Title
Discretizing Distributions with Exact Moments: Error Estimate and Convergence Analysis

Journal
SIAM Journal on Numerical Analysis, 53(5)

ISSN
0036-1429 1095-7170

Authors
Tanaka, Ken'ichiro
Toda, Alexis Akira

Publication Date
2015

DOI
10.1137/140971269

Peer reviewed

DISCRETIZING DISTRIBUTIONS WITH EXACT MOMENTS: ERROR ESTIMATE AND CONVERGENCE ANALYSIS*

KEN'ICHIRO TANAKA† AND ALEXIS AKIRA TODA‡

Abstract. The maximum entropy principle is a powerful tool for solving underdetermined inverse problems. This paper considers the problem of discretizing a continuous distribution, which arises in various applied fields. We obtain the approximating distribution by minimizing the Kullback-Leibler information (relative entropy) of the unknown discrete distribution relative to an initial discretization based on a quadrature formula subject to some moment constraints. We study the theoretical error bound and the convergence of this approximation method as the number of discrete points increases. We prove that (i) the theoretical error bound of the approximate expectation of any bounded continuous function has at most the same order as the quadrature formula we start with, and (ii) the approximate discrete distribution weakly converges to the given continuous distribution. Moreover, we present some numerical examples that show the advantage of the method and apply to numerically solving an optimal portfolio problem.

Key words. probability distribution, discrete approximation, generalized moment, quadrature formula, Kullback-Leibler information, Fenchel duality

AMS subject classifications. 41A25, 41A29, 62E17, 62P20, 65D30, 65K99

1. Introduction.

This paper has two goals. First, we propose a numerical method to approximate continuous probability distributions by discrete ones. Second, we study the convergence of the method and derive error estimates. Discretizing continuous distributions is important in applied numerical analysis [26, 35]. To motivate the problem, we list a few concrete examples, many of which come from economics.

Optimal portfolio problem. Suppose that there are J assets indexed by $j = 1, \dots, J$. Asset j has gross return R_j (which is a random variable), which means that a dollar invested in asset j will give a total return of R_j dollars over the investment horizon. Let θ_j be the fraction of an investor's wealth invested in asset j (so $\sum_{j=1}^J \theta_j = 1$) and $\theta = (\theta_1, \dots, \theta_J)$ be the portfolio. Then the gross return on the portfolio is the weighted average of each asset return, $R(\theta) := \sum_{j=1}^J R_j \theta_j$. Assume that the investor wishes to maximize the risk-adjusted expected returns $\mathrm{E}\left[\frac{1}{1-\gamma} R(\theta)^{1-\gamma}\right]$, where $\gamma > 0$ is the degree of relative risk aversion.[1] Since in general this expectation has no closed-form expression in the parameter θ, in order to maximize it we need to carry out a numerical integration. This problem is equivalent to finding an approximate discrete distribution of the returns $(R_1, \dots, R_J)$.

Optimal consumption-portfolio problem. In the above example the portfolio choice was a one time decision. But we can consider a dynamic version, where the investor chooses the optimal consumption and portfolio over time (Samuelson [37] and Merton [34] are classic examples that admit closed-form solutions. Almost all practical problems, however, admit no closed-form solutions). Since we need to numerically solve a portfolio problem for each time period, we face one more layer of complexity.

*March 13, 2015

†Corresponding author: School of Systems Information Science, Future University Hakodate (ketanaka@fun.ac.jp).

‡Department of Economics, University of California San Diego (atoda@ucsd.edu).

[1] Readers unfamiliar with economic concepts need not worry here. The point is that we want to maximize an expectation that is a function of some parameter. The case $\gamma = 1$ corresponds to the log utility $\mathrm{E}[\log R(\theta)]$.

General equilibrium problem. Instead of solving the optimization problem of a single investor, we can consider a model of the whole economy, and might want to determine the asset prices that make demand equal to supply. This is called a general equilibrium problem in economics. Since we need to solve a dynamic optimal portfolio problem given asset prices, and we need to find asset prices that clear markets, we face yet another layer of complexity. Such problems are especially computationally intensive [22, 4, 29], and it is important to discretize a continuous distribution with only a small number of support points in order to mitigate the 'curse of dimensionality'.

Discretizing stochastic processes. In many fields, including decision analysis, economic modeling, and option pricing, it is necessary to discretize stochastic processes such as diffusions or autoregressive processes to generate scenario trees [21, 36] or finite state Markov chain approximations [43, 44, 19].

The above examples have the following common feature. The researcher is given a continuous density f (transition densities in the case of a stochastic processes), which comes from either some theoretical model or data (say the kernel density estimation). The density f is used to solve a complicated model. In order to reduce the complexity of the problem, the researcher wishes to discretize the density.

Since the density is ultimately used to compute expectations, a natural idea is to start with some quadrature formula

$$\mathrm{E}[g(X)] = \int_{\mathbb{R}^K} g(x)f(x)\,\mathrm{d}x \approx \sum_{i=1}^{I_M} w_{i,M}\, g(x_{i,M}) f(x_{i,M}), \tag{1.1}$$

where M is an index of the quadrature formula (for example, the grid size in each dimension), I_M is the number of integration points $\{x_{i,M}\}$, and $w_{i,M}$ is the weight on the point $x_{i,M}$. Then the probability mass function $q_{i,M} \propto w_{i,M} f(x_{i,M})$ is a valid discrete approximation of the continuous density f. Popular quadrature formulas—such as the Newton-Cotes type or the Gauss type formulas (see [14] for a standard textbook treatment)—may not be suitable for our purpose, however. First, the choice of the quadrature formula automatically determines the discrete points $\{x_{i,M}\}$. However, in practice the researcher often wishes to choose the set of points $D \subset \mathbb{R}^K$ at his or her will, depending on the specific application. Second, for a general distribution f, the quadrature formula (1.1) may not even match low order moments such as the mean or the variance, *i.e.*,

$$\mathrm{E}[X^l] = \int_{-\infty}^{\infty} x^l f(x)\,\mathrm{d}x \neq \sum_{i=1}^{I_M} w_{i,M} x_{i,M}^l f(x_{i,M})$$

for $l = 1, 2$ (in the one-dimensional case).[2] Although the Newton-Cotes and the Gauss formulas (in the one-dimensional case) match low order moments exactly, the multi-dimensional case is not trivial. This is a major disadvantage because matching moments is known to improve the accuracy of the analysis [39]. Third, some quadrature formulas have *negative* weights $w_{i,M}$ when the order is high. Since we want to obtain probabilities, which are necessarily nonnegative, such formulas are inappropriate. Thus it is desirable to obtain a discretization method that (i) has a flexible choice

[2] Matching moments of a single density with a quadrature formula is not hard. The difficulty arises when discretizing a stochastic process. If we approximate a stochastic process by a finite-state Markov chain, when the process hits the lower or upper boundary of the support points, the conditional moments under the true transition density and the approximate probabilities implied by the quadrature formula will be vastly different, making the approximation poor.

of the integration points, (ii) matches at least low order moments, and (iii) always assigns positive weights.

Several methods for discrete approximations of continuous distributions have been proposed in the literature. Given a continuous probability distribution, Tauchen [43] and Adda and Cooper [3] adopt simple partitions of the support of the distribution and assign the true probability to a representative point of each partitioned domain. Although their methods are intuitive, simple, and work in any dimension, their methods are not so accurate, for they generate discrete distributions with only approximate moments. Miller and Rice [35], Tauchen and Hussey [44], and Smith [39] discretize the density function using the weights and the points of the Gaussian quadrature, and Devuyst and Preckel [16] consider its generalization to multi-dimensions. Although their methods are often more accurate and can match prescribed polynomial moments exactly, they do not allow for the restriction $\{x_{i,M}\} \subset D$ and cannot be applied to non-polynomial moments. Furthermore, the multi-dimensional method by Devuyst and Preckel [16] is computationally intensive and does not have a theoretical guarantee for the existence of the discretization, error bounds, or convergence.

As a remedy, in Tanaka and Toda [42] we proposed an approximation method based on the maximum entropy principle (MaxEnt) that exactly matches prescribed moments of the distribution. Starting with any discretization, we "fine-tune" the given probabilities by minimizing the Kullback-Leibler information (relative entropy) of the unknown probabilities subject to some moment constraints. In that paper we proved the existence and the uniqueness of the solution to the minimization problem, showed that the solution can be easily computed by solving the dual problem, and presented some numerical examples that show that the approximation method is satisfactorily accurate. The method is computationally very simple and works on any discrete set D of any dimension with any prescribed moments (not necessarily polynomials). However, up to now the theoretical approximation error and the convergence of this method remain unknown.

This paper gives a theoretical error bound for this approximation method and shows its convergence. We first evaluate the theoretical error of our proposed method. It turns out that the order of the theoretical error estimate is at most that of the initial discretization, and actually improves if the integrand is well-approximated by the moment defining function. Thus our proposed method does not compromise the order of the error at the expense of matching moments. Second, as a theoretical consequence of the error estimate, we show the weak convergence of the discrete distribution generated by the method to the given continuous distribution. This means that for any bounded continuous function g, the expectation of g under the approximating discrete distribution converges to that under the exact distribution as the number of integration points increases. This property is practically important because it guarantees that the approximation method never generates a pathological discrete distribution with exact moments which has extremely different probability from the given distribution on some domain, at least when the discrete set is large enough. In addition, we present some numerical examples (including a numerical solution to an optimal portfolio problem) that show the advantage of our proposed method.

The idea of using the maximum entropy principle to obtain a solution to underdetermined inverse problems (such as the Hausdorff moment problem) is similar to that of Jaynes [24] and Mead and Papanicolaou [33]. There is an important distinction, however. In typical inverse problems, one studies the convergence of the approximating solution to the true one when *the number of moment constraints* tends to infinity. In contrast, in this paper we study the convergence when *the number of approximating points* tends to infinity, fixing the moments. Thus the two problems are quite different. The literature on the foundations, implementations, and the applications of maximum entropy methods is immense: any literature review is necessarily partial. The maximum entropy principle (as an inference method in general, not necessarily restricted to physics) was proposed by Jaynes [23]. For axiomatic approaches, see Shore and Johnson [38], Jaynes [25], Caticha and Giffin [11], and Knuth and Skilling [28]. For the relation to Bayesian inference, see Van Campenhout and Cover [48] and Csiszár [13]. For the duality theory of entropy maximization, see [7, 15, 20]. For numerical algorithms for computing maximum entropy densities, see [1, 2, 5]. Budišić and Putinar [10] study the opposite problem of ours, namely transforming a probability measure to one that is absolutely continuous with respect to the Lebesgue measure. Applications of maximum entropy methods can be found in economics [18, 45, 46], statistics and econometrics [6, 27, 49], finance [40, 41, 9], among many other fields.

2. The approximation method. This section reviews the discrete approximation method of continuous distributions proposed in [42].

Let f be a probability density function on $\mathbb{R}^K$ with some generalized moments

$$\bar{T} = \int_{\mathbb{R}^K} f(x)T(x)\,\mathrm{d}x, \tag{2.1}$$

where $T : \mathbb{R}^K \to \mathbb{R}^L$ is a continuous function. (Below, we sometimes refer to this function as the "moment defining function".) For instance, if we are interested in the first and second polynomial moments (*i.e.*, mean and variance), T becomes

$$T(x) = (x_1, \dots, x_K, x_1^2 \dots, x_k x_l, \dots, x_K^2).$$

In this case, we have K expectations, K variances, and $\frac{K(K-1)}{2}$ covariances. Hence the total number of moment constraints (the dimension of the range space of T) is

$$L = K + K + \frac{K(K-1)}{2} = \frac{K(K+3)}{2}. \tag{2.2}$$

In general, the components of $T(x)$ need not be polynomials.

Moreover, for each positive integer M, assume that a finite discrete set

$$D_M = \{x_{i,M} \mid i = 1, \dots, I_M\} \subset \mathbb{R}^K$$

is given, where I_M is the number of discrete points. An example of D_M is the lattice

$$D_M = \{(m_1 h, m_2 h, \dots, m_K h) \mid m_1, m_2 \dots, m_K = 0, \pm 1, \dots, \pm M\},$$

where $h > 0$ is the grid size, in which case $I_M = (2M+1)^K$. Our aim is to find a discrete probability distribution

$$P_M = \{p_{i,M} \mid x_{i,M} \in D_M\}$$

on D_M with exact moments $\bar{T}$ that approximates f (in the sense of the weak topology, that is, convergence in distribution).

To match the moments $\bar{T}$ with $P_M = \{p_{i,M} \mid x_{i,M} \in D_M\}$, it suffices to assign probabilities $\{p_{i,M}\}$ such that

$$\sum_{i=1}^{I_M} p_{i,M} T(x_{i,M}) = \bar{T}. \tag{2.3}$$

Note that the solution to this equation is generally underdetermined because the number of unknowns—$p_{i,M}$'s—of which there are I_M, is typically much larger than the number of equations (moments), $L+1$.[3] To obtain P_M that approximates f and satisfies (2.3), we first choose an arbitrary discretization of f with *not necessarily exact moments*, which we denote by $Q_M = \{q_{i,M}\}$. For example, if we already have a numerical quadrature formula

$$\int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x \approx \sum_{i=1}^{I_M} w_{i,M} f(x_{i,M}) g(x_{i,M}) \tag{2.4}$$

with positive weights $w_{i,M}$ $(i = 1, 2, \dots, I_M)$, where g is an arbitrary function that we want to compute the expectation with respect to the density f, then it is natural to set $q_{i,M}$ proportional to $w_{i,M} f(x_{i,M})$, *i.e.*,

$$q_{i,M} = \frac{w_{i,M} f(x_{i,M})}{\sum_{i=1}^{I_M} w_{i,M} f(x_{i,M})}.$$

In the following, we do not address how to choose Q_M (or the quadrature formula (2.4)) but take it as given.

Now the approximation method is defined as follows. Let g be any bounded continuous function. Then the approximation error of the expectation under P_M is

$$\begin{aligned} &\left| \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x - \sum_{i=1}^{I_M} p_{i,M} g(x_{i,M}) \right| \\ &\le \left| \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x - \sum_{i=1}^{I_M} q_{i,M} g(x_{i,M}) \right| + \left| \sum_{i=1}^{I_M} q_{i,M} g(x_{i,M}) - \sum_{i=1}^{I_M} p_{i,M} g(x_{i,M}) \right| \\ &\le \left| \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x - \sum_{i=1}^{I_M} q_{i,M} g(x_{i,M}) \right| + \|g\|_\infty \sum_{i=1}^{I_M} |p_{i,M} - q_{i,M}|, \end{aligned} \tag{2.5}$$

where $\|g\|_\infty = \sup_{x \in \mathbb{R}^K} |g(x)|$ is the sup norm. Since the first term depends only on the initial discretization Q_M, we focus on the second term. By Pinsker's inequality,[4] the second term can be bounded as follows:

$$\sum_{i=1}^{I_M} |p_{i,M} - q_{i,M}| \le \sqrt{2H(P_M; Q_M)}, \tag{2.6}$$

where

$$H(P_M; Q_M) = \sum_{i=1}^{I_M} p_{i,M} \log \frac{p_{i,M}}{q_{i,M}} \tag{2.7}$$

[3]The "+1" comes from accounting the probabilities $\sum_{i=1}^{I_M} p_{i,M} = 1$.

[4]See [12, 30] for proofs of Pinsker's inequality and [17] and references therein for refinements. The appendix of this paper gives a short proof of Pinsker's inequality.

is the Kullback-Leibler information [31].

In order to make the approximation error (2.5) small, given Pinsker's inequality (2.6), it is quite natural to obtain the approximate discrete distribution $P_M = \{p_{i,M} \mid x_{i,M} \in D_M\}$ as the solution to the following optimization problem:

$$\text{(P)} \qquad \min_{\{p_{i,M}\}} \sum_{i=1}^{I_M} p_{i,M} \log \frac{p_{i,M}}{q_{i,M}}$$
$$\text{subject to } \sum_{i=1}^{I_M} p_{i,M} T(x_{i,M}) = \bar{T},\ \sum_{i=1}^{I_M} p_{i,M} = 1,\ p_{i,M} \geq 0.$$

The first constraint matches the moments $\int f(x)T(x)\,\mathrm{d}x = \bar{T}$ exactly. $\sum_{i=1}^{I_M} p_{i,M} = 1$ and $p_{i,M} \geq 0$ ensure that $\{p_{i,M}\}$ is a probability mass function. The problem (P) has a unique solution if $\bar{T} \in \operatorname{co} T(D_M)$, where $\operatorname{co} T(D_M)$ is the convex hull of $T(D_M)$ defined by

$$\operatorname{co} T(D_M) = \left\{ \sum_{i=1}^{I_M} \alpha_{i,M} T(x_{i,M}) \,\middle|\, \sum_{i=1}^{I_M} \alpha_{i,M} = 1 \text{ and } \alpha_{i,M} \geq 0 \right\}, \tag{2.8}$$

because in that case the constraint set is nonempty, compact, convex, and the objective function is continuous (by adopting the convention $0 \log 0 = 0$) and strictly convex.

To characterize the solution of (P), we consider the Fenchel dual[5] of (P),

$$\text{(D)} \qquad \max_{\lambda \in \mathbb{R}^L} \left[\left\langle \lambda, \bar{T} \right\rangle - \log \left(\sum_{i=1}^{I_M} q_{i,M} \mathrm{e}^{\langle \lambda, T(x_{i,M}) \rangle} \right) \right],$$

where $\langle \cdot , \cdot \rangle$ denotes the usual inner product in $\mathbb{R}^L$. Tanaka and Toda [42] show that we can obtain the solution to (P) as fine-tuned values of $q_{i,M}$, and that the minimum value of (P) and the maximum value of (D) coincide. Although these properties are routine exercises in convex duality theory (see for example [8] for a textbook treatment), we present them nevertheless in order to make the paper self-contained.

Theorem 1. *Suppose that* $\bar{T} \in \operatorname{int} \operatorname{co} T(D_M)$, *where* int *denotes the interior. Then the following is true.*

1. *The dual problem* (D) *has a unique solution* λ_M.
2. *The probability distribution* $P_M = \{p_{i,M} \mid x_{i,M} \in D_M\}$ *defined by*

$$p_{i,M} = \frac{q_{i,M} \mathrm{e}^{\langle \lambda_M, T(x_{i,M}) \rangle}}{\sum_{i=1}^{I_M} q_{i,M} \mathrm{e}^{\langle \lambda_M, T(x_{i,M}) \rangle}} \tag{2.9}$$

   is the unique solution to (P).
3. *The duality* $H(P_M; Q_M) = \min \text{(P)} = \max \text{(D)}$ *holds.*

Our proposed approximation method has two advantages. The first is the computational simplicity. The dual problem (D) is an *unconstrained* optimization problem with typically a small number of unknowns (L), whereas the primal problem (P) is a *constrained* optimization problem with typically a large number of unknowns (I_M). With existing methods such as Gospodinov and Lkhagvasuren [19] that target the first

[5] See [7] for an application of the Fenchel duality to entropy-like minimization problems.

and second moments with a quadratic loss function, the optimization problem remains high dimensional and constrained. For example, if we discretize a 3-dimensional distribution with 10 discrete points in each dimension and match the mean and the variances, then there will be $10^3 = 1,000$ unknowns, 9 moment constraints (according to (2.2)), and 1,000 nonnegativity constraints. On the other hand, our dual problem (D) is unconstrained and involves only 9 variables. The second advantage is that the resulting weights on the points (2.9) are automatically positive, as they should be since they are probabilities.

3. Error bound and convergence. Let $g : \mathbb{R}^K \to \mathbb{R}$ be a bounded continuous function and

$$E^{(Q)}_{g,M} = \left| \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x - \sum_{i=1}^{I_M} q_{i,M}\, g(x_{i,M}) \right|, \tag{3.1}$$

$$E^{(P)}_{g,M} = \left| \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x - \sum_{i=1}^{I_M} p_{i,M}\, g(x_{i,M}) \right| \tag{3.2}$$

be the approximation errors under the initial discretization Q_M and $P_M = \{p_{i,M}\}$ in (2.9). In this section we estimate the error $E^{(P)}_{g,M}$ and prove the weak convergence[6] of $P_M = \{p_{i,M}\}$ to f, *i.e.*, $E^{(P)}_{g,M} \to 0$ as $M \to \infty$ for any g. Using the definition of the errors (3.1) and (3.2), the triangle inequality (2.5), and Pinsker's inequality (2.6), we obtain the following error estimate:

$$E^{(P)}_{g,M} \le E^{(Q)}_{g,M} + \|g\|_\infty \sqrt{2H(P_M; Q_M)}, \tag{3.3}$$

where $H(P_M; Q_M)$ is the Kullback-Leibler information (2.7).

We consider the error estimate and the convergence analysis under the following two assumptions. The first assumption states that the moment defining function T has no degenerate components and the moment $\bar{T}$ can also be expressed as an expectation on the discrete set D_M.

Assumption 1. *The components of the moment defining function T are affine independent on $\mathbb{R}^L \cap \operatorname{supp} f$: for any $0 \neq (\lambda, \mu) \in \mathbb{R}^L \times \mathbb{R}$, there exists $x \in \operatorname{supp} f$ such that $\langle \lambda, T(x) \rangle + \mu \neq 0$. Furthermore, $\bar{T} \in \operatorname{int}(\operatorname{co} T(D_M))$ for any M.*

The second assumption concerns the convergence of the initial discretization $Q_M = \{q_{i,M}\}$.

Assumption 2. *The initial discretization weakly converges: for any bounded continuous function g on $\mathbb{R}^K$, we have*

$$\lim_{M\to\infty} \sum_{i=1}^{I_M} q_{i,M}\, g(x_{i,M}) = \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x. \tag{3.4}$$

Furthermore, the targeted moments converge as well:

$$\lim_{M\to\infty} \sum_{i=1}^{I_M} q_{i,M}\, T(x_{i,M}) = \int_{\mathbb{R}^K} f(x)T(x)\,\mathrm{d}x. \tag{3.5}$$

[6]For readers unfamiliar with probability theory, a sequence of probability measures $\{\mu_n\}$ is said to *weakly converge* to μ if $\lim_{n\to\infty} \int g\,\mathrm{d}\mu_n = \int g\,\mathrm{d}\mu$ for every bounded continuous function g. In particular, by choosing g as an indicator function, we have $\mu_n(B) \to \mu(B)$ for any Borel set B, so the probability distribution μ_n approximates μ.

Assumption 2 is natural since any discrete approximation should become accurate as we increase the number of grid points. Under this assumption, we have $E^{(Q)}_{g,M} \to 0$ as $M \to \infty$, so by (3.3) we have $E^{(P)}_{g,M} \to 0$ whenever $H(P_M; Q_M) \to 0$ as $M \to \infty$. Below, we focus on estimating $H(P_M; Q_M)$.

Lemma 2. *Define the function $J_M : \mathbb{R}^L \to \mathbb{R}$ by*

$$J_M(\lambda) = \sum_{i=1}^{I_M} q_{i,M} \mathrm{e}^{\left\langle \lambda, T(x_{i,M}) - \bar{T} \right\rangle}. \tag{3.6}$$

Then

$$H(P_M; Q_M) = -\log\left(\min_{\lambda \in \mathbb{R}^L} J_M(\lambda)\right). \tag{3.7}$$

Proof. By Theorem 1 and the definition of J_M, we obtain

$$\begin{aligned} H(P_M; Q_M) &= \max_{\lambda \in \mathbb{R}^L} \left[\left\langle \lambda, \bar{T} \right\rangle - \log\left(\sum_{i=1}^{I_M} q_{i,M} \mathrm{e}^{\langle \lambda, T(x_{i,M}) \rangle} \right) \right] && (\because \text{(D), Theorem 1}) \\ &= \max_{\lambda \in \mathbb{R}^L} \left[-\log\left(\sum_{i=1}^{I_M} q_{i,M} \mathrm{e}^{\left\langle \lambda, T(x_{i,M}) - \bar{T} \right\rangle} \right) \right] && (\because \textstyle\sum q_{i,M} = 1) \\ &= \max_{\lambda \in \mathbb{R}^L} [-\log(J_M(\lambda))] = -\log\left(\min_{\lambda \in \mathbb{R}^L} J_M(\lambda)\right), && (\because (3.6)) \end{aligned}$$

which is (3.7). ∎

By Lemma 2, in order to estimate the Kullback-Leibler information $H(P_M; Q_M)$ from above, it suffices to bound $J_M(\lambda)$ from below. To this end let

$$E^{(Q)}_{T,M} = \left\| \int_{\mathbb{R}^K} f(x) T(x) \, \mathrm{d}x - \sum_{i=1}^{I_M} q_{i,M} T(x_{i,M}) \right\| \tag{3.8}$$

be the initial approximation error for the moments T,

$$C_\alpha = \inf_{\lambda \in \mathbb{R}^L, \|\lambda\| = 1} \int_{\mathbb{R}^K} f(x) \left(\max\left\{0, \min\left\{ \left\langle \lambda, T(x) - \bar{T} \right\rangle, \alpha \right\}\right\} \right)^2 \mathrm{d}x \tag{3.9}$$

for any $\alpha > 0$, and $C_\infty = \lim_{\alpha \to \infty} C_\alpha$. The following lemma gives a quadratic lower bound of J_M.

Lemma 3. *Let Assumptions 1 and 2 be satisfied. Then*

1. $0 < C_\alpha \leq C_\infty \leq \infty$.
2. *For any C with $0 < C < C_\infty$, there exists a positive integer M_C such that for any $M \geq M_C$, we have*

$$J_M(\lambda) \geq 1 - E^{(Q)}_{T,M} \|\lambda\| + \frac{1}{2} C \|\lambda\|^2. \tag{3.10}$$

Proof. First we prove $0 < C_\alpha \leq C_\infty \leq \infty$. Since the integrand in (3.9) is nonnegative and increasing in $\alpha > 0$, clearly $C_\alpha \leq C_\infty \leq \infty$. Due to the presence of the max and min operators, the integrand is positive if and only if $x \in \operatorname{supp} f$

and $\left\langle \lambda, T(x) - \bar{T} \right\rangle > 0$. Since the components of T are affine independent on $\operatorname{supp} f$ (Assumption 1) and

$$\begin{aligned}\bar{T} = \int_{\mathbb{R}^K} f(x)T(x)\,\mathrm{d}x &\iff \int_{\mathbb{R}^K} f(x)(T(x)-\bar{T})\,\mathrm{d}x = 0\\ &\iff (\forall \lambda) \int_{\mathbb{R}^K} f(x)\left\langle \lambda, T(x)-\bar{T}\right\rangle\,\mathrm{d}x = 0,\end{aligned}$$

for any $\lambda \neq 0$ there exists a region in $\operatorname{supp} f$ for which $\left\langle \lambda, T(x) - \bar{T} \right\rangle > 0$. Since T is continuous, the integrand in (3.9) is positive on a set with positive measure and continuous with respect to λ, so $C_\alpha > 0$.

Next we prove (3.10). For this purpose, we use the inequality

$$\mathrm{e}^z \geq 1 + z + \frac{1}{2}\left(\max\left\{0, \min\left\{z, a\right\}\right\}\right)^2 \tag{3.11}$$

for any $z, a \in \mathbb{R}$. (3.11) follows from

$$\max\left\{0, \min\left\{z,a\right\}\right\} = \begin{cases} 0, & (z \leq 0 \text{ or } a \leq 0)\\ z, & (0 \leq z \leq a)\\ a, & (0 \leq a \leq z)\end{cases}$$

$\mathrm{e}^z \geq 1 + z$ if $z \leq 0$, and $\mathrm{e}^z \geq 1 + z + \frac{1}{2}z^2$ if $z \geq 0$.

Let $z = \left\langle \lambda, T(x_{i,M}) - \bar{T}\right\rangle$ and $a = \|\lambda\|\,\alpha$ for $\alpha > 0$ in the inequality (3.11). Then

$$\mathrm{e}^{\left\langle \lambda, T(x_{i,M})-\bar{T}\right\rangle} \geq 1 + \left\langle \lambda, T(x_{i,M}) - \bar{T}\right\rangle + \frac{1}{2}\left(\max\left\{0, \min\left\{\left\langle \lambda, T(x_{i,M}) - \bar{T}\right\rangle, \|\lambda\|\,\alpha\right\}\right\}\right)^2.$$

Multiplying both sides by $q_{i,M} \geq 0$, letting $\lambda^* = \frac{\lambda}{\|\lambda\|}$, and summing over i, we obtain

$$J_M(\lambda) \geq 1 + \langle \lambda, B_M\rangle + \frac{1}{2}C_{M,\alpha}(\lambda^*)\,\|\lambda\|^2, \tag{3.12}$$

where

$$B_M = \sum_{i=1}^{I_M} q_{i,M}(T(x_{i,M}) - \bar{T}), \tag{3.13}$$

$$C_{M,\alpha}(\lambda^*) = \sum_{i=1}^{I_M} q_{i,M}\left(\max\left\{0, \min\left\{\left\langle \lambda^*, T(x_{i,M}) - \bar{T}\right\rangle, \alpha\right\}\right\}\right)^2. \tag{3.14}$$

Letting $M \to \infty$ in (3.14), we obtain

$$\begin{aligned}C_{M,\alpha}(\lambda^*) &\to \int f(x)\left(\max\left\{0, \min\left\{\left\langle \lambda^*, T(x) - \bar{T}\right\rangle, \alpha\right\}\right\}\right)^2\,\mathrm{d}x && (\because \text{Assumption 2})\\ &\geq C_\alpha. && (\because (3.9))\end{aligned}$$

Since $\lim_{\alpha\to\infty} C_\alpha = C_\infty > C > 0$, we can take $\alpha > 0$ large enough such that $C_\alpha > C$. Since $\lim_{M\to\infty} C_{M,\alpha}(\lambda^*) \geq C_\alpha > C$, we can take M_C such that for any $M \geq M_C$ we have $C_{M,\alpha}(\lambda^*) \geq C$. Then by (3.12) and the Cauchy-Schwarz inequality, we obtain

$$J_M(\lambda) \geq 1 - \|B_M\|\,\|\lambda\| + \frac{1}{2}C\,\|\lambda\|^2.$$

Since $\|B_M\| = E^{(Q)}_{T,M}$ by (3.8), $\int_{\mathbb{R}^K} f(x)T(x)\,\mathrm{d}x = \bar{T}$, and (3.13), we obtain the conclusion (3.10). ◻

Combining the above lemmas, we obtain the following estimate of $E^{(P)}_{g,M}$.

Theorem 4. *Let Assumptions 1 and 2 be satisfied and g be a bounded continuous function. Then, for any $C > 0$ be with $0 < C < C_\infty = \lim_{\alpha\to\infty} C_\alpha$, there exists a positive integer M_C such that for any M with $M \geq M_C$, we have*

$$E^{(P)}_{g,M} \leq E^{(Q)}_{g,M} + \|g\|_\infty \sqrt{-2\log\left(1 - \frac{(E^{(Q)}_{T,M})^2}{2C}\right)}, \tag{3.15}$$

where the right-hand side is interpreted as ∞ if the inside of logarithm is negative.

Proof. For notational simplicity let $E := E^{(Q)}_{T,M}$. Minimizing the right-hand side of (3.10) analytically, we obtain

$$J_M(\lambda) \geq 1 - \frac{E^2}{2C}. \tag{3.16}$$

If $E^2 \geq 2C$, (3.16) is not tight since $J_M \geq 0$ always. In this case, we obtain the trivial estimate $E^{(P)}_{g,M} \leq \infty$. Otherwise,

$$\begin{aligned} E^{(P)}_{g,M} &\leq E^{(Q)}_{g,M} + \|g\|_\infty \sqrt{-2\log\left(\min_\lambda J_M(\lambda)\right)} && (\because (3.3), (3.7)) \\ &\leq E^{(Q)}_{g,M} + \|g\|_\infty \sqrt{-2\log(1 - E^2/2C)}, && (\because (3.16)) \end{aligned}$$

which is (3.15). ◻

Note that $E^{(P)}_{g,M}$ is bounded by a formula consisting of $E^{(Q)}_{g,M}$ and $E^{(Q)}_{T,M}$, which are the errors of the initial discretization Q_M for the functions g and T. Since both of them tend to zero as $M \to \infty$ by Assumption 2, it follows from (3.15) and $-\log(1-t) \approx t$ for small t that

$$E^{(P)}_{g,M} = \mathrm{O}\left(\max\left\{E^{(Q)}_{g,M}, E^{(Q)}_{T,M}\right\}\right) \quad (M \to \infty). \tag{3.17}$$

The equality (3.17) shows that the error $E_{g,M}$ is at most of the same order as the error of the initial discretization. Thus our method does not compromise the order of the error at the expense of matching moments.

Using Theorem 4, we can prove our main result, the weak convergence of the approximating discrete distribution $P_M = \{p_{i,M}\}$ to f.

Theorem 5. *Let Assumptions 1 and 2 be satisfied. Then, for any bounded continuous function g, we have*

$$\lim_{M\to\infty} \sum_{i=1}^{I_M} p_{i,M}\, g(x_{i,M}) = \int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x, \tag{3.18}$$

i.e.*, the discrete distribution P_M weakly converges to the exact distribution f.*

Proof. By the definition of $E^{(P)}_{g,M}$ in (3.2), (3.17), and Assumption 2, we get

$$\left|\int_{\mathbb{R}^K} f(x)g(x)\,\mathrm{d}x - \sum_{i=1}^{I_M} p_{i,M}\, g(x_{i,M})\right| = E^{(P)}_{g,M} = \mathrm{O}\left(\max\left\{E^{(Q)}_{g,M}, E^{(Q)}_{T,M}\right\}\right) \to 0$$

as $M \to \infty$, which is (3.18). ∎

The following theorem gives a tighter error estimate when $E^{(Q)}_{T,M}$ is large.

Theorem 6. *Let everything be as in Theorem 4. Then*

$$E^{(P)}_{g,M} \le E^{(Q)}_{g,M} + \|g\|_\infty \frac{2}{\sqrt{C}} E^{(Q)}_{T,M}. \tag{3.19}$$

Proof. By (3.3) it suffices to show

$$H(P_M; Q_M) \le \frac{2}{C}(E^{(Q)}_{T,M})^2. \tag{3.20}$$

Let $\lambda_M \in \mathbb{R}^L$ be the solution to the dual problem (D). By Theorem 1, we have

$$\begin{aligned} H(P_M;Q_M) &= -\log\left(\sum_{i=1}^{I_M} q_{i,M} \mathrm{e}^{\left\langle \lambda_M, T(x_{i,M}) - \bar{T}\right\rangle}\right) \\ &\le -\sum_{i=1}^{I_M} q_{i,M} \log\left(\mathrm{e}^{\left\langle \lambda_M, T(x_{i,M}) - \bar{T}\right\rangle}\right) && (\because -\log(\cdot) \text{ is convex}) \\ &= -\left\langle \lambda_M, \sum_{i=1}^{I_M} q_{i,M} T(x_{i,M}) - \bar{T}\right\rangle && (\because \textstyle\sum q_{i,M} = 1) \\ &\le \|\lambda_M\| \, E^{(Q)}_{T,M}. && (\because \text{Cauchy-Schwarz, (3.8)}) \end{aligned} \tag{3.21}$$

Since λ_M solves (D), and hence minimizes J_M, we obtain

$$\begin{aligned} \lambda_M = \arg\min_{\lambda \in \mathbb{R}^L} J_M(\lambda) &\subset \{\lambda \mid J_M(\lambda) \le J_M(0)\} \\ &\subset \left\{\lambda \,\middle|\, 1 - E^{(Q)}_{T,M}\|\lambda\| + \frac{1}{2}C\|\lambda\|^2 \le 1\right\}. && (\because \text{(3.10)}, J_M(0) = 1) \end{aligned}$$

Solving the inequality, we obtain $\|\lambda_M\| \le \frac{2E^{(Q)}_{T,M}}{C}$. (3.20) follows from (3.21). ∎

Since $-\log(1-t) \approx t < 2t$ for small $t > 0$, the error bound (3.15) is tighter than (3.19) for large M by setting $t = \frac{(E^{(Q)}_{T,M})^2}{2C}$.

Theorems 4 and 6 show that the order of the theoretical error of our approximation method is at most that of the initial quadrature formula, but does not say that the error actually improves. Next we show that our method improves the accuracy of the integration in some situation.

Assume that the density f has a compact support. Let $T(x) = (T_1(x), \ldots, T_L(x))$ be the moment defining function, which are bounded on $\operatorname{supp} f$. Consider approximating the integrand g using the components of T as basis functions:

$$g(x) \approx b_{g,T}(x) = \sum_{l=1}^{L} \beta_l T_l(x),$$

where $\beta_1, \ldots, \beta_L$ are coefficients. Let

$$r_{g,T} = \frac{g(x) - b_{g,T}(x)}{\|g - b_{g,T}\|_\infty} \tag{3.22}$$

be the normalized residual of the approximation. Clearly $\|r_{g,T}\|_\infty \le 1$.

Theorem 7. *Let everything be as in Theorem 6 and* $\operatorname{supp} f$ *be compact. Then*

$$E^{(P)}_{g,M} \le \|g - b_{g,T}\|_\infty \left(E^{(Q)}_{r_{g,T},M} + \frac{2}{\sqrt{C}} E^{(Q)}_{T,M} \right). \tag{3.23}$$

Proof. Since the moments $T_1, \dots, T_L$ are exact under $P_M = \{p_{i,M}\}$, we have $E^{(P)}_{b_{g,T},M} = 0$. Therefore by the triangle inequality we obtain

$$\begin{aligned} E^{(P)}_{g,M} &\le E^{(P)}_{g-b_{g,T},M} + E^{(P)}_{b_{g,T},M} = E^{(P)}_{g-b_{g,T},M} \\ &\le E^{(Q)}_{g-b_{g,T},M} + \|g - b_{g,T}\|_\infty \frac{2}{\sqrt{C}} E^{(Q)}_{T,M} && (\because \text{Theorem 6}) \\ &= \|g - b_{g,T}\|_\infty \left(E^{(Q)}_{r_{g,T},M} + \frac{2}{\sqrt{C}} E^{(Q)}_{T,M} \right), && (\because (3.22)) \end{aligned}$$

which is (3.23). ∎

A similar bound can be obtained if we apply Theorem 4 instead of 6. Theorem 7 shows that by matching the moments $T(x)$, the error improves by the factor $\|g - b_{g,T}\|_\infty$, which is the approximation error of the integrand g by a linear combination of the moments $T(x)$.

For example, suppose that the problem is one dimensional with $\operatorname{supp} f = [c, d]$ and we match the polynomial moments by setting the moment defining function $T_l(x) = x^l$ ($l = 1, \dots, L$). Then for $b_{g,T}$ above we can adopt the Chebyshev interpolating polynomial of g, which is an L-degree polynomial that coincides with g at the points $x_j = \frac{c+d}{2} + \frac{d-c}{2} \cos(j\pi/L)$, where $j = 0, 1, \dots, L$. The Chebyshev interpolating polynomial can be easily computed and is known to be a nearly optimal approximating polynomial of a continuous function on a finite interval [47, Ch. 16].[7]

4. Numerical experiments. In this section, we present some numerical examples that compare the accuracy of the approximate expectations computed by an initial quadrature formula and its modifications by our proposed method. All computations in this section are done by MATLAB programs with double precision floating point arithmetic on a PC.

4.1. Beta and uniform distributions. For simplicity, we consider continuous probability distributions on the finite interval $[0, 1]$, specifically the beta distributions and the uniform distribution. The exact density function of the beta distribution is

$$f(x) = x^{a-1}(1 - x)^{b-1}/B(a, b),$$

where $B(\cdot, \cdot)$ is the beta function. We set $(a, b) = (1, 3)$ and $(2, 4)$ in the following.

4.1.1. Initial discretization and implementation. We adopt the trapezoidal formula and Simpson's formula as initial quadrature formulas. Consider the discrete set $D_M = \{m h_M \mid m = 0, 1, \dots, 2M\}$, where $h_M = \frac{1}{2M}$ is the distance between the

[7] Although the theoretical error estimate of the Chebyshev interpolation is well-known [47, Ch. 7], we do not use the estimate in the next section because it is not so tight to explain the improvement by our method when L is small. Instead we use the actual computed values of the error $\|g - b_{g,T}\|_\infty$ as an improvement factor.

points and $M = 1, 2, \dots, 12$. The number of points is $I_M = 2M+1$ and the integration points are $x_{i,M} = (i-1)h_M$, where $i = 1, \dots, I_M$. The integration weights are

$$\text{Trapezoidal:} \qquad w_{i,M} = \begin{cases} h_M, & (i \neq 1, I_M) \\ h_M/2, & (i = 1, I_M) \end{cases}$$

$$\text{Simpson:} \qquad w_{i,M} = \begin{cases} 4h_M/3, & (i \neq 1, I_M \text{ and } i \text{ is even}) \\ 2h_M/3, & (i \neq 1, I_M \text{ and } i \text{ is odd}) \\ h_M/3. & (i = 1, I_M) \end{cases}$$

Below, we compute the approximate expectation $\mathrm{E}[g(X)]$ of a test function $g(x)$ using eight formulas: the trapezoidal formula, Simpson's formula, and their modifications by our proposed method with exact polynomial moments $\mathrm{E}[X^l]$ up to 2nd order ($l = 1, 2$), 4th order ($l = 1, \dots, 4$), and 6th order ($l = 1, \dots, 6$).

We numerically solve the dual problem (D) as follows. First, note that in order for (P) to have a solution, it is necessary that there are at least as many unknown variables ($p_{i,M}$'s, so in total I_M) as the number of constraints (L moment constraints and $+1$ for probabilities to add up to 1, so $L+1$). Thus we need $I_M \geq L+1$.[8] A sufficient condition for the existence of a solution is $\bar{T} \in \operatorname{co} T(D)$ (Theorem 1), which can be easily verified in the current application.

Second, note that (D) is equivalent to the minimization of $J_M(\lambda)$ in (3.6), which is a strictly convex function of λ. In order to minimize J_M, we apply a variant of the Newton-Raphson algorithm. Starting with $\lambda_0 = 0$, we iterate

$$\lambda_{n+1} = \lambda_n - [\kappa I + \nabla^2 J_M(\lambda_n)]^{-1} \nabla J_M(\lambda_n) \tag{4.1}$$

over $n = 0, 1, \dots,$ where $\kappa > 0$ is a small number, I is the L-dimensional identity matrix, and $\nabla J_M, \nabla^2 J_M$ denote the gradient and the Hessian of J_M. Such an algorithm is advocated in [32]. The Newton-Raphson algorithm corresponds to setting $\kappa = 0$ in (4.1). Since the Hessian $\nabla^2 J_M$ is often nearly singular, the presence of $\kappa > 0$ stabilizes the iteration (4.1). Below we set $\kappa = 10^{-7}$ and terminate the iteration (4.1) when $\|\lambda_{n+1} - \lambda_n\| < 10^{-10}$.

4.1.2. Results for the beta distributions. For the test function, we pick $g(x) = \mathrm{e}^x$ for $x \in [0, 1]$. The exact expectations are $\mathrm{E}[g(X)] = 3(-5 + 2\mathrm{e})$ for $X \sim \mathrm{Be}(1, 3)$ and $\mathrm{E}[g(X)] = 20(49 - 18\mathrm{e})$ for $X \sim \mathrm{Be}(2, 4)$.

Figure 4.1 shows the results. According to the figures, our proposed method excels the trapezoidal and Simpson's formula in the accuracy. The errors basically decrease as the order of the matched moments increases, consistent with Theorem 7. Note that the curves marked "Theoretical" in Figure 4.1 express not exact error estimate but only the order of the theoretical error ($\mathrm{O}(M^{-2})$ for trapezoidal and $\mathrm{O}(M^{-4})$ for Simpson), which is same for Figures 4.3 and 4.4 below.

The reason why the error curves for our proposed method are not necessarily parallel to those for the initial quadrature formula when M is very small is because when the number of constraints $L+1$ is large relative to the number of unknown variables I_M, the method generates pathological probability distributions at the expense of matching many moments. For instance, Figure 4.2 shows the discrete approximations

[8] Since the beta density is zero at $x = 0, 1$, which are included in $x_{i,M}$'s, we necessarily have $p(x_{i,M}) = 0$ for $i = 1, I_M$. Thus, the number of unknown variables is $I_M - 2$, so we need $I_M - 2 \geq L + 1 \iff I_M \geq L + 3$ in the case of the beta distribution.

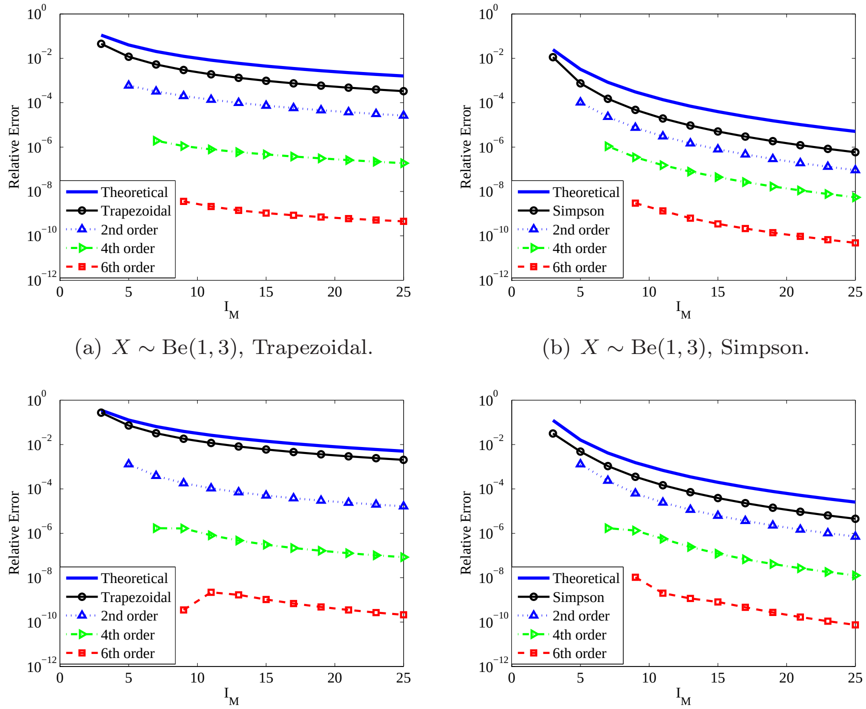

(a) $X \sim \mathrm{Be}(1,3)$, Trapezoidal. (b) $X \sim \mathrm{Be}(1,3)$, Simpson.

(c) $X \sim \mathrm{Be}(2,4)$, Trapezoidal. (d) $X \sim \mathrm{Be}(2,4)$, Simpson.

FIG. 4.1. *Relative errors of the computed values of* $\mathrm{E}[g(X)]$*, where* $g(x) = \mathrm{e}^x$*. The legends "2nd order", "4th order", and "6th order" represent those by our method with exact polynomial moments* $\mathrm{E}[X^l]$ *up to 2nd order* ($l = 1, 2$)*, 4th order* ($l = 1, \dots, 4$)*, and 6th order* ($l = 1, \dots, 6$)*, respectively.*

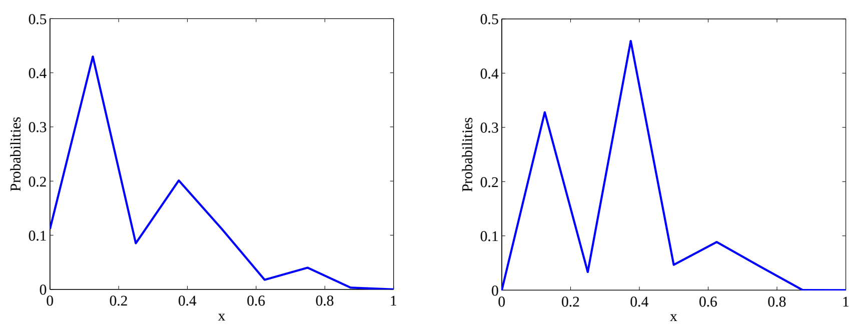

(a) $\mathrm{Be}(1,3)$. (b) $\mathrm{Be}(2,4)$.

FIG. 4.2. *6th order discrete approximation from the trapezoidal formula with* $I_M = 9$ *grid points.*

of the beta distributions $\mathrm{Be}(1,3)$ and $\mathrm{Be}(2,4)$ with $L = 6$ exact polynomial moments and $I_M = 9$ grid points. Clearly the discrete approximations do not resemble the continuous counterparts. This pathological behavior is rarely an issue, however. As long as there are twice as many grid points as constraints ($I_M \geq 2(L+1)$), the discrete approximation is well-behaved.

4.1.3. Results for the uniform distribution. We choose the test functions $g_1(x) = x^{\frac{9}{2}}$, $g_2(x) = \frac{1}{1+x}$, $g_3(x) = \sin(\pi x)$, and $g_4(x) = \log(1+x)$ for the uniform distribution on $[0,1]$.

Figures 4.3 and 4.4 show the numerical results. In all cases, our method excels the initial quadrature formula. The improvement in the accuracy is significant (of the order 10^{-4} for the trapezoidal formula and 10^{-2} for Simpson's formula), consistent with Theorem 7.

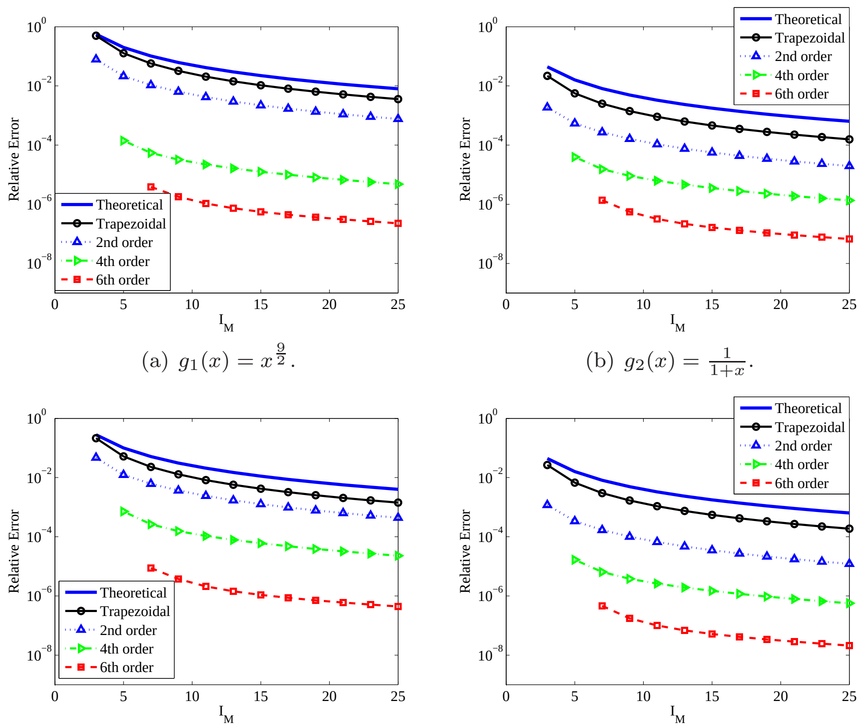

(a) $g_1(x) = x^{\frac{9}{2}}$. (b) $g_2(x) = \frac{1}{1+x}$.

(c) $g_3(x) = \sin(\pi x)$. (d) $g_4(x) = \log(1+x)$.

FIG. 4.3. *Relative errors of the computed values of $\int_0^1 g(x)\,\mathrm{d}x$. The legends "2nd order", "4th order", and "6th order" represent those by our method with exact polynomial moments $\mathrm{E}[X^l]$ up to 2nd order ($l = 1,2$), 4th order ($l = 1,\dots,4$), and 6th order ($l = 1,\dots,6$), respectively.*

4.1.4. Discussion. According to Figures 4.1, 4.3, and 4.4, the rate of convergence (the slope of the relative error with respect to the number of grid points I_M) is almost the same for the initial quadrature formula and our method. Specifically, the curves of the relative error are quite similar to the translations of the graphs of the theoretical errors, which are $\mathrm{O}(M^{-2})$ for the trapezoidal formula and $\mathrm{O}(M^{-4})$ for Simpson's formula when the integrands are in $C^2[0,1]$ and $C^4[0,1]$, respectively [14, Ch. 2]. This observation is consistent with Theorem 7, which shows that the error estimate is of the same order as the quadrature formula but improves by the error of the Chebyshev approximation.

Table 4.1 shows the errors of the Chebyshev approximations of the functions $g, g_1, \dots, g_4$. Since the trapezoidal formula does not even match the second moment, the theoretical improvement of our method (in $\log_{10}$) should be the numbers in Table 4.1. The actual improvements in Figures 4.1 and 4.3 are in line with these numbers.

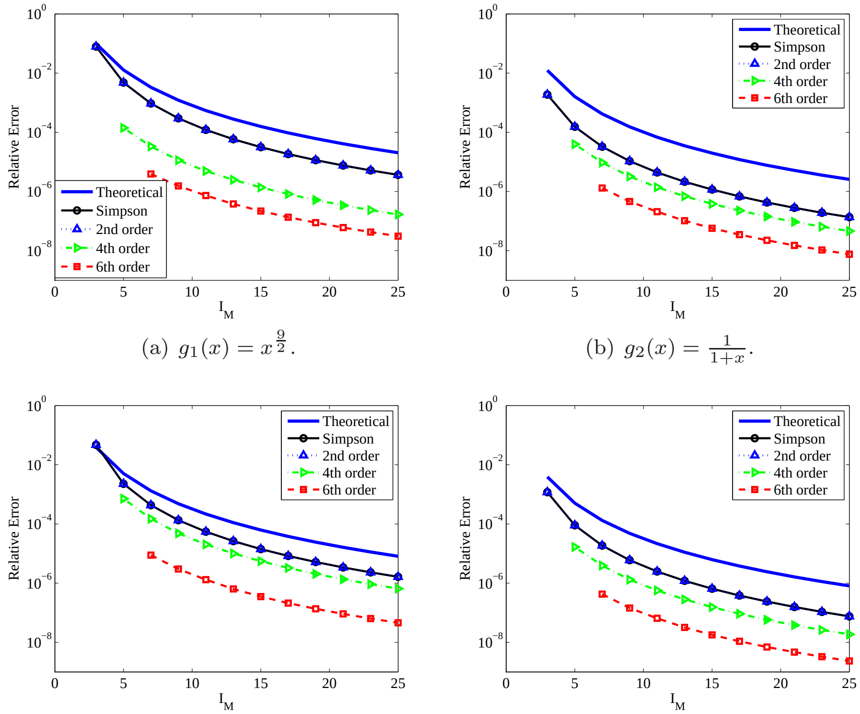

(a) $g_1(x) = x^{\frac{9}{2}}$.

(b) $g_2(x) = \frac{1}{1+x}$.

(c) $g_3(x) = \sin(\pi x)$.

(d) $g_4(x) = \log(1 + x)$.

FIG. 4.4. *Relative errors of the computed values of* $\int_0^1 g(x)\,\mathrm{d}x$. *The legends "2nd order", "4th order", and "6th order" represent those by our method with exact polynomial moments* $\mathrm{E}[X^l]$ *up to 2nd order* ($l = 1, 2$)*, 4th order* ($l = 1, \dots, 4$)*, and 6th order* ($l = 1, \dots, 6$)*, respectively.*

Since Simpson's formula gives exact quadratures for 2nd order polynomials (hence our method with $L = 2$ coincides with Simpson), the theoretical improvement of our method should be the *difference* between the numbers in Table 4.1 and those in the row corresponding to $L = 2$. For example, when we match $L = 6$ moments, the improvement should be approximately 10^{-3}, which is similar to what we observe in Figures 4.1 and 4.4.

TABLE 4.1
Errors of the Chebyshev approximations of $g, g_1, \dots, g_4$ *in* $\log_{10}$.

	Test function				
Degree	e^x	$x^{\frac{9}{2}}$	$(1+x)^{-1}$	$\sin(\pi x)$	$\log(1+x)$
2	−1.841	−0.847	−1.874	−1.251	−2.221
4	−4.285	−2.869	−3.363	−3.031	−3.918
6	−7.102	−5.048	−4.895	−4.872	−5.592

In summary, our method seems to be particularly suited for fine-tuning a quadrature formula with a small number of integration points. For instance, we can construct a highly accurate compound rule by subdividing the interval and applying our method to each subinterval.

4.2. Optimal portfolio problem. In this section we numerically solve the optimal portfolio problem briefly discussed in the introduction (see [42] for more details). Suppose that there are two assets, stock and bond, with gross returns R_1, R_2. Asset 1 (stock) is stochastic and lognormally distributed: $\log R_1 \sim N(\mu, \sigma^2)$, where μ is the expected return and σ is the volatility. Asset 2 (bond) is risk-free and $\log R_2 = r$, where r is the (continuously compounded) interest rate. The optimal portfolio θ is determined by the optimization

$$U = \max_{\theta} \frac{1}{1-\gamma} \mathrm{E}[(R_1\theta + R_2(1-\theta))^{1-\gamma}], \tag{4.2}$$

where $\gamma > 0$ is the relative risk aversion coefficient. We set the parameters such that $\gamma = 3$, $\mu = 0.07$, $\sigma = 0.2$, and $r = 0.01$. We numerically solve the optimal portfolio problem (4.2) in two ways, applying the trapezoidal formula and our proposed method. (We also tried Simpson's method but it was similar to the trapezoidal method.) To approximate the lognormal distribution, let $I_M = 2M+1$ be the number of grid points (M is the number of positive grid points) and $D_M = \{mh \mid m = 0, \pm 1, \dots, \pm M\}$, where $h = 1/\sqrt{M}$ is the grid size. Let $p(x)$ be the approximating discrete distribution of $N(0,1)$ as before (trapezoidal or the proposed method with various moments). Then we put probability $p(x)$ on the point $\mathrm{e}^{\mu+\sigma x}$ for each $x \in D_M$ to obtain the approximate stock return R_1.

Table 4.2 shows the optimal portfolio θ and its relative error for various moments L and number of points $I_M = 2M+1$. The result is somewhat surprising. Even with 3 approximating points ($M = 1$), our proposed method derives an optimal portfolio that is off by only 0.5% to the true value, whereas the trapezoidal method is off by 127%. While the proposed method virtually obtains the true value with 9 points ($M = 4$, especially when the 4th moment is matched), the trapezoidal method still has 23% of error.

Table 4.2

Optimal portfolio and relative error for the trapezoidal method and our method. M: number of positive grid points, $I_M = 2M+1$: total number of grid points, L: maximum order of moments.

# of grid points		$L = 0$ (trapezoidal)		$L = 2$		$L = 4$	
M	I_M	θ	Error (%)	θ	Error (%)	θ	Error (%)
$1^2 = 1$	3	1.5155	127	0.6717	0.54	-	-
$2^2 = 4$	9	0.8246	23.4	0.6694	0.20	0.6680	-0.015
$3^2 = 9$	19	0.6830	2.24	0.6684	0.044	0.6681	0
$4^2 = 16$	33	0.6687	0.088	0.6682	0.015	0.6681	0
$5^2 = 25$	51	0.6681	0	0.6681	0	0.6681	0

The reason why the trapezoidal method gives poor results when the number of approximating points is small is because the moments are not matched. To see this, taking the first-order condition for the optimal portfolio problem (4.2), we obtain $\mathrm{E}[(\theta X + R_2)^{-\gamma} X] = 0$, where $X = R_1 - R_2$ is the excess return on the stock. Using the Taylor approximation $x^{-\gamma} \approx a^{-\gamma} - \gamma a^{-\gamma-1}(x-a)$ for $x = \theta X + R_2$ and $a = \mathrm{E}[\theta X + R_2]$ and solving for θ, after some algebra we get

$$\theta = \frac{R_2 \mathrm{E}[X]}{\gamma \mathrm{Var}[X] - \mathrm{E}[X]^2}.$$

Therefore the (approximate) optimal portfolio depends on the first and second moments of the excess return X. Our method is accurate precisely because we match the moments. In complex economic problems, oftentimes we cannot afford to use many integration points, in which case our method might be useful to obtain accurate results.

Acknowledgments. We thank Jonathan Borwein, seminar participants at the 33rd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MaxEnt 2013) and 2014 Info-Metrics Institute Conference, and two anonymous referees for comments and feedback that greatly improved the paper. KT was partially supported by JSPS KAKENHI Grant Number 24760064.

Appendix A. Proof of Pinsker's inequality. This appendix proves Pinsker's inequality $\frac{1}{2}\left\|P-Q\right\|_1^2 \le H(P;Q)$, where $P=\{p_n\}_{n=1}^N$ and $Q=\{q_n\}_{n=1}^N$ are probability distributions and $\left\|\cdot\right\|_1$ denotes the L^1 norm. Let $N_+=\{n \mid p_n>q_n\}$, $N_-=\{n \mid p_n\le q_n\}$, $p=\sum_{n\in N_+}p_n$, and $q=\sum_{n\in N_+}q_n$. Then

$$\begin{aligned}\left\|P-Q\right\|_1=\sum_{n=1}^N\left|p_n-q_n\right|&=\sum_{n\in N_+}(p_n-q_n)-\sum_{n\in N_-}(p_n-q_n)\\&=2\sum_{n\in N_+}(p_n-q_n)-\sum_{n=1}^N(p_n-q_n)=2(p-q),\end{aligned}$$

where we have used $\sum_{n=1}^N p_n=\sum_{n=1}^N q_n=1$. By the convexity of $-\log(\cdot)$, we get

$$\begin{aligned}\sum_{n=1}^N p_n\log\frac{p_n}{q_n}&=p\sum_{n\in N_+}\frac{p_n}{p}\left(-\log\frac{q_n/p}{p_n/p}\right)+(1-p)\sum_{n\in N_-}\frac{p_n}{1-p}\left(-\log\frac{q_n/(1-p)}{p_n/(1-p)}\right)\\&\ge -p\log\left(\sum_{n\in N_+}\frac{q_n}{p}\right)-(1-p)\log\left(\sum_{n\in N_-}\frac{q_n}{1-p}\right)\\&=p\log\frac{p}{q}+(1-p)\log\frac{1-p}{1-q}.\end{aligned}$$

Therefore it suffices to show

$$h(p,q):=p\log\frac{p}{q}+(1-p)\log\frac{1-p}{1-q}-2(p-q)^2\ge 0.$$

Fix q and regard the left-hand side as a function of p alone. Then

$$\begin{aligned}h'(p,q)&=\log\frac{p}{q}-\log\frac{1-p}{1-q}-4(p-q),\\h''(p,q)&=\frac{1}{p}+\frac{1}{1-p}-4=\frac{(1-2p)^2}{p(1-p)}\ge 0,\end{aligned}$$

so h is convex. Clearly $h'(q,q)=0$, so it follows that $h(p,q)\ge h(q,q)=0$.

REFERENCES

[1] Rafail V. Abramov. An improved algorithm for the multidimensional moment-constrained maximum entropy problem. *Journal of Computational Physics*, 226(1):621–644, September 2007.

[2] Rafail V. Abramov. The multidimensional maximum entropy moment problem: A review of numerical methods. *Communications in Mathematical Sciences*, 8(2):377–392, 2010.
[3] Jérôme Adda and Russel W. Cooper. *Dynamic Economics: Quantitative Methods and Applications*. MIT Press, Cambridge, MA, 2003.
[4] S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684, 1994.
[5] Graham W. Alldredge, Cory D. Hauck, Dianne P. O'Leary, and André L. Tits. Adaptive change of basis in entropy-based moment closures for linear kinetic equations. *Journal of Computational Physics*, 258(1):489–508, February 2014.
[6] Andrew R. Barron and Chyong-Hwa Sheu. Approximation of density functions by sequences of exponential families. *Annals of Statistics*, 19(3):1347–1369, 1991.
[7] Jonathan M. Borwein and Adrian S. Lewis. Duality relationships for entropy-like minimization problems. *SIAM Journal on Control and Optimization*, 29(2):325–338, March 1991.
[8] Jonathan M. Borwein and Adrian S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Canadian Mathematical Society Books in Mathematics. Springer, New York, 2nd edition, 2006.
[9] Peter W. Buchen and Michael Kelly. The maximum entropy distribution of an asset inferred from option prices. *Journal of Financial and Quantitative Analysis*, 31(1):143–159, 1996.
[10] Marko Budišić and Mihai Putinar. Conditioning moments of singular measures for entropy optimization. I. *Indagationes Mathematicae*, 23(4):848–883, December 2012.
[11] Ariel Caticha and Adom Giffin. Updating probabilities. In Ali Mohammad-Djafari, editor, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 872 of *AIP Conference Proceedings*, pages 31–42, 2006.
[12] Imre Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Scientiarum Mathematicarum Hungarica*, 2:299–318, 1967.
[13] Imre Csiszár. Sanov property, generalized I-projection and a conditional limit theorem. *Annals of Probability*, 12(3):768–793, August 1984.
[14] Philip J. Davis and Philip Rabinowitz. *Methods of Numerical Integration*. Academic Press, Orlando, FL, second edition, 1984.
[15] Andrée Decarreau, Danielle Hilhorst, Claude Lemaréchal, and Jorge Navaza. Dual methods in entropy maximization. Application to some problems in crystallography. *SIAM Journal on Optimization*, 2(2):173–197, 1992.
[16] Eric A. DeVuyst and Paul V. Preckel. Gaussian cubature: A practitioner's guide. *Mathematical and Computer Modelling*, 45(7-8):787–794, April 2007.
[17] Alexei A. Fedotov and Peter Harremoës. Refinements of Pinsker's inequality. *IEEE Transactions on Information Theory*, 49(6):1491–1498, June 2003.
[18] Duncan K. Foley. A statistical equilibrium theory of markets. *Journal of Economic Theory*, 62(2):321–345, April 1994.
[19] Nikolay Gospodinov and Damba Lkhagvasuren. A moment-matching method for approximating vector autoregressive processes by finite-state Markov chains. *Journal of Applied Econometrics*, 29(5):843–859, August 2014.
[20] Cory D. Hauck, C. David Levermore, and André L. Tits. Convex duality and entropy-based moment closures: Characterizing degenerate densities. *SIAM Journal on Control and Optimization*, 47(4):1977–2015, 2008.
[21] Kjetil Høyland and Stein W. Wallace. Generating scenario trees for multistage decision problems. *Management Science*, 47(2):295–307, February 2001.
[22] Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969, September-November 1993.
[23] Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, May 1957.
[24] Edwin T. Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, 1982.
[25] Edwin T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, U.K., 2003. Edited by G. Larry Bretthorst.
[26] Donald L. Keefer and Samuel E. Bodily. Three-point approximations for continuous random variables. *Management Science*, 29(5):595–609, May 1983.
[27] Yuichi Kitamura and Michael Stutzer. An information-theoretic alternative to generalized method of moments estimation. *Econometrica*, 65(4):861–874, July 1997.
[28] Kevin H. Knuth and John Skilling. Foundations of inference. *Axioms*, 1(1):38–73, 2012.
[29] Per Krusell and Anthony A. Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, October 1998.
[30] Solomon Kullback. A lower bound for discrimination information in terms of variation. *IEEE*

Transactions on Information Theory, 13(1):126–127, January 1967.
[31] Solomon Kullback and Richard A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, March 1951.
[32] David G. Luenberger and Yinyu Ye. *Linear and Nonlinear Programming*. International Series in Operations Research and Management Science. Springer, NY, third edition, 2008.
[33] Lawrence R. Mead and Nikos Papanicolaou. Maximum entropy in the problem of moments. *Journal of Mathematical Physics*, 25(8):2404–2417, August 1984.
[34] Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, December 1971.
[35] Allen C. Miller, III and Thomas R. Rice. Discrete approximations of probability distributions. *Management Science*, 29(3):352–362, March 1983.
[36] Georg Ch. Pflug. Scenario tree generation for multiperiod financial optimization by optimal discretization. *Mathematical Programming*, 89(2):251–271, January 2001.
[37] Paul A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3):239–246, August 1969.
[38] John E. Shore and Rodney W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, 26(1):26–37, January 1980.
[39] James E. Smith. Moment methods for decision analysis. *Management Science*, 39(3):340–358, March 1993.
[40] Michael Stutzer. A Bayesian approach to diagnosis of asset pricing models. *Journal of Econometrics*, 68(2):367–397, August 1995.
[41] Michael Stutzer. A simple nonparametric approach to derivative security valuation. *Journal of Finance*, 51(5):1633–1652, December 1996.
[42] Ken'ichiro Tanaka and Alexis Akira Toda. Discrete approximations of continuous distributions by maximum entropy. *Economics Letters*, 118(3):445–450, March 2013.
[43] George Tauchen. Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986.
[44] George Tauchen and Robert Hussey. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica*, 59(2):371–396, March 1991.
[45] Alexis Akira Toda. Existence of a statistical equilibrium for an economy with endogenous offer sets. *Economic Theory*, 45(3):379–415, 2010.
[46] Alexis Akira Toda. Bayesian general equilibrium. *Economic Theory*, 58(2):375–411, 2015.
[47] Lloyd N. Trefethen. *Approximation Theory and Approximation Practice*. SIAM, Philadelphia, 2013.
[48] Jan M. Van Campenhout and Thomas M. Cover. Maximum entropy and conditional probability. *IEEE Transactions on Information Theory*, 27(4):483–489, July 1981.
[49] Ximing Wu. Calculation of maximum entropy densities with application to income distribution. *Journal of Econometrics*, 115(2):347–354, August 2003.