

Epi-convergent discretizations of stochastic programs via integration quadratures

Teemu Pennanen* and Matti Koivu
Department of Management Science
Helsinki School of Economics
PL 1210,
00101 Helsinki
Finland
[pennanen, koivu]@hkkk.fi

July 25, 2003

Abstract

Modern integration quadratures are designed to produce finitely supported approximations of a given (probability) measure. This makes them well suited for discretization of stochastic programs. We give conditions that guarantee the epi-convergence of resulting objectives to the original one. Our epi-convergence result is closely related to some of the existing ones but it is easier to apply to discretizations. As examples, we will verify the conditions for discretizations of three different models of portfolio management and we study the behavior of various discretizations numerically. In our tests, modern quadratures clearly outperform crude Monte Carlo sampling in discretization of stochastic programs.

1 Introduction

Let X and Ξ be complete separable metric spaces, and Σ the Borel σ -algebra on Ξ . Let P be a probability measure on (Ξ, Σ) , and f an extended real-valued function on $X \times \Xi$, with $f(x, \cdot)$ measurable for every $x \in X$. Consider the problem

$$\underset{x \in X}{\text{minimize}} \quad E^P f(x) := \int_{\Xi} f(x, \xi) P(d\xi), \quad (SP)$$

where the integral is interpreted as $+\infty$ when $f(x, \cdot) \notin L^1(\Xi, \Sigma, P)$. The decision variable x is not a function of ξ , so (SP) represents a *static* (one-stage) stochastic program. We emphasize that by allowing f to take on the value $+\infty$ we can incorporate constraints into the objective, which makes (SP) a very general model for static decision making problems under uncertainty.

In solving (SP) , one often replaces P by a finitely supported measure of the form

$$P_\nu = \sum_{i=1}^{\nu} p_i^\nu \delta_{\xi_i^\nu},$$

*The work of this author was partially supported by The Finnish Foundation for Economic Education under grant number 21599

where $\delta_{\xi_i^\nu}$ denotes the unit mass located at ξ_i^ν . This results in the problem

$$\underset{x \in X}{\text{minimize}} \quad E^{P^\nu} f(x) := \sum_{i=1}^{\nu} p_i^\nu f(x, \xi_i^\nu), \quad (SP^\nu)$$

which, for moderate values of ν , is usually easier to solve than (SP) . Several approaches for constructing the measures P^ν have been considered in the literature. In general, the aim is to choose P^ν so that (SP^ν) is a good approximation of (SP) and that the number ν of support points of P^ν is small enough to allow for numerical solution of (SP^ν) . The simplest and perhaps the most popular choice in stochastic programming applications is to use empirical measures of the form $P^\nu = \sum_{i=1}^{\nu} \frac{1}{\nu} \delta_{\xi_i^\nu}$, where $\{\xi_i^\nu\}_{i=1}^{\nu}$ is a random sample from P . It is clear that a random sample can lead to a bad approximation of P , which in turn, may lead to an equally bad approximation of the optimization problem. Indeed, in practice, the optimal values and solutions corresponding to empirical measures are often highly dependent on the sample. There have been attempts to improve the accuracy of crude Monte Carlo sampling by using ideas from importance sampling technique; see Infanger [19] and Dempster and Thompson [10]. Høyland and Wallace [18] proposed to use moment matching where P^ν is constructed so that it has the first few moments of the original distribution; see also [17]. In barycentric approximation, one constructs P^ν so that, under some convexity properties of the function $f(\cdot, \cdot)$, the optimum value of (SP^ν) provides an upper/lower bound to that of (SP) ; see Frauendorfer [14]. Pflug [28] proposed to construct discrete measures P^ν so that they are as close as possible to P in the sense of the so called Wasserstein-distance.

This paper studies the use of modern integration quadratures in constructing the discretizations (SP^ν) . Such quadratures have the attractive feature that they have been designed to give discrete measures that approximate a given measure as well as possible. Moreover, they are just as easy to use as the Monte Carlo method and they are fast compared to alternative approaches such as barycentric approximation, moment matching or the method in [28]. We study the corresponding discretizations both analytically and numerically.

Since we are dealing with minimization problems, a natural framework for analyzing approximations is *epi-convergence* of the objectives; see Attouch [3] or Rockafellar and Wets [31] for introduction to epi-convergence. Epi-convergence of stochastic programs with respect to perturbations in the probability measure has been studied, for example, by Birge and Wets [8], Robinson and Wets [30], Dupacova and Wets [11], Kall, Ruszczyński, and Frauendorfer [20], Lucchetti and Wets [26], Artstein and Wets [1, 2], Zervos [35], and Schultz [32]. In these studies, *weak convergence* of the approximating measures P^ν to the original measure P has been found an important property. In numerical integration, weak convergence corresponds to *consistency* which is a minimal requirement for any integration quadrature. We derive an epi-convergence result which is closely related to the ones in the above references but it is easier to apply to discretizations. As examples, we verify these conditions in three different models of portfolio management. Numerical tests on these models clearly favor modern quadratures over crude Monte Carlo. Indeed, they seem to result in faster convergence of the approximations thus allowing for smaller values of ν and cheaper computations.

The rest of this paper is organized as follows. Section 2 gives a brief review of modern integration quadratures and their use in generation of weakly convergent probability measures. In Section 3, we derive an epi-convergence result for $E^{P^\nu} f$. In Section 4 we combine results from Sections 2 and 3 to construct epi-convergent discretizations of some particular problems, and we study the stability of the corresponding optimal values numerically.

2 Constructing weakly convergent probability measures

Based on the importance of weak convergence in studying epi-convergence of stochastic programs, it is natural to try to choose the measures P_ν in (SP^ν) so that they converge weakly to P as $\nu \nearrow \infty$. Recall that weak convergence, which will be denoted by $P_\nu \rightarrow P$, means that

$$E^{P_\nu} \varphi \rightarrow E^P \varphi, \quad (1)$$

for all bounded and continuous functions φ ; see Billingsley [7]. The literature of numerical integration is full of methods for generating such sequences.

2.1 Low dimensional spaces: Gaussian quadratures

Gaussian quadratures are usually very efficient in one-dimensional integration [29]. For different choices of the integration limits a and b and a weight (density) function w , they yield approximations

$$\int_a^b \varphi(\eta)w(\eta)d\eta \approx \sum_{i=1}^{\nu} w_i^{\nu} \varphi(\eta_i^{\nu}), \quad (2)$$

where the quadrature points η_i^{ν} and weights w_i^{ν} are chosen so that the quadrature has an optimal order of accuracy: *a ν point Gaussian quadrature is exact for all polynomials of degree $2\nu - 1$ or less.* Given a , b and w , the values of η_i and w_i can be computed numerically. For certain choices of a , b and w , the computation of the quadrature points and weights is particularly easy.

In the case $a = -\infty$, $b = \infty$, $w(\eta) = \exp(-\eta^2)$, (2) is known as *Gauss-Hermite quadrature*. A C-routine for computing the the points η_i^{ν} and the weights w_i^{ν} of the Gauss-Hermite quadrature can be found in [29]. Gauss-Hermite quadrature can be used to approximate the expectation under the normal distribution P as

$$E^P \varphi = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{\xi^2}{2}} \varphi(\xi) d\xi = \int_{-\infty}^{\infty} \frac{1}{\sqrt{\pi}} e^{-\eta^2} \varphi(\sqrt{2}\eta) d\eta \approx \sum_{i=1}^{\nu} \frac{1}{\sqrt{\pi}} w_i^{\nu} \varphi(\sqrt{2}\eta_i^{\nu}).$$

It turns out that $w_i^{\nu} > 0$ and $\sum_{i=1}^{\nu} \frac{w_i^{\nu}}{\sqrt{\pi}} = 1$, the latter following from the fact that the quadrature is exact for constant functions. Thus,

$$P_{\nu} := \sum_{i=1}^{\nu} \frac{w_i^{\nu}}{\sqrt{\pi}} \delta_{\sqrt{2}\eta_i^{\nu}}$$

defines a probability measure which satisfies $E^{P_{\nu}} \varphi = E^P \varphi$ for all polynomials φ of order $2\nu - 1$ or less. (In particular, P_{ν} matches the first $2\nu - 1$ moments of the normal distribution; compare with [17]). Weierstrass approximation theorem then implies that the measures P_{ν} converge weakly to P as $\nu \nearrow \infty$.

Gaussian quadratures do not directly extend to higher dimensions. The generation of integration quadratures of optimal order in arbitrary dimension is an open problem [29]. If a random variable can be expressed as $\xi = G(\zeta)$, where G is almost everywhere continuous, and $\zeta = (\zeta_1, \dots, \zeta_d)$ for independent random variables ζ_i with densities w_i , we can approximate the distribution of each ζ_i with a Gaussian quadrature to get discrete measures Q_{ν}^i , and then construct the measure

$$P_{\nu} = (Q_{\nu}^1, \dots, Q_{\nu}^d) G^{-1}; \quad (3)$$

see Theorem 1 below. Such approximations work well in dimensions 1, 2 or 3, but in higher dimensions they suffer from the ‘‘curse of dimensionality’’: if one approximates each Q_i with an k -point quadrature, a 10-dimensional approximation of the above form would have k^{10} quadrature points. Fortunately, there are better integration quadratures for high dimensional spaces.

2.2 Higher dimensions: low discrepancy point sets and the method of inversion

In the scalar case, a usual thing to do is to approximate the uniform distribution on $[0, 1]$ and to transform each point with the inverse of the distribution function of the desired distribution. This is known as the *method of inversion*. The same idea works whenever $P = QG^{-1}$, where Q is the multivariate uniform distribution and G is Q -a.s. continuous, in other words, whenever

$$\xi = G(u),$$

where u is uniformly distributed in the unit cube $[0, 1]^d$, and $G : [0, 1]^d \rightarrow \Xi$ is almost everywhere continuous. This is based on the following very useful result from Billingsley [7] where U is any metric space with Borel algebra \mathcal{B} .

Theorem 1 (Billingsley) *Let $G : (U, \mathcal{B}) \rightarrow (\Xi, \Sigma)$ be a measurable function and Q a probability distribution on (U, \mathcal{B}) . Then $QG^{-1}(A) := Q(G^{-1}A)$ defines a probability measure on (Ξ, Σ) , and if G is Q -a.s. continuous, then*

$$Q_\nu \rightarrow Q \implies Q_\nu G^{-1} \rightarrow QG^{-1}.$$

Given a Q -a.s. continuous G and a discrete approximation $Q_\nu = \sum_{i=1}^\nu p_i^\nu \delta_{u_i^\nu}$ of Q , Theorem 1 says that the discrete measures

$$P_\nu := Q_\nu G^{-1} = \sum_{i=1}^\nu p_i^\nu \delta_{G(u_i^\nu)}$$

converge weakly to $P = QG^{-1}$ whenever $Q_\nu \rightarrow Q$. It is then natural to try to choose discrete approximations Q_ν which are as close as possible to the uniform distribution Q . Modern methods of numerical integration do exactly this; see the books of Niederreiter [27] and Sloan and Joe [33]. Much of this theory has evolved around the following notion of distance from Q .

Definition 2 *The star-discrepancy of a point set $\{u_1, \dots, u_\nu\} \subset [0, 1]^d$ is defined as*

$$D^*(u_1, \dots, u_\nu) = \sup_{C \in \mathcal{C}_0} |Q_\nu(C) - Q(C)|, \quad (4)$$

where

$$Q_\nu = \sum_{i=1}^\nu \frac{1}{\nu} \delta_{u_i},$$

and \mathcal{C}_0 is the set of rectangles $C \subset [0, 1]^d$ with $0 \in C$.

The following is a direct consequence of Corollary 11 in Lucchetti, Salinetti and Wets [25].

Proposition 3 *For each ν , let $\{u_1^\nu, \dots, u_\nu^\nu\}$ be point sets in the unit cube. The measures*

$$Q_\nu = \sum_{i=1}^\nu \frac{1}{\nu} \delta_{u_i^\nu}$$

converge weakly to the uniform distribution if and only if $D^*(u_1^\nu, \dots, u_\nu^\nu) \rightarrow 0$.

Thus, if we can find point sets whose star-discrepancy approaches zero as $\nu \nearrow \infty$, we obtain weakly convergent discrete approximations of the uniform distribution. If $P = QG^{-1}$, we can then use the method of inversion to get weakly convergent discretizations of P . In the literature of numerical integration, many methods have been proposed that are aimed at producing point sets that have as small star-discrepancy as possible. It is thus natural to employ them in the construction of discrete measures P^ν and the corresponding approximations (SP^ν). This is what the present paper is about.

This approach to discretization of stochastic programs is close in spirit to the method proposed in Pflug [28], where the aim is to find discrete measures P^ν that are as close as possible to P in the sense of the so called *Wasserstein-distance*. In general, the problem of finding a discrete measure that minimizes a distance from a given measure can be very hard. Fortunately, in the case of star-discrepancy, many efficient methods are already available.

Example 4 (low discrepancy sequences) *Low discrepancy sequences are infinite sequences whose first ν points have low discrepancy for all ν . Examples are*

1. *Faure sequence [12]. A FORTRAN 77-routine for Faure sequence has been implemented by Fox as ACM Algorithm 647 [13].*
2. *Sobol sequence [34]. A C-routine for Sobol sequence is available in GSL (Gnu Scientific Library, www.gnu.org/software/gsl/gsl.html).*

3. Niederreiter sequence [9]. This is also available in GSL.

These satisfy

$$D^*(u_1, \dots, u_\nu) \leq C \frac{(\log \nu)^d}{\nu} \quad \forall \nu, \quad (5)$$

for a constant C independent of ν . These examples fall in the general class of (t, s) -sequences; see [27, Chapter 4]. Figure 1 displays the first 15 and 127 points for Faure and Sobol sequences in the 2-dimensional unit cube.

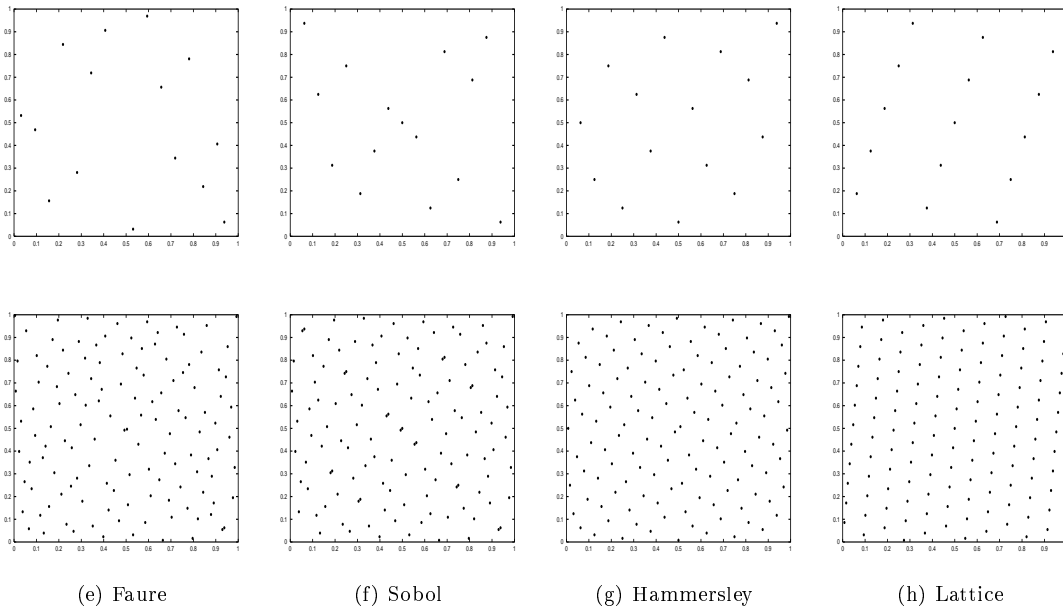


Figure 1: Discretizations of the uniform distribution by 15 and 127 points

In direct numerical integration, infinite low discrepancy sequences are useful in that after evaluating a ν -point quadrature, one can continue to compute the next $(\nu + 1)$ -point quadrature simply by evaluating the function at one new point. In stochastic programming this advantage is lost since, in general, the solution x and thus the integrand $f(x, \cdot)$ changes every time a new point (scenario) is added to the problem. This raises the question whether it is possible to obtain more accurate quadratures if it is not required that ν points of a $(\nu + 1)$ -point quadrature are the points of the ν -point quadrature. This is indeed possible.

Example 5 (low discrepancy point sets) A set of points $\{u_1, \dots, u_\nu\}$ in the unit cube is called a low discrepancy point set if it has low discrepancy. Examples are

1. Hammersley point sets [16]. Hammersley point sets can be obtained quite easily from the Halton sequence [15] that has been implemented in [13].
2. (t, m, s) -nets are a general class of low discrepancy point sets that are described in detail in [27, Chapter 4].

These satisfy

$$D^*(u_1, \dots, u_\nu) \leq C \frac{(\log \nu)^{d-1}}{\nu}, \quad (6)$$

for a constant C independent of ν . Figure 1 displays 15 and 127 Hammersley points in a 2-dimensional unit cube.

There is another class of quadratures designed to take advantage of additional regularity properties of integrands.

Example 6 (lattice rules) *Lattice rules are a general family of methods for generating point sets with low discrepancy; see for example [27, Chapter 5], [33] and L'Ecuyer and Lemieux [24]. For each d and ν there exist lattice rules that give point sets satisfying (5); see [27, page 115]. This is not as good as (6), but for certain classes of functions much tighter error bounds can be derived.*

Figure 1 displays lattices of 15 and 127 points produced by the so called Korobov lattice rule [23]. Parameters required by the method were provided by Pierre L'Ecuyer and Christiane Lemieux (personal communication).

It can be shown that if $Q_\nu = \sum_{i=1}^{\nu} \frac{1}{\nu} \delta_{u_i^\nu}$, then

$$\left| E^{Q_\nu} \varphi - E^Q \varphi \right| \leq V(\varphi) D^*(u_1^\nu, \dots, u_\nu^\nu),$$

where $V(\varphi)$ is the variation of φ in the sense of Hardy and Krause; see [27, Section 2.2]. In direct integration, the above methods thus achieve the asymptotic convergence rate of ν^{-1} , whereas for pure Monte Carlo methods the rate is $\nu^{-\frac{1}{2}}$; more precisely, in Monte Carlo integration, the standard deviation of the integration error is $\sigma(\varphi)\nu^{-\frac{1}{2}}$, where $\sigma(\varphi)$ is the standard deviation of φ . It should be noticed, however, that in Monte Carlo, the error bound is independent of the dimension of the space, whereas the bounds in the above examples may depend on the dimension so that the actual error bound achieved in practice is much greater than $\sigma(\varphi)\nu^{-\frac{1}{2}}$. In numerical tests, however, low discrepancy point sets and sequences are often much more efficient than pure Monte Carlo; see for example [33]. In the tests of Section 4, one can see a similar effect in discretizations of stochastic programs.

3 Epi-convergence of the objectives

Given efficient procedures for constructing finitely supported measures P_ν that converge weakly to P , our next step is to find conditions that guarantee the epi-convergence of $E^{P_\nu} f$ to $E^P f$. Recall that the *lower epi-limit* of a sequence $\{F^\nu\}$ of functions is the lsc function given by

$$(\text{e-lim inf } F^\nu)(x) = \inf_{x_\nu \rightarrow x} \liminf_{\nu \rightarrow \infty} F^\nu(x_\nu)$$

and the *upper epi-limit* is the lsc function given by

$$(\text{e-lim sup } F^\nu)(x) = \inf_{x_\nu \rightarrow x} \limsup_{\nu \rightarrow \infty} F^\nu(x_\nu).$$

If $\text{e-lim inf } F^\nu = \text{e-lim sup } F^\nu$, then the common limit, denoted $\text{e-lim } F^\nu$, is called the *epi-limit* of $\{F^\nu\}$ and the sequence is said to *epi-converge* to it. Epi-convergence has many important implications in studying approximations of minimization problems. The following is one of them; see Attouch [3, Section 2.2].

Theorem 7 *If $\text{e-lim } F^\nu = F$, then*

$$\limsup_{\nu \rightarrow \infty} \inf F^\nu \leq \inf F,$$

and if there is a convergent sequence $x_k \rightarrow x$ such that $x_k \in \text{argmin } F^{\nu_k}$ for some subsequence $\{\nu_k\}_{k=1}^\infty$, then $x \in \text{argmin } F$ and $\inf F^{\nu_k} \rightarrow \inf F$. In particular, if there is a compact set C such that $\text{argmin } F^\nu \cap C \neq \emptyset$ for all ν , then $\inf F^\nu \rightarrow \inf F$.

Our proof of epi-convergence for (SP^ν) is based on ideas from Artstein and Wets [2], where $\{P_\nu\}_{\nu=1}^\infty$ was a sequence of empirical measures, and the main tools were the strong law of large numbers and an approximation algorithm due to Beer [4]. In our case, $\{P_\nu\}_{\nu=1}^\infty$ is a weakly convergent non-random sequence, and our main tools are the algorithm of Beer and Theorem 8 below. Recall that a sequence of functions $\{\varphi_\nu\}_{\nu=1}^\infty$ converges continuously to a function φ_0 at $\xi \in \Xi$ if

$$\varphi_\nu(\xi_\nu) \rightarrow \varphi_0(\xi),$$

whenever $\xi_\nu \rightarrow \xi$. The following is based on [1, Remark 4.3] and [32, pp. 67–68].

Theorem 8 *If $P_\nu \rightarrow P_0$, $\varphi_\nu \rightarrow \varphi_0$ continuously at P_0 -almost every $\xi \in \Xi$, and if for each $\epsilon > 0$ there exists a measurable set $K_\epsilon \subset \Xi$ and a bound $b_\epsilon \in \mathbb{R}$, such that for each $\nu = 0, 1, 2, \dots$*

1. $|\varphi_\nu(\xi)| \leq b_\epsilon$ for P_ν -almost every $\xi \in K_\epsilon$,
2. $\int_{\Xi \setminus K_\epsilon} |\varphi_\nu(\xi)| P_\nu(d\xi) < \epsilon$,

then

$$\lim_{\nu \rightarrow \infty} E^{P_\nu} \varphi_\nu = E^{P_0} \varphi_0.$$

Proof. Choose an $\epsilon > 0$ and let $b_\epsilon \in \mathbb{R}$ and $K_\epsilon \subset \Xi$ be the bound and the measurable set, respectively, provided by the last condition. For each ν , let φ_ν^ϵ be the bounded function whose value at a point ξ is the projection of $\varphi_\nu(\xi)$ to the interval $[-b_\epsilon, b_\epsilon]$. Then $|\varphi_\nu^\epsilon| \leq |\varphi_\nu|$, $\varphi_\nu^\epsilon(\xi) = \varphi_\nu(\xi)$ for P_ν -almost every $\xi \in K_\epsilon$, and $\varphi_\nu^\epsilon \rightarrow \varphi_0^\epsilon$ continuously P_0 -a.s. We have

$$|E^{P_\nu} \varphi_\nu - E^{P_0} \varphi_0| \leq |E^{P_\nu} \varphi_\nu - E^{P_\nu} \varphi_\nu^\epsilon| + |E^{P_\nu} \varphi_\nu^\epsilon - E^{P_0} \varphi_0^\epsilon| + |E^{P_0} \varphi_0^\epsilon - E^{P_0} \varphi_0|,$$

where the second term on the right converges to zero by [6, Theorem 5.5], and for $\nu = 0, 1, \dots$,

$$|E^{P_\nu} \varphi_\nu - E^{P_\nu} \varphi_\nu^\epsilon| = \left| \int_{\Xi \setminus K_\epsilon} [\varphi_\nu(\xi) - \varphi_\nu^\epsilon(\xi)] P_\nu(d\xi) \right| \leq 2 \int_{\Xi \setminus K_\epsilon} |\varphi_\nu(\xi)| P_\nu(d\xi) \leq 2\epsilon.$$

Since $\epsilon > 0$ was arbitrary, the result follows. \square

In particular, if φ is P_0 -a.s. continuous and bounded, then

$$E^{P_\nu} \varphi \rightarrow E^{P_0} \varphi.$$

Note that this is also implied directly by Theorem 1.

Recall that the *domain* of an extended real-valued function g is the set $\text{dom } g = \{x \mid g(x) < \infty\}$, and its *lower closure* is the function

$$(\text{cl } g)(x) = \liminf_{y \rightarrow x} g(y).$$

A function is called *lower semicontinuous* (lsc) if it is equal to its lower closure. We can now state our epi-convergence result for $E^{P_\nu} f$.

Theorem 9 *Let $P_\nu \rightarrow P_0$. If for each $x \in X$,*

1. *there is an open set $N \ni x$ such that f is bounded from below on $N \times \Xi$,*
2. $(\text{cl } f)(x, \cdot) = f(x, \cdot)$ P_0 -a.s.,

then

$$\text{e-lim inf } E^{P_\nu} f \geq E^P f.$$

If for each $x \in \text{dom } E^{P_0} f$,

3. *there is a sequence $x_\nu \rightarrow x$ such that P_ν and $\varphi_\nu := f(x_\nu, \cdot)$ satisfy the conditions of Theorem 8,*

then

$$\text{e-lim sup } E^{P_\nu} f \leq E^P f.$$

Proof. To verify the first claim, fix an $x \in X$ and let $x_\nu \rightarrow x$. According to [4] (see also the proof of [5, Theorem 7.14]), the first condition implies that there exists a sequence $\{f_k\}$ of bounded from below Lipschitz functions on $N \times \Xi$ such that $f_k \nearrow \text{cl} f$. The functions $g_k = \min\{f_k, k\}$ are also Lipschitz with $g_k \nearrow \text{cl} f$ but, in addition, they are bounded. Then

$$E^{P_\nu} f(x_\nu) \geq E^{P_\nu} g_k(x_\nu) \geq E^{P_\nu} g_k(x) - L_k d(x_\nu, x) \quad \forall \nu, k$$

where L_k is the Lipschitz constant of g_k and d is the metric on X . Since $P_\nu \rightarrow P$, we get

$$\liminf_{\nu \rightarrow \infty} E^{P_\nu} f(x_\nu) \geq \liminf_{\nu \rightarrow \infty} E^{P_\nu} g_k(x) = E^{P_0} g_k(x) \quad \forall k,$$

where $E^{P_0} g_k(x) \nearrow E^{P_0}(\text{cl} f)(x)$ as $k \nearrow \infty$ by the monotone convergence theorem. Since by the second condition, $E^{P_0}(\text{cl} f)(x) = E^{P_0} f(x)$, and since $x \in X$ and $x_\nu \rightarrow x$ were arbitrary, the claim follows.

For the second claim, it suffices to consider points $x \in \text{dom} E^{P_0} f$. The result then follows from the third condition and Theorem 8. \square

Note that choosing $P_\nu = P_0$ for all ν , the first claim shows that under conditions 1 and 2 $E^{P_0} f$ is lsc.

Theorem 9 is close to the epi-convergence results in [26, 1, 35], but it has some advantages. The results of [35] concerned real-valued functions which do not allow modeling constraints as infinite penalties. In [26, 1], the tightness-like conditions are much stronger than condition 3 above. In particular, condition 3 only asks for a measurable set K_ϵ instead of a compact one, and instead of all sequences $x_\nu \rightarrow x$, it only involves *one* sequence for each $x \in \text{dom} E^{P_0} f$. This is important since it is the tightness-like conditions that are usually hard to check for discretizations in practice. Also, requiring condition 3 to hold at all $x \in X$ would imply $\text{dom} E^{P_0} f = X$. In a sense, we have traded the stronger tightness-like conditions for conditions 1 and 2, which are often much easier to check; see Section 4. The lower-boundedness property in condition 1 holds in many applications arising in practice. Condition 2 holds in particular if f is lsc, which was assumed in [5, Section 8.3] and [26]. According to the remark after Theorem 7, we thus have the following simplified version, which is often sufficient in applications.

Corollary 10 *Let $P_\nu \rightarrow P_0$ and assume that f is lsc. If*

1. *for each $x \in X$, there is an open set $N \ni x$ such that f is bounded from below on $N \times \Xi$,*
2. *for each $x \in \text{dom} E^{P_0} f$, $f(x, \cdot)$ is P_0 -a.s. continuous and bounded,*

then the functions $E^{P_\nu} f$ both pointwise and epi-converge to $E^P f$.

4 Numerical tests

4.1 Markowitz model

We start the testing with a model which can be solved exactly. Of course, discretization is unnecessary in such cases but here we get to compare the approximate solutions with the exact one. We will study the following mean-variance model

$$\begin{aligned} & \underset{x \in \mathbb{R}^n}{\text{minimize}} && E^{P_0}(r \cdot x - \bar{r} \cdot x)^2 && (MP) \\ & \text{subject to} && \bar{r} \cdot x \geq w, \\ & && \sum_{i=1}^n x_i \leq 1, \\ & && x \in C, \end{aligned}$$

where $x = (x_1, \dots, x_n)$ is a portfolio of assets, $r = (r_1, \dots, r_n)$ is the vector of returns (that is, r_i is the ratio of the final and initial price of asset i), $r \cdot x = \sum_{j=1}^n r_j x_j$ is the terminal wealth, w is the required level of expected wealth and C is the set of feasible portfolios. The components of the return vector r

are random variables with joint distribution P_0 and expectation \bar{r} . As is well-known, the expectation in (MP) can be computed explicitly as

$$E^{P_0}(r \cdot x - \bar{r} \cdot x)^2 = E^{P_0}[(r - \bar{r}) \cdot x]^2 = E^{P_0}[x \cdot (r - \bar{r})(r - \bar{r})^T x] = x \cdot Vx,$$

where $V = E^{P_0}[(r - \bar{r})(r - \bar{r})^T]$ is the covariance matrix of r . If V and \bar{r} are known, (MP) can then be solved without discretization with standard solvers yielding the optimal value and optimal solution.

To test the performance of integration quadratures in discretization, we will approximate problem (MP) by the discretizations

$$\begin{aligned} & \underset{x \in \mathbb{R}^n}{\text{minimize}} && \sum_{i=1}^{\nu} p_i^{\nu} (r_i^{\nu} \cdot x - \bar{r} \cdot x)^2 && (MP^{\nu}) \\ & \text{subject to} && \bar{r} \cdot x \geq w, \\ & && \sum_{i=1}^n x_i \leq 1, \\ & && x \in C. \end{aligned}$$

Under mild conditions, convergence of optimal values and solutions can be guaranteed. Recall that the *support*, $\text{supp } P$, of a measure P is the intersection of all closed sets of full measure. For a Borel probability measure $\text{supp } P$ is well defined and unique with $P(\text{supp } P) = 1$.

Proposition 11 *Assume that $\text{supp } P_0$ is bounded, C is closed, and that the measures*

$$P_{\nu} = \sum_{i=1}^{\nu} p_i^{\nu} \delta_{r_i^{\nu}}$$

converge weakly to P_0 and satisfy $\text{supp } P_{\nu} \subset \text{supp } P_0$. If the feasible set is bounded, then the optimal values of (MP^{ν}) converge to that of (MP) and the cluster points of the solutions of (MP^{ν}) are solutions of (MP) .

Proof. This fits the format of (SP) with $\xi = r$ and

$$f(x, r) = (r \cdot x - \bar{r} \cdot x)^2 + \delta_{C'}(x),$$

where $C' = \{x \in C \mid \bar{r} \cdot x \geq w, \sum_{i=1}^n x_i \leq 1\}$. So by Theorem 7, it suffices to verify the conditions of Corollary 10. Lower semicontinuity and condition 1 are clear. Since $\text{supp } P_{\nu} \subset \text{supp } P_0$, we can assume that $\Xi = \text{supp } P_0$, and then condition 2 holds by boundedness of $\text{supp } P_0$. \square

In our test, the number of assets $n = 10$ and

$$r = \bar{r} + 12L(u - \frac{1}{2}e),$$

where u is uniformly distributed in the 10-dimensional unit cube, L is a 10×10 matrix and e is a vector of ones. Then $\text{supp } P_0$ is bounded, r has mean \bar{r} and variance $V = LL^T$. We can then solve (MP) exactly by standard QP-solvers and the discretizations (MP^{ν}) are easily generated by the integration quadratures described in Section 2.2. Note that the objective of (MP^{ν}) can be written as $x \cdot V^{\nu} x$, where $V^{\nu} = E^{P_{\nu}}[(r - \bar{r})(r - \bar{r})^T] = \sum_{i=1}^{\nu} p_i^{\nu} (r_i^{\nu} - \bar{r})(r_i^{\nu} - \bar{r})^T$, so (MP^{ν}) can also be solved with a QP-solver. We chose $C = \mathbb{R}_+^n$, which means that ‘‘short selling’’ is prohibited.

With our choices of \bar{r} and V , the optimal value in the original problem (MP) turned out to be 1.9221. Figure 2(a) shows the development of the optimal value of (MP^{ν}) as a function of the number of quadrature points $\nu = 100, 200, \dots, 10000$ for five quadratures. In our implementation, the number of points in the lattice rule is restricted to powers of 2. Each quadrature produces discretizations whose optimal value seems to converge toward the exact value 1.9221. The objective values corresponding to Halton sequence seem to behave most stably whereas Niederreiter exhibits largest variations.

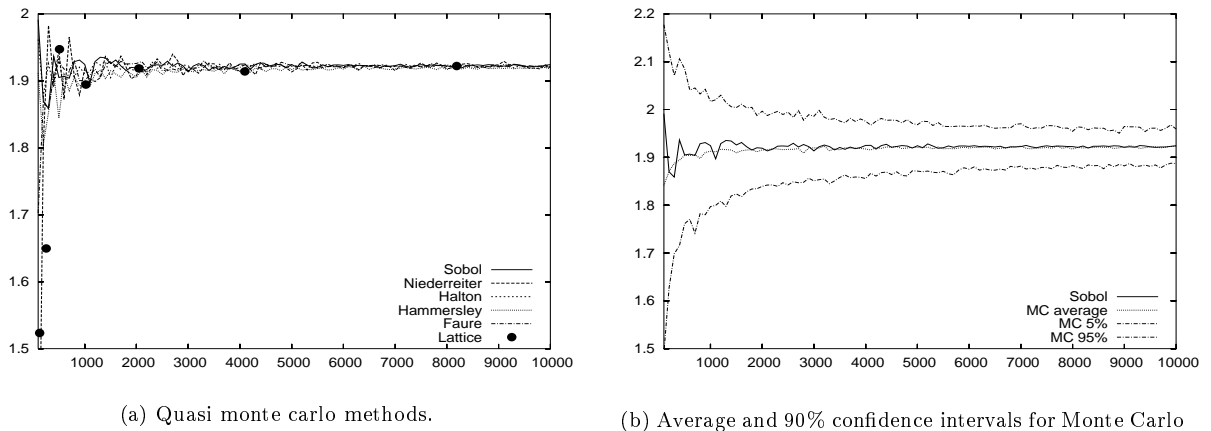


Figure 2: Optimal values of (MP^ν) as a function of ν .

For comparison, we discretized the problem also with Monte Carlo sampling. Almost sure epi-convergence of such discretizations have been established under quite general conditions in [2]. These conditions are strictly weaker than those in Theorem 9, which guarantees *sure* (not just almost sure) epi-convergence. For each $\nu = 100, 200, \dots, 10000$, we generated 250 discretizations and computed the average and the 90% confidence interval of the corresponding optimal values. The results are displayed in Figure 2(b). The optimal values obtained with Sobol sequence are repeated for reference. The average of the Monte Carlo values seems to converge towards the correct value but the convergence of the confidence interval seems slow.

Figure 3(a) displays the logarithmic error in the optimal value for Sobol discretizations as a function of $\ln \nu$. Figure 3(b) does the same for Monte Carlo. Interestingly, in both cases, there seems to be an upper bound on the log-error which is linear in $\ln \nu$. For Sobol, the slope is roughly -1 whereas for MC it seems to be close to -1/2. This kind of quantitative behavior, which may have to do with the strong convexity of the objective, is not explained by the results of Section 3. An explanation would require a quantitative stability analysis with respect to the measures, but that is beyond the scope of this paper.

4.2 Utility maximization

The objective in the Marowitz model penalizes for exceeding the expected wealth $\bar{r} \cdot x$. When the distribution of r is symmetric, this does not matter, but in practice, the distribution of r is usually nonsymmetric since $r \geq 0$. The following utility maximization problem still makes sense

$$\begin{aligned}
 & \underset{x \in \mathbb{R}^n}{\text{maximize}} && E^{P_0} u(r \cdot x) && (UP) \\
 & \text{subject to} && \sum_{i=1}^n x_i \leq w_0, \\
 & && x \in C.
 \end{aligned}$$

Here x , r and C are as in the previous example and u measures the utility from terminal wealth. The components of the return vector r are nonnegative random variables with joint distribution P_0 .

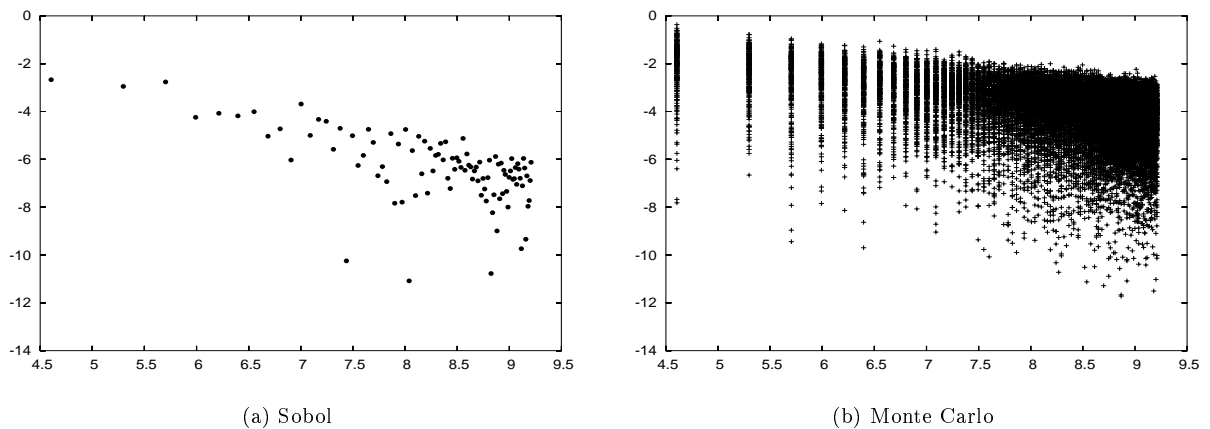


Figure 3: Log of objective errors as a function of $\ln \nu$.

In general, (UP) cannot be solved analytically, so we will consider the discretizations

$$\begin{aligned}
 & \underset{x \in \mathbb{R}^n}{\text{maximize}} && \sum_{i=1}^{\nu} p_i^{\nu} u(r_i^{\nu} \cdot x) && (UP^{\nu}) \\
 & \text{subject to} && \sum_{i=1}^n x_i \leq w_0, \\
 & && x \in C.
 \end{aligned}$$

Proposition 12 Assume $\text{supp } P_0 \subset \mathbb{R}_+^n$, u is continuous and bounded on \mathbb{R}_+ , C is closed and contained in \mathbb{R}_+^n (short selling is not allowed) and that the measures

$$P_{\nu} = \sum_{i=1}^{\nu} p_i^{\nu} \delta_{r_i^{\nu}}$$

converge weakly to P_0 and satisfy $\text{supp } P_{\nu} \subset \mathbb{R}_+^n$. Then the optimal values of (UP^{ν}) converge to that of (UP) and the cluster points of the solutions of (UP^{ν}) are solutions of (UP) .

Proof. This fits the format of (SP) with $\Xi = \mathbb{R}_+^n$, $\xi = r$, and

$$f(x, r) = -u(r \cdot x) + \delta_{C'}(x),$$

where $C' = \{x \in C \mid \sum_{i=1}^n x_i \leq w_0\}$. So by Theorem 7, it suffices to verify the conditions of Corollary 10. Since u is continuous and C is closed, f is lsc. Condition 1 follows from the boundedness of u on \mathbb{R}_+ , and the fact that $r \cdot x \in \mathbb{R}_+$ for all $r \geq 0$ and $x \in C \subset \mathbb{R}_+^n$. Condition 2 follows from the boundedness and continuity of u on \mathbb{R}_+ . \square

Note that many familiar utility functions, like the exponential utility, are bounded on \mathbb{R}_+ . More general utility functions are easily modified to be bounded on \mathbb{R}_+ in a way that does not affect computations in practice.

In our test, the number of assets $n = 10$, r is log-normally distributed, $u(w) = -\exp(-w)$ and $C = \mathbb{R}_+^n$. Figure 4(a) shows the development of the optimal value of (UP^{ν}) as a function of the number of quadrature points $\nu = 100, 110, \dots, 2500$ for six quadratures. Again, the quadratures seem to converge to a common value, but this time, Halton and Hammersley result in largest oscillations of the optimal

values whereas Niederreiter seems to work best. Note, however, that in this problem, we do not know the exact solution of (UP) . Figure 4(b) depicts the development of the average and the 90% confidence interval for the optimal values obtained with 250 Monte Carlo samples for each value of ν . The values obtained with Sobol are shown for reference. The average seems to converge to the same value as the optimal values obtained by the quadratures but the confidence interval narrows down very slowly.

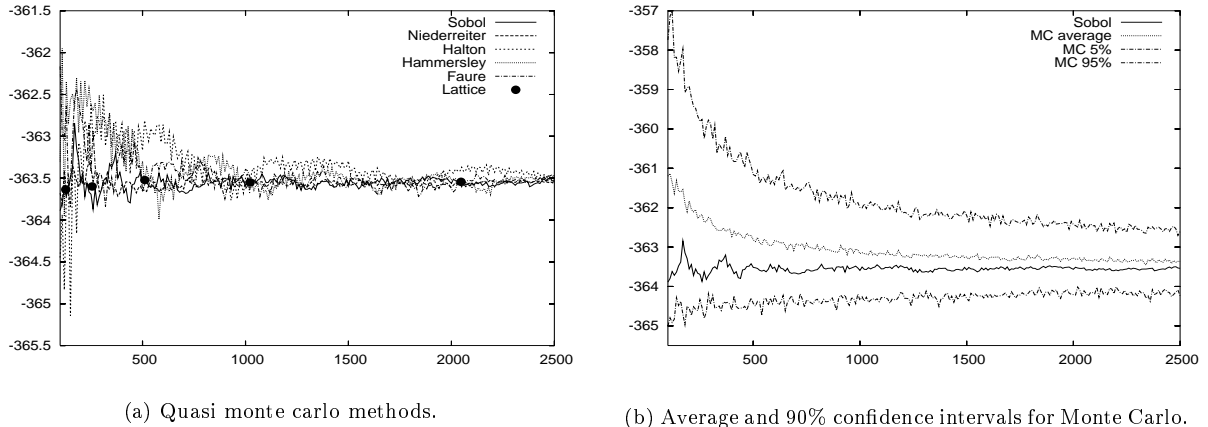


Figure 4: Optimal values of (UP^ν) as a function of ν .

4.3 Super-replication of contingent claims

Consider the problem

$$\begin{aligned}
 & \underset{V, \theta}{\text{minimize}} && V && (PP) \\
 & \text{subject to} && S_0 \cdot \theta \leq V, \\
 & && S \cdot \theta \geq F, \quad P_0\text{-a.s.} \\
 & && \theta \in C,
 \end{aligned}$$

where V is the wealth invested in a portfolio $\theta = (\theta_1, \dots, \theta_J)$ of assets that have prices $S_0 = (S_0^1, \dots, S_0^J)$ at the beginning and $S = (S^1, \dots, S^J)$ at the end of a holding period and F is a cash-flow at the end of the holding period. S and F are random variables with joint distribution P_0 . (PP) can be interpreted as a pricing problem of a seller of the cash-flow F ; see for example King [21] and Korf [22, Section 7]. The seller tries to find the least amount of initial wealth that can be used to buy a portfolio that is almost surely worth at least F at the end of the holding period. The optimum value of (PP) is called the seller's price of F and the corresponding portfolio θ is called a seller's hedge.

(PP) is a semi-infinite linear programming problem and, in general, impossible to solve analytically. Replacing P_0 by a discrete measure $P_\nu = \sum_{i=1}^\nu p_i^\nu \delta_{(S_i^\nu, F_i^\nu)}$ with $p_i^\nu > 0$, for all $i = 1, \dots, \nu$ gives the

problem

$$\begin{aligned}
& \underset{V, \theta}{\text{minimize}} && V && (PP^\nu) \\
& \text{subject to} && S_0 \cdot \theta \leq V, \\
& && S_i^\nu \cdot \theta \geq F_i^\nu, \quad i = 1, \dots, \nu, \\
& && \theta \in C,
\end{aligned}$$

which is an LP problem for which many solvers are available.

Proposition 13 *Assume that the points $\{(S_i^\nu, F_i^\nu)\}_{i=1}^\nu$ are all contained in $\text{supp } P_0$ and that for some $\{p_i^\nu\}_{i=1}^\nu$, $\nu = 0, 1, 2, \dots$, with $p_i^\nu > 0$, for all $i = 1, \dots, \nu$, the measures*

$$P_\nu = \sum_{i=1}^{\nu} p_i^\nu \delta_{(S_i^\nu, F_i^\nu)}$$

converge weakly to P_0 . If the feasible set is bounded, then the optimal values to (PP^ν) converge to the seller's price of F and the cluster points of the solutions of (PP^ν) are seller's hedges for F .

Proof. This fits the format of (SP) with $x = (V, \theta)$, $\xi = (S, F)$ and

$$f(V, \theta, S, F) = V + \delta_C(\theta) + \delta_{C_0}(V, \theta) + \delta_{C_1}(\theta, S, F),$$

where

$$C_0 = \{(V, \theta) \mid S_0 \cdot \theta \leq V\}$$

and

$$C_1 = \{(\theta, S, F) \mid S \cdot \theta \geq F\}.$$

Since C_0 and C_1 are closed, f is lsc. It is also clear that condition 1 of Corollary 10 holds. To verify condition 2, note first that for each $(V, \theta) \in \text{dom } E^{P_0} f$, $f(V, \theta, \cdot)$ is the constant function V on the set

$$C_1(\theta) = \{(S, F) \mid S \cdot \theta \geq F\}$$

which is of full measure. Since $C_1(\theta)$ is closed, we must have $\text{supp } P_0 \subset C_1(\theta)$ for every $(V, \theta) \in \text{dom } E^{P_0} f$. Thus, condition 2 holds if we let $\Xi = \text{supp } P_0$, which is legitimate since $\text{supp } P_\nu \subset \text{supp } P_0$. \square

In our test, the set of assets consists of cash, SP500 index and 28 European call and put options on the index with maturity of 17 calendar days. The cash-flow F is taken to be that of a call option with the same maturity but different strike than any other call included in S . In this case, the value of S is fully determined by the value of the index at the maturity which is assumed to be log-normally distributed.

Figure 5(a) displays the objective values obtained with Sobol along with the averages and 90% confidence intervals obtained with Monte Carlo from 250 samples for each value of $\nu = 100, 110, \dots, 2500$.

Our random variable being one-dimensional in the current problem suggests using Gaussian quadratures. The use of Gauss-Hermite quadrature for discretizing the normal distribution has been already described in Subsection 2.1. We will also utilize Gauss-Legendre quadrature that gives convergent discretizations of the one-dimensional uniform distribution on $[0, 1]$. From this we obtain discretizations of the normal distribution by mapping each point through the inverse of the normal distribution function. The results are shown in Figure 5(b). With $\nu = 60$, the optimal values obtained with Gauss-Legendre quadrature have converged to the same value as the optimal values obtained with Sobol after 2500 points. Gauss-Hermite is almost as good but it results in slight oscillations.

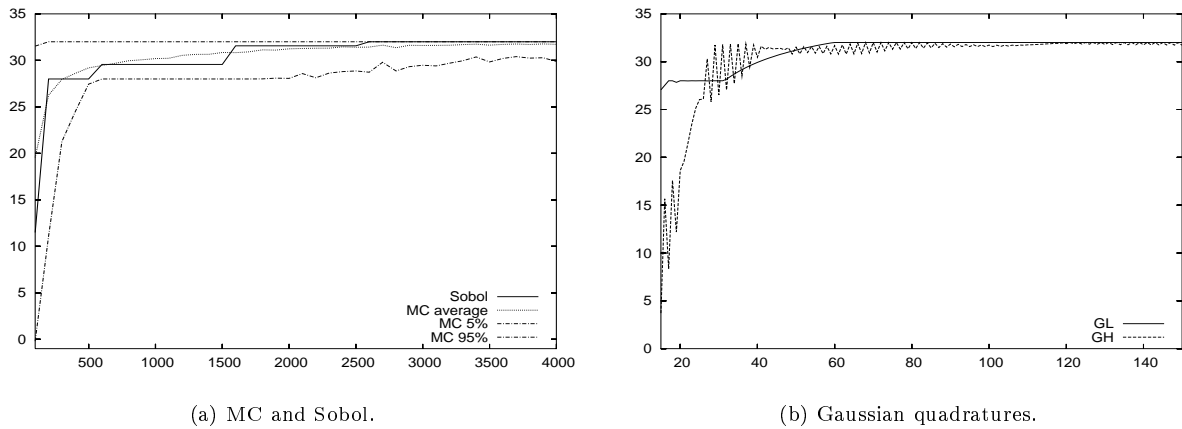


Figure 5: Optimal values of (PP^ν) as a function of ν .

References

- [1] Z. Artstein and R. J.-B. Wets. Stability results for stochastic programs and sensors, allowing for discontinuous objective functions. *SIAM J. Optim.*, 4(3):537–550, 1994.
- [2] Z. Artstein and R. J.-B. Wets. Consistency of minimizers and the SLLN for stochastic programs. *J. Convex Anal.*, 2(1-2):1–17, 1995.
- [3] H. Attouch. *Variational convergence for functions and operators*. Pitman (Advanced Publishing Program), Boston, MA, 1984.
- [4] G. Beer. A geometric algorithm for approximating semicontinuous function. *J. Approx. Theory*, 49(1):31–40, 1987.
- [5] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control*, volume 139 of *Mathematics in Science and Engineering*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York, 1978. The discrete time case.
- [6] P. Billingsley. *Convergence of probability measures*. John Wiley & Sons Inc., New York, 1968.
- [7] P. Billingsley. *Convergence of probability measures*. John Wiley & Sons Inc., New York, second edition, 1999. A Wiley-Interscience Publication.
- [8] J. R. Birge and R. J.-B. Wets. Designing approximation schemes for stochastic optimization problems, in particular for stochastic programs with recourse. *Math. Programming Stud.*, (27):54–102, 1986. Stochastic programming 84. I.
- [9] P. Bratley, B. L. Fox, and H. Niederreiter. Implementation and tests of low-discrepancy sequences. *ACM Trans. Model. Comput. Simul.*, 2(3):195–213, 1992.
- [10] M. A. H. Dempster and R. T. Thompson. EVPI-based importance sampling solution procedures for multistage stochastic linear programmes on parallel MIMD architectures. *Ann. Oper. Res.*, 90:161–184, 1999.
- [11] J. Dupačová and R. J.-B. Wets. Asymptotic behavior of statistical estimators and of optimal solutions of stochastic optimization problems. *Ann. Statist.*, 16(4):1517–1549, 1988.
- [12] H. Faure. Discrépance de suites associées à un système de numération (en dimension s). *Acta Arith.*, 41(4):337–351, 1982.

- [13] B. L. Fox. Algorithm 647: Implementation and relative efficiency of quasirandom sequence generators. *ACM Transactions on Mathematical Software*, 12(4):362–376, 1986.
- [14] K. Frauendorfer. *Stochastic two-stage programming*, volume 392 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, 1992. Habilitationsschrift, University of Zürich, Zürich, 1992.
- [15] J. H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.*, 2:84–90, 1960.
- [16] J. M. Hammersley. Monte Carlo methods for solving multivariable problems. *Ann. New York Acad. Sci.*, 86:844–874 (1960), 1960.
- [17] K. Høyland, M. Kaut, and S. W. Wallace. Heuristic for moment-matching scenario generation. *Computational Optimization and Applications*, 24(2–3):169–185, 2003.
- [18] K. Høyland and S. W. Wallace. Generating scenario trees for multistage decision problems. *Management Science*, 47(2):295–307, 2001.
- [19] G. Infanger. Monte Carlo (importance) sampling within a Benders decomposition algorithm for stochastic linear programs. *Ann. Oper. Res.*, 39(1-4):69–95 (1993), 1992.
- [20] P. Kall, A. Ruszczyński, and K. Frauendorfer. Approximation techniques in stochastic programming. In *Numerical techniques for stochastic optimization*, volume 10 of *Springer Ser. Comput. Math.*, pages 33–64. Springer, Berlin, 1988.
- [21] A. J. King. Duality and martingales: a stochastic programming perspective on contingent claims. *Math. Program.*, 91(3, Ser. B):543–562, 2002. ISMP 2000, Part 1 (Atlanta, GA).
- [22] L. Korf. Stochastic programming duality: l^∞ multipliers for unbounded constraints with an application to mathematical finance. *submitted*, 2002.
- [23] N. M. Korobov. Approximate evaluation of repeated integrals. *Dokl. Akad. Nauk SSSR*, 124:1207–1210, 1959.
- [24] P. L’Ecuyer and C. Lemieux. Variance reduction via lattice rules. *Management Science*, 46(2):1214–1235, 2000.
- [25] R. Lucchetti, G. Salinetti, and R. J.-B. Wets. Uniform convergence of probability measures: topological criteria. *J. Multivariate Anal.*, 51(2):252–264, 1994.
- [26] R. Lucchetti and R. J.-B. Wets. Convergence of minima of integral functionals, with applications to optimal control and stochastic optimization. *Statist. Decisions*, 11(1):69–84, 1993.
- [27] H. Niederreiter. *Random number generation and quasi-Monte Carlo methods*, volume 63 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
- [28] G. Ch. Pflug. Scenario tree generation for multiperiod financial optimization by optimal discretization. *Math. Program.*, 89(2, Ser. B):251–271, 2001. Mathematical programming and finance.
- [29] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C*. Cambridge University Press, Cambridge, second edition, 1992. The art of scientific computing.
- [30] S. M. Robinson and R. J.-B. Wets. Stability in two-stage stochastic programming. *SIAM J. Control Optim.*, 25(6):1409–1416, 1987.
- [31] R. T. Rockafellar and R. J.-B. Wets. *Variational analysis*, volume 317 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1998.
- [32] R. Schultz. Some aspects of stability in stochastic programming. *Ann. Oper. Res.*, 100:55–84 (2001), 2000. Research in stochastic programming (Vancouver, BC, 1998).
- [33] I. H. Sloan and S. Joe. *Lattice methods for multiple integration*. Oxford Science Publications. The Clarendon Press Oxford University Press, New York, 1994.

- [34] I. M. Sobol'. The distribution of points in a cube and the approximate evaluation of integrals. *U.S.S.R. Computational Math. And Math. Phys.*, (4):86–112, 1967.
- [35] M. Zervos. On the epiconvergence of stochastic optimization problems. *Math. Oper. Res.*, 24(2):495–508, 1999.