Transportation Research Part B 109 (2018) 238–256

Contents lists available at ScienceDirect

Transportation Research Part B

journal homepage: www.elsevier.com/locate/trb

New matrix-based methods for the analytic evaluation of the multivariate cumulative normal distribution function

Chandra R. Bhat[a,b,*]

[a] The University of Texas at Austin, Department of Civil, Architectural and Environmental Engineering, 301 E. Dean Keeton St. Stop C1761, Austin, TX 78712, USA
[b] The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong

ARTICLE INFO

Article history:
Received 7 August 2017
Revised 17 January 2018
Accepted 18 January 2018

Keywords:
Multivariate normal cumulative distribution function
Multinomial probit
Discrete choice models
Econometric models

ABSTRACT

In this paper, we develop a new matrix-based implementation of the Mendell and Elston (ME) analytic approximation to evaluate the multivariate normal cumulative distribution (MVNCD) function, using an LDLT decomposition method followed by a rank 1 update of the LDLT factorization. Our implementation is easy to code for individuals familiar with matrix-based coding. Further, our new matrix-based implementation for the ME algorithm allows us to efficiently write the analytic matrix-based gradients of the approximated MVNCD function with respect to the abscissae and correlation parameters, an issue that is important in econometric model estimation. In addition, we propose four new analytic methods for approximating the MVNCD function. The paper then evaluates the ability of the multiple approximations for individual MVNCD evaluations as well as multinomial probit model estimation. As expected, in our tests for evaluating individual MVNCD functions, we found that the traditional GHK approach degrades rapidly as the dimensionality of integration increases. Concomitant with this degradation in accuracy is a rapid increase in computational time. The analytic approximation methods are also much more stable across different numbers of dimensions of integration, and even the simplest of these methods is superior to the GHK-500 beyond seven dimensions of integration. Based on all the evaluation results in this paper, we recommend the new Two-Variate Bivariate Screening (TVBS) method proposed in this paper as the evaluation approach for MVNCD function evaluation.

1. Introduction

Many statistical and econometric applications require the evaluation of the multivariate normal cumulative distribution (MVNCD) function. For example, in consumer choice analysis in general, and transportation and marketing analysis in particular, the estimation of such models as the multinomial probit (MNP) model, the multivariate binary and ordered-response models, and the multiple discrete-continuous model all require the computation of the MVNCD function. The computation of the MVNCD function also features in applications in a variety of other stochastic programming fields such as defense and environmental economics, geography, and water management.

Earlier studies have developed many ways to evaluate the MVNCD function. While these methods are general and can be used in a variety of situations, most earlier studies have examined the MVNCD function evaluation in the context of

---

[*] Corresponding author at: The University of Texas at Austin, Department of Civil, Architectural and Environmental Engineering, 301 E. Dean Keeton St. Stop C1761, Austin, TX 78712, USA.

E-mail address: bhat@mail.utexas.edu

https://doi.org/10.1016/j.trb.2018.01.011

MNP estimation. Some of these earlier MVNCD evaluations are based on simulation techniques and others on analytic approximations. Among the simulation methods, the best known approach within a frequentist estimation framework is the GHK probability simulator, named after Geweke (1991), Hajivassiliou (Hajivassiliou and McFadden, 1998), and Keane (1990, 1994). The GHK approach starts with transforming the *correlated* random terms into linear functions of *uncorrelated* standard normal deviates using the Cholesky decomposition of the correlation matrix in the MVNCD evaluation. Doing so helps in recasting the MVNCD as a recursive product of univariate (conditional) cumulative normal distributions (UCCNCD). Each UCCNCD involves the integral over a single-sided truncated normal, which is achieved in the GHK through simulation of pseudo-random draws from a truncated normal. Bhat et al. (2010) embed the Halton approach (rather than the pseudo-random approach) to draw from the truncated normal, because of the better coverage of the Halton draws over the integration space (see Bhat, 2001). Alternative GHK procedures involve the use of multivariate quadrature using sparse grid integration (SGI) (see Heiss and Winschel, 2008) or the use of Efficient Importance Sampling (EIS) within the GHK simulator (see Heiss, 2010). In addition to the above frequentist approach, efficient MCMC methods for models involving the MVNCD function evaluation have also been proposed through data augmentation techniques that make it possible to use standard Bayesian regression techniques (see McCulloch et al., 2000; Imai and van Dyk, 2005).

Among the analytic approximation techniques, one of the first approaches was that proposed by Clark (1961). Unfortunately, this approximation does not perform well for MVNCD evaluations when the random variables are highly correlated or have different variances. In another analytic approximation study, Mendell and Elston (1974) (ME) use the same univariate conditioning approach that formed the basis later for the GHK, except they replace draws from the truncated normal at each conditioning step with approximations of the first two moments of the truncated variables at earlier conditioning steps. This method has also been used by many other authors since, including Rice et al. (1979), Kamakura (1989), and Hutmacher and French (2011). Yet another MVNCD analytic approximation was first proposed by Solow (1990) based on Switzer (1977), and then refined by Joe (1995). This procedure entails the decomposition of the multivariate integral into a product of conditional probabilities. At each step, the conditional probability is approximated based on replacing the conditional events by binary variables and the conditional probability itself as an expectation in a linear regression (with the binary variables for the conditional events serving as exogenous variables, with known covariances amongst themselves based on the correlation matrix of the MVNCD evaluation).

With the many simulation and analytic approaches, a couple of recent studies have examined the accuracy and precision offered by the many approaches in the context of MNP model estimation. Patil et al. (2017) compared all the simulation techniques mentioned earlier with Bhat's (2011) maximum approximate composite marginal likelihood (MACML) approach, which combines the Switzer–Solow–Joe (SSJ) approximation for the MVNCD function with the composite marginal likelihood (CML) inference approach for MNP models. The focus of Patil et al. (2017) study was on the accuracy and precision of MNP parameter recovery in a five-alternative choice context. They find that, among all the simulation-based techniques, the GHK-Halton performs best for MNP estimation. However, among the simulation and MACML approaches, the overall winner in terms of accuracy and precision of underlying parameter recovery, as well as computational time, is the MACML procedure with but one permutation of the ordering of the random term abscissae in its embedded analytic approximation for the MVNCD function. Connors et al. (2014) focused on the analytic approximations corresponding to the ME method and the SSJ method, though they also included the GHK and a couple of other simulation approaches for reference reasons. Unlike Patil et al. (2017), Connors et al. (2014) focused on the ability of the methods to recover the probabilities for individual observations in an MNP setting rather than the underlying choice process parameters at the end of the choice model estimation. They tested four different numbers of alternatives (5, 7, 9, and 15 alternatives) as well as a range of utility values and correlation structures. Their results indicated that, for estimating the probabilities for individual observations, an optimally ordered version of the ME method (to be discussed later in this paper) does much better than the SSJ method with even as many as ten permutations of the abscissae. In addition, they found that the ME method is an order faster than the typical GHK approach in computing the choice probabilities while providing at least the same level of accuracy. A few earlier studies have also done a relatively limited comparison, including Kamakura (1989) who, using a three-to-five alternative set-up in an MNP model, evaluated the ME method with Clark's (1961) approximation and another method proposed by Langdon (1984). He found that the ME method works best relative to the other two, in both evaluating the MVNCD function (as reflected in individual choice probabilities) as well as the underlying MNP model parameters. Joe (1995) tested the SSJ approximation of two different orders (a first order one that entails the evaluation of univariate and bivariate cumulative normal distributions and a more accurate second order one that entails the evaluation of trivariate and quadrivariate cumulative normal distributions) in the context of MVNCD evaluations (rather than MNP parameter recovery). He observed that the SSJ approximation of the first order does better than the ME as well as simulated versions for up to 20 dimensions, though his SSJ approximation is based on averaging over the results of up to 2000 permutations of the abscissae (or all permutations of abscissae if this is less than 2000) for each MVNCD evaluation.

In the current paper, we first propose a streamlined and matrix-based version of the ME method that relies on a single-sided truncation of a multivariate normal distribution in which some variables are truncated while others are not. A number of recent papers have focused on such multivariate distributions and studied the properties of the resulting distributions (see, for example, Kim and Kim, 2015), making use of results related to the moments of truncated multivariate normal distributions (Manjunath and Wilhelm, 2012; Kan and Robotti, 2017) and using a regression technique (see Kotz et al., 2000, p. 70) to obtain the mean and covariance matrices of the untruncated variables from the moments of the truncated variables. We use this approach, except propose a new way to implement this approach using an LDLT decomposition method for