbf{Q}_{-t} + \mathbf{Q}_t^{\text{NEW}})^{-1} = \boldsymbol{\Sigma}_{h_t}, \end{cases}$$

are satisfied by the updated parameters $\mathbf{r}_t^{\text{NEW}}$ and $\mathbf{Q}_t^{\text{NEW}}$. This translates into

$$\begin{cases} \mathbf{r}_t^{\text{NEW}} = (\mathbf{Q}_{-t} + \mathbf{Q}_t^{\text{NEW}}) \boldsymbol{\mu}_{h_t} - \mathbf{r}_{-t} \\ \mathbf{Q}_t^{\text{NEW}} = \boldsymbol{\Sigma}_{h_t}^{-1} - \mathbf{Q}_{-t}. \end{cases}$$

Adapting the derivations in [13], where results for a broader class of models reported in [11] are specified for the classical probit model, the application of Woodbury's identity to $\boldsymbol{\Sigma}_{h_t}^{-1}$ and some algebra lead to the equalities

$$\mathbf{Q}_t^{\text{NEW}} = k_t^{\text{NEW}} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top, \quad \mathbf{r}_t^{\text{NEW}} = m_t^{\text{NEW}} \tilde{\mathbf{x}}_t,$$

with $k_t^{\text{NEW}} = -\zeta_2(\tau_t) / (1 + \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t + \zeta_2(\tau_t) \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t)$ and $m_t^{\text{NEW}} = \zeta_1(\tau_t) s_t + k_t^{\text{NEW}} \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \mathbf{r}_{-t} + k_t^{\text{NEW}} \zeta_1(\tau_t) s_t \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t \tilde{\mathbf{x}}_t$. Thus, the updates of the multidimensional parameters \mathbf{r}_t and \mathbf{Q}_t are fully determined by the updates of the scalar quantities k_t and m_t , $t = 1, \dots, n$. These are usually initialized to zero so that the initial EP approximation $q(\boldsymbol{\theta}_{1:n})$ coincides with the prior distribution for $\boldsymbol{\theta}_{1:n}$. In order to be able to implement the updates of k_t and m_t , $t = 1, \dots, n$, one needs to compute $\mathbf{\Omega}_t = \mathbf{Q}_{-t}^{-1}$. This can be done by avoiding direct matrix inversions during the iterations, exploiting Woodbury's identity. Indeed, it holds

$$\mathbf{\Omega}_t = \mathbf{Q}_{-t}^{-1} = (\mathbf{Q} - k_t \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top)^{-1} = \mathbf{Q}^{-1} + \frac{k_t}{1 - k_t \tilde{\mathbf{x}}_t^\top \mathbf{Q}^{-1} \tilde{\mathbf{x}}_t} \mathbf{Q}^{-1} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \mathbf{Q}^{-1},$$

with \mathbf{Q}^{-1} computed explicitly in the precomputation step and then updated at each step after the site-specific parameters are computed, exploiting the equality $(\mathbf{Q}^{\text{NEW}})^{-1} = (\mathbf{Q}_{-t} + \mathbf{Q}_t^{\text{NEW}})^{-1} = \boldsymbol{\Sigma}_{h_t} = \mathbf{\Omega}_t + \zeta_2(\tau_t) s_t^2 \mathbf{\Omega}_t \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top \mathbf{\Omega}_t$. The results presented above form the EP implementation presented in Algorithm 1. Such

implementation does not involve direct matrix inversions across iterations. Nevertheless, it requires storing, updating and multiplying the $pn \times pn$ -dimensional matrices $\mathbf{\Omega}_t$, $t = 1, \dots, n$, and \mathbf{Q}^{-1} . This may result inefficient when either p or n (or both) are large. We then show in Section 3.2 how the EP routine can be implemented without storing and manipulating such $pn \times pn$ -dimensional matrices, but only lower dimensional ones.

3.2 Implementation without $pn \times pn$ matrix updates

As shown in Section 3.1, the parameters \mathbf{r}_t and \mathbf{Q}_t , $t = 1, \dots, n$, are fully characterized by the known vector $\tilde{\mathbf{x}}_t$ and the scalar quantities k_t and m_t , which are iteratively updated in the EP algorithm. A close inspection of Algorithm 1 shows that the quantities \mathbf{Q}^{-1} and $\mathbf{\Omega}_t$, $t = 1, \dots, n$, are not explicitly needed to perform such updates, but it is enough to store and update the pn -dimensional vectors $\mathbf{w}_t = \mathbf{\Omega}_t \tilde{\mathbf{x}}_t = \mathbf{Q}_t^{-1} \tilde{\mathbf{x}}_t$ and $\mathbf{v}_t = \mathbf{Q}^{-1} \tilde{\mathbf{x}}_t$, $t = 1, \dots, n$. Thus, in order to be able to implement the EP routine in terms of these quantities, their update rules need to be derived and, most importantly, they must be computable in an efficient way. We show here that they allow an efficient formulation, which may make the new implementation preferable, from a computational point of view, to Algorithm 1. Adapting the derivations in [13], applying Woodbury's identity to $\mathbf{Q}_t^{-1} = (\mathbf{Q} - \mathbf{Q}_t)^{-1}$, after some algebra one gets $\mathbf{w}_t = d_t \mathbf{v}_t$, with $d_t = (1 - k_t \tilde{\mathbf{x}}_t^\top \mathbf{v}_t)^{-1}$. Thus, when updating the parameters for site t , $t = 1, \dots, n$, \mathbf{w}_t can be computed immediately from \mathbf{v}_t . On the other hand, each time a site t is updated also the EP covariance matrix \mathbf{Q}^{-1} changes. Thus, at the end of the update of each site t , all the \mathbf{v}_j 's (not only \mathbf{v}_t) must be updated to reflect this

Algorithm 2: EP for dynamic probit - no $pn \times pn$ matrix inversions or updates

Initialization: $\mathbf{r} = \mathbf{0}$; $k_t = 0$ and $m_t = 0$ for $t = 1, \dots, n$; $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n] = \mathbf{\Omega} \tilde{\mathbf{X}}^\top$.

for s from 1 until convergence **do**

for t from 1 to n **do**

$$\mathbf{w}_t = (1 - k_t \tilde{\mathbf{x}}_t^\top \mathbf{v}_t)^{-1} \mathbf{v}_t$$

$$\mathbf{r}_{-t} = \mathbf{r} - m_t \tilde{\mathbf{x}}_t$$

$$s_t = (2y_t - 1)(1 + \tilde{\mathbf{x}}_t^\top \mathbf{w}_t)^{-1/2}$$

$$\tau_t = s_t \mathbf{w}_t^\top \mathbf{r}_{-t}$$

$$k_t^{\text{NEW}} = -\zeta_2(\tau_t) / (1 + \tilde{\mathbf{x}}_t^\top \mathbf{w}_t + \zeta_2(\tau_t) \tilde{\mathbf{x}}_t^\top \mathbf{w}_t)$$

$$m_t = \zeta_1(\tau_t) s_t + k_t^{\text{NEW}} \mathbf{w}_t^\top \mathbf{r}_{-t} + k_t^{\text{NEW}} \zeta_1(\tau_t) s_t \tilde{\mathbf{x}}_t^\top \mathbf{w}_t$$

$$k_t = k_t^{\text{NEW}}$$

$$\mathbf{r} = \mathbf{r}_{-t} + m_t \tilde{\mathbf{x}}_t$$

$$\mathbf{V} = \mathbf{V} - \mathbf{v}_t [(k_t^{\text{NEW}} - k_t) / (1 + (k_t^{\text{NEW}} - k_t) \tilde{\mathbf{x}}_t^\top \mathbf{v}_t)] \tilde{\mathbf{x}}_t^\top \mathbf{V}$$

$$\mathbf{Q}^{-1} = \mathbf{\Omega} - \mathbf{V} \mathbf{K} \tilde{\mathbf{X}} \mathbf{\Omega}$$

Output: $q(\boldsymbol{\theta}_{1:n}) = \phi_p(\boldsymbol{\theta}_{1:n} - \mathbf{Q}^{-1} \mathbf{r}; \mathbf{Q}^{-1})$

change. After some algebra, exploiting again Woodbury's identity, one gets

$$\begin{aligned}\mathbf{v}_j^{\text{NEW}} &= (\mathbf{Q}^{\text{NEW}})^{-1}\tilde{\mathbf{x}}_j = (\mathbf{Q} - \mathbf{Q}_t + \mathbf{Q}_t^{\text{NEW}})^{-1}\tilde{\mathbf{x}}_j = [\mathbf{Q} + (k_t^{\text{NEW}} - k_t)\tilde{\mathbf{x}}_t\tilde{\mathbf{x}}_t^\top]^{-1}\tilde{\mathbf{x}}_j \\ &= [\mathbf{Q}^{-1} - (k_t^{\text{NEW}} - k_t)[1 + (k_t^{\text{NEW}} - k_t)\tilde{\mathbf{x}}_t^\top\mathbf{Q}^{-1}\tilde{\mathbf{x}}_t]^{-1}\mathbf{Q}^{-1}\tilde{\mathbf{x}}_t\tilde{\mathbf{x}}_t^\top\mathbf{Q}^{-1}]\tilde{\mathbf{x}}_j \\ &= \mathbf{Q}^{-1}\tilde{\mathbf{x}}_j - [(k_t^{\text{NEW}} - k_t)^{-1} + \tilde{\mathbf{x}}_t^\top\mathbf{v}_t]^{-1}\mathbf{v}_t\tilde{\mathbf{x}}_t^\top\mathbf{v}_j = \mathbf{v}_j - c_t(\tilde{\mathbf{x}}_t^\top\mathbf{v}_j)\mathbf{v}_t,\end{aligned}$$

where $c_t = (k_t^{\text{NEW}} - k_t)/(1 + (k_t^{\text{NEW}} - k_t)\tilde{\mathbf{x}}_t^\top\mathbf{v}_t)$. Storing the \mathbf{v}_j 's in the $pn \times n$ matrix $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n]$, these updates translate in the matrix update $\mathbf{V}^{\text{NEW}} = \mathbf{V} - c_t\mathbf{v}_t\tilde{\mathbf{x}}_t^\top\mathbf{V}$. Finally, after the algorithm has converged, one has to derive the form of the EP covariance matrix \mathbf{Q}^{-1} from the quantities used in the updates. It holds $\mathbf{Q}^{-1} = (\sum_{t=0}^n \mathbf{Q}_t)^{-1} = (\boldsymbol{\Omega}^{-1} + \sum_{t=1}^n \mathbf{Q}_t)^{-1} = (\boldsymbol{\Omega}^{-1} + \tilde{\mathbf{X}}^\top\mathbf{K}\tilde{\mathbf{X}})^{-1}$, having defined $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n)^\top$ and $\mathbf{K} = \text{diag}(k_1, \dots, k_n)$. Consequently, calling $\boldsymbol{\Lambda} = (\mathbf{I}_n + \mathbf{K}\tilde{\mathbf{X}}\boldsymbol{\Omega}\tilde{\mathbf{X}}^\top)^{-1}$, exploiting Woodbury's identity, it trivially follows that $\mathbf{Q}^{-1} = \boldsymbol{\Omega} - \boldsymbol{\Omega}\tilde{\mathbf{X}}^\top\boldsymbol{\Lambda}\mathbf{K}\tilde{\mathbf{X}}\boldsymbol{\Omega}$. From this, one has $\mathbf{V} = \mathbf{Q}^{-1}\tilde{\mathbf{X}}^\top = \boldsymbol{\Omega}\tilde{\mathbf{X}}^\top[\mathbf{I}_n - \boldsymbol{\Lambda}\mathbf{K}\tilde{\mathbf{X}}\boldsymbol{\Omega}\tilde{\mathbf{X}}^\top] = \boldsymbol{\Omega}\tilde{\mathbf{X}}^\top\boldsymbol{\Lambda}[\boldsymbol{\Lambda}^{-1} - \mathbf{K}\tilde{\mathbf{X}}\boldsymbol{\Omega}\tilde{\mathbf{X}}^\top] = \boldsymbol{\Omega}\tilde{\mathbf{X}}^\top\boldsymbol{\Lambda}$ so that $\mathbf{Q}^{-1} = \boldsymbol{\Omega} - \mathbf{V}\mathbf{K}\tilde{\mathbf{X}}\boldsymbol{\Omega}$, showing that \mathbf{Q}^{-1} can be computed without any direct matrix inversion. Algorithm 2 reports all the steps needed to perform such efficient EP implementation.

3.3 Computational costs

Calling \tilde{p} the total number of parameters, [11] and [13] noted that, in the case of static probit regression, the cost of each EP update for the corresponding versions of Algorithm 1 and 2 scales as $O(n\tilde{p}^2)$ and $O(n^2\tilde{p})$, respectively. Adapting those derivations to the current dynamic setting, where $\tilde{p} = pn$, one obtains that the cost of each EP update for Algorithms 1 and 2 scales respectively as $O(n^3p^2)$ and $O(n^3p)$, making the latter always more efficient, differently from the static case, where Algorithm 2 becomes advantageous only in high-dimensional scenarios, with $p > n$. The same consideration holds true even accounting for operations required to initialize and conclude the main EP routines. Calling N_{EP} the number of iterations needed to reach convergence, it is straightforward to show that Algorithms 1 and 2 scale respectively as $O(n^3p^2N_{\text{EP}}) = O(n^3p(p \cdot N_{\text{EP}}))$ and $O(n^3p(p + N_{\text{EP}}))$. This result makes our contribution even more significant, as the proposed EP implementation leads to computational improvements regardless of being applied to high-dimensional data \mathbf{x}_t or not. Finally, it is worth mentioning that the sparse nature of $\tilde{\mathbf{x}}_t$ allows us to further reduce the computational costs of different steps in Algorithms 1 and 2. Albeit reducing the number of maximum-cost operations, this does not alter the overall scaling of the EP routines.

4 Financial illustration

We demonstrate the performance of the EP approximation, derived in Section 3 in a financial application that was previously used in [3] for comparisons of some state-of-the-art algorithms to approximate the smoothing distribution (6). More




Fig. 1. $\mathbb{E}[\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}]$ (—) and $\mathbb{E}[\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}] \pm \sqrt{\text{var}[\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}]}$ (- - -) for the i.i.d. sampler (IID) and the EP, PFM-VB and MF-VB approximations.

precisely, we consider a dynamic probit regression for the daily opening directions of the French CAC40 stock market index, spanning from January 4th, 2018, to December 28th, 2018, comprising a total of $n = 241$ observations. In this study, the binary response variable y_t is defined as 1 if the opening value of the CAC40 on the day t exceeds the corresponding closing value from the previous day and 0 otherwise. We incorporate two covariates: the intercept and the opening direction of the NIKKEI225, treated as a binary covariate and denoted as ξ_t . As the Japanese market opens before the French one, ξ_t is available prior to y_t , making it a valid predictor for each day t . Therefore, in reference to model (1)-(2), we have $p = 2$ and $\mathbf{x}_t = (1, \xi_t)^\top$. Additionally, we set $\mathbf{W}_t = \text{diag}(0.01, 0.01)$ for all t and $\mathbf{P}_0 = \text{diag}(3, 3)$. Detailed information about the hyperparameters' setting can be found in [2].

Inference about the smoothing distribution (6) is conducted using as benchmark the results arising from 10^4 samples from the i.i.d. sampler in [2]. We then compare the accuracy of three approximate methods in recovering functionals of interest of the smoothing distribution. More precisely, we consider the EP approximation implemented as in Algorithm 2, the PFM-VB introduced in [3] briefly described in Section 2, and a mean-field variational Bayes (MF-VB) approximation which adapts [16] to the dynamic setting. These results are illustrated in Figure 1, where we plot $\mathbb{E}[\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}]$ and the bands $\mathbb{E}[\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}] \pm \sqrt{\text{var}[\boldsymbol{\theta}_{1:n} | \mathbf{y}_{1:n}]}$. It can be seen that, although all approximate methods reach a reasonable accuracy, MF-VB shows some over-shrinkage of posterior moments towards zero, while EP is slightly more precise than PFM-VB. This is made clearer in Figure 2, where we show the boxplots of the differences between the smoothing means $\mathbb{E}[\theta_{j1} | \mathbf{y}_{1:n}], \dots, \mathbb{E}[\theta_{jn} | \mathbf{y}_{1:n}]$ and log standard deviations $\log\left(\sqrt{\text{var}[\theta_{j1} | \mathbf{y}_{1:n}]}\right), \dots, \log\left(\sqrt{\text{var}[\theta_{jn} | \mathbf{y}_{1:n}]}\right)$, $j = 1, 2$, obtained via the i.i.d. sampler and the ones resulting from the approximate methods. From there, the above-mentioned over-shrinkage of the MF-VB approximation is immediate to notice, while the improvements of EP over PFM-VB are also clarified.




Fig. 2. Boxplot of the differences of the $\mathbb{E}[\theta_{j1} \mid \mathbf{y}_{1:n}], \dots, \mathbb{E}[\theta_{jn} \mid \mathbf{y}_{1:n}]$ and $\log(\sqrt{\text{var}[\theta_{j1} \mid \mathbf{y}_{1:n}]}, \dots, \log(\sqrt{\text{var}[\theta_{jn} \mid \mathbf{y}_{1:n}]})$, $j = 1, 2$, obtained with the EP the PFM-VB, and the MF-VB solutions, using the inferences obtained via i.i.d. sampling from the exact SUN as benchmark.

Finally, we note that the approximated methods enable the computation of the desired moments in less than half a second (on a MacBook Pro 14-inch, 2023 the EP takes 0.43 secs, the PFM-VB takes 0.27 secs, and the MF-VB takes 0.20 secs). In contrast, the i.i.d. sampler requires a significantly longer time of 36.28 seconds. Code for reproducing all the results can be accessed at the following link: <https://github.com/augustofasano/Dynamic-Probit-EP>.

5 Discussion

In this contribution, we have shown how EP can be effectively used to perform approximate inference for the smoothing distribution in dynamic probit models. The EP Gaussian approximation of the smoothing distribution allows estimating functionals of interest with computational times that are orders of magnitude smaller than the ones of exact sampling methods. We have shown in a financial application that the EP approximate moments come with great accuracy, despite being endowed with fewer theoretical guarantees than alternative approximation schemes [10]. This is in line with recent literature about models for partially-observed Gaussian variables [11,12,13], where EP was empirically shown to lead to accurate approximations of posterior quantities of interest. Such results constitute empirical guarantees about the use of EP as an approximate method for Bayesian inference and could motivate future research on further theoretical guarantees.

References

1. Andrieu, C. and Doucet, A. (2002). Particle filtering for partially observed Gaussian state space models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **64**, 827–836.
2. Fasano, A., Rebaudo, G., Durante D., and Petrone S. (2021). A closed-form filter for binary time series. *Statistics and Computing*, **31**:47, 1–20.
3. Fasano, A. and Rebaudo, G. (2021). Variational inference for the smoothing distribution in dynamic probit models. *Book of Short Papers - SIS 2021*, 1076–1081.
4. He, Z. (2023). A dynamic binary probit model with time-varying parameters and shrinkage prior. *Journal of Business & Economic Statistics*, in press.
5. Arellano-Valle, R. B. and Azzalini, A. (2006). On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics*, **33**, 561–574.
6. Vehtari, A., Gelman A., Sivula, T., Jylänki, P., Tran D., Sahai, S., Blomstedt, P., Cunningham J., Schiminovich, D., and Robert C. (2020). Expectation Propagation as a Way of Life: A Framework for Bayesian Inference on Partitioned Data. *Journal of Machine Learning Research*, **21**, 1–53.
7. Bishop, C. M. (2006). *Pattern Recognition and Machine Learning* Springer.
8. Durante, D. (2019). Conjugate Bayes for probit regression via unified skew-normal distributions. *Biometrika*, **106**, 765–779.
9. Botev, Z. (2017). The normal law under linear restrictions: simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* **79**, 125–148.
10. Fasano, A., Durante, D., and Zanella, G. (2022). Scalable and accurate variational Bayes for high-dimensional binary regression models. *Biometrika*, **109**, 901–919.
11. Anceschi, N., Fasano, A., Durante, D., and Zanella, G. (2023). Bayesian conjugacy in probit, tobit, multinomial probit and extensions: a review and new results. *Journal of the American Statistical Association*, **118**, 1451–1469.
12. Chopin, N. and Ridgway, J. (2017). Leave Pima Indians alone: binary regression as a benchmark for Bayesian computation. *Statistical Science*, **32**, 64–87.
13. Fasano, A., Anceschi N., Franzolini, B., and Rebaudo G. (2023). Efficient expectation propagation for posterior approximation in high-dimensional probit models. *Book of Short Papers - SIS 2023*, 1133–1138.
14. Minka, T. P. (2001). Expectation propagation for approximate Bayesian inference. *Proceedings of Uncertainty in Artificial Intelligence*, **17**, 362–369.
15. Azzalini, A. and Capitanio, A. (2014). *The Skew-normal and Related Families*. Cambridge University Press.
16. Consonni, G. and Marin, J. M. (2007). Mean-field variational approximate Bayesian inference for latent variable models. *Computational Statistics & Data Analysis*, **52**, 790–798.