



HAL
open science

Sequential quasi Monte Carlo

Mathieu Gerber, Nicolas Chopin

► **To cite this version:**

Mathieu Gerber, Nicolas Chopin. Sequential quasi Monte Carlo. Journal of the Royal Statistical Society. Series B. Methodological, 2015, 77 (3), pp.509-579. 10.1111/rssb.12104 . hal-01409255

HAL Id: hal-01409255

<https://hal.science/hal-01409255v1>

Submitted on 25 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sequential Quasi-Monte Carlo

Mathieu Gerber*

Nicolas Chopin†

We derive and study SQMC (Sequential Quasi-Monte Carlo), a class of algorithms obtained by introducing QMC point sets in particle filtering. SQMC is related to, and may be seen as an extension of, the array-RQMC algorithm of L'Ecuyer et al. (2006). The complexity of SQMC is $\mathcal{O}(N \log N)$, where N is the number of simulations at each iteration, and its error rate is smaller than the Monte Carlo rate $\mathcal{O}_P(N^{-1/2})$. The only requirement to implement SQMC is the ability to write the simulation of particle \mathbf{x}_t^n given \mathbf{x}_{t-1}^n as a deterministic function of \mathbf{x}_{t-1}^n and a fixed number of uniform variates. We show that SQMC is amenable to the same extensions as standard SMC, such as forward smoothing, backward smoothing, unbiased likelihood evaluation, and so on. In particular, SQMC may replace SMC within a PMCMC (particle Markov chain Monte Carlo) algorithm. We establish several convergence results. We provide numerical evidence that SQMC may significantly outperform SMC in practical scenarios.

Key-words: Array-RQMC; Low discrepancy; Particle filtering; Quasi-Monte Carlo; Randomized Quasi-Monte Carlo; Sequential Monte Carlo

1. Introduction

Sequential Monte Carlo (SMC, also known as particle filtering) is a class of algorithms for computing recursively Monte Carlo approximations of a sequence of distributions $\pi_t(d\mathbf{x}_t)$, $t \in 0:T$, $0:T = \{0, \dots, T\}$. The initial motivation of SMC was the filtering of state-space models (also known as hidden Markov models); that is, given a latent Markov process (\mathbf{x}_t) , observed imperfectly as e.g. $\mathbf{y}_t = f(\mathbf{x}_t) + \epsilon_t$, recover at every time t the distribution of \mathbf{x}_t given the data $\mathbf{y}_{0:t} = (\mathbf{y}_0, \dots, \mathbf{y}_t)$. SMC's popularity stems from the fact it is the only realistic approach for filtering and related problems outside very specific cases (such as the linear Gaussian model). Recent research has further increased interest in SMC, especially in Statistics, in at least two directions. First, several papers (Neal, 2001; Chopin, 2002; Del Moral et al., 2006) have extended SMC to non-sequential problems; that is, to sample from distribution π , one applies SMC to some artificial

*Université de Lausanne, and CREST

†CREST-ENSAE (corresponding author, nicolas.chopin@ensae.fr)

sequence π_t that ends up at $\pi_T = \pi$. In certain cases, such an approach outperforms MCMC (Markov chain Monte Carlo) significantly. Second, the seminal paper of Andrieu et al. (2010) established that SMC may be used as a proposal mechanism within MCMC, leading to so called PMCMC (particle MCMC) algorithms. While not restricted to such problems, PMCMC is the only possible approach for inference in state-space models such that the transition kernel of (\mathbf{x}_t) may be sampled from, but does not admit a tractable density. Excitement about PMCMC is evidenced by the 30 papers or so that have appeared in the last two years on possible applications and extensions.

Informally, the error rate of SMC at iteration t is $C_t N^{-1/2}$, where C_t is some function of t . There has been a lot of work on SMC error rates (e.g. Del Moral, 2004), but it seems fair to say that most of it has focussed on the first factor C_t ; that is, whether to establish that the error rate is bounded uniformly in time, $C_t \leq C$, or to reduce C_t through more efficient algorithmic designs, such as better proposal kernels or resampling schemes.

In this work, we focus on the second factor $N^{-1/2}$, i.e. we want the error rate to converge quicker relative to N than the standard Monte Carlo rate $N^{-1/2}$. To do so, we adapt to the SMC context ideas borrowed from QMC (Quasi-Monte Carlo); that is, the idea of replacing random numbers by low discrepancy point sets.

The following subsections contain very brief introductions to SMC and QMC, with an exclusive focus on the concepts that are essential to follow this work. For a more extensive presentation of SMC, the reader is referred to the books of Doucet et al. (2001), Del Moral (2004) and Cappé et al. (2005), while for QMC and RQMC, see Chapter 5 of Glasserman (2004), Chapters 5 and 6 of Lemieux (2009), and Dick and Pillichshammer (2010).

1.1. Introduction to SMC

As already mentioned, the initial motivation of SMC is the sequential analysis of state-space models; that, is models for a Markov chain (\mathbf{x}_t) in $\mathcal{X} \subseteq \mathbb{R}^d$,

$$\mathbf{x}_0 \sim f_0^X(\mathbf{x}_0), \quad \mathbf{x}_t | \mathbf{x}_{t-1} \sim f^X(\mathbf{x}_t | \mathbf{x}_{t-1}),$$

which is observed only indirectly through some \mathbf{y}_t , with density $\mathbf{y}_t | \mathbf{x}_t \sim f^Y(\mathbf{y}_t | \mathbf{x}_t)$.

This kind of model arises in many areas of science: in tracking for instance, \mathbf{x}_t may be the position of a ship (in two dimensions) or a plane (in three dimensions), and \mathbf{y}_t may be a noisy angular observation (radar). In Ecology, \mathbf{x}_t would be the size of a population of bats in a cave, and \mathbf{y}_t would be \mathbf{x}_t plus noise. And so on.

The most standard inferential task for such models is that of *filtering*; that is, to recover iteratively in time t , $p(\mathbf{x}_t | \mathbf{y}_{0:t})$, the distribution of \mathbf{x}_t , given the data collected up time t , $\mathbf{y}_{0:t} = (\mathbf{y}_0, \dots, \mathbf{y}_t)$. One may also be interested in smoothing, $p(\mathbf{x}_{0:t} | \mathbf{y}_{0:t})$, or likelihood evaluation, $p(\mathbf{y}_{0:t})$, notably when the model depends on a fixed parameter θ which should be learnt from the data.

A simple Monte Carlo approach to filtering is sequential importance sampling: choose an initial distribution $m_0(d\mathbf{x}_0)$, a sequence of Markov kernels $m_t(\mathbf{x}_t | \mathbf{x}_{t-1}, d\mathbf{x}_t)$, $t \geq 1$, then simulate N times iteratively from these m_t 's, $\mathbf{x}_0^n \sim m_0(d\mathbf{x}_0)$, $\mathbf{x}_t^n | \mathbf{x}_{t-1}^n \sim m_t(\mathbf{x}_t^n | \mathbf{x}_{t-1}^n, d\mathbf{x}_t)$,

and reweight ‘particle’ (simulation) \mathbf{x}_t^n as follows: $w_0^n = G_0(\mathbf{x}_0^n)$, $w_t^n = w_{t-1}^n \times G_t(\mathbf{x}_{t-1}^n, \mathbf{x}_t^n)$, where the weight functions G_t are defined as

$$G_0(\mathbf{x}_0) = \frac{f^Y(\mathbf{y}_0|\mathbf{x}_0)f_0^X(\mathbf{x}_0)}{m_0(\mathbf{x}_0)}, \quad G_t(\mathbf{x}_{t-1}, \mathbf{x}_t) = \frac{f^Y(\mathbf{y}_t|\mathbf{x}_t)f^X(\mathbf{x}_t|\mathbf{x}_{t-1})}{m_t(\mathbf{x}_t|\mathbf{x}_{t-1})}, \quad (1)$$

and $m_t(\mathbf{x}_t|\mathbf{x}_{t-1})$ in the denominator denotes the conditional probability density associated to kernel $m_t(\mathbf{x}_{t-1}, d\mathbf{x}_t)$. Then it is easy to check that the weighted average $\sum_{n=1}^N w_t^n \varphi(\mathbf{x}_t^n) / \sum_{n=1}^N w_t^n$ is a consistent estimate of the filtering expectation $\mathbb{E}[\varphi(\mathbf{x}_t)|\mathbf{y}_{0:t}]$, as $N \rightarrow +\infty$. However, it is well known that, even for carefully chosen proposal densities m_t , sequential importance sampling quickly degenerates: as time progresses, more and more particles get a negligible weight.

Surprisingly, there is a simple solution to this degeneracy problem: one may *resample* the particles; that is, draw N times with replacement from the set of particles, with probabilities proportional to the weights w_t^n . In this way, particles with low weight gets quickly discarded, while particles with large weight may get many children at the following iteration. Empirically, the impact of resampling is dramatic: the variance of filtering estimates typically remains stable over time, while without resampling it diverges exponentially fast.

The idea of using resampling may be traced back to Gordon et al. (1993), and has initiated the whole field of particle filtering. See Algorithm 1 for a summary of a basic PF (particle filter). The price to pay for introducing resampling is that it creates non-trivial dependencies between the particles, which complicates the formal study of such algorithms. In particular, establishing convergence (as $N \rightarrow +\infty$) is non-trivial, although the error rate of SMC is known to be $\mathcal{O}_P(N^{-1/2})$; see e.g. the central limit theorems of Del Moral and Guionnet (1999), Chopin (2004) and Künsch (2005). We shall see that it is also the resampling step that makes the introduction of Quasi-Monte Carlo into SMC non-trivial.

Algorithm 1 Basic particle filter

At time $t = 0$,

- (a) Generate $\mathbf{x}_0^n \sim m_0(d\mathbf{x}_0)$ for all $n \in 1:N$.
- (b) Compute $w_0^n = G_0(\mathbf{x}_0^n)$ and $W_0^n = w_0^n / \sum_{m=1}^N w_0^m$ for all $n \in 1:N$.

From time $t = 1$ to time T ,

- (a) Generate $a_{t-1}^n \sim \mathcal{M}(W_{t-1}^{1:N})$ for all $n \in 1:N$, the multinomial distribution that produces outcome m with probability W_{t-1}^m . See Algorithm 2.
 - (b) Generate $\mathbf{x}_t^n \sim m_t(\mathbf{x}_{t-1}^{a_{t-1}^n}, d\mathbf{x}_t)$ for all $n \in 1:N$.
 - (c) Compute $w_t^n = G_t(\mathbf{x}_{t-1}^{a_{t-1}^n}, \mathbf{x}_t^n)$, and $W_t^n = w_t^n / \sum_{m=1}^N w_t^m$ for all $n \in 1:N$.
-

The complexity of SMC is $\mathcal{O}(N)$. In particular, to implement the resampling step in

$\mathcal{O}(N)$ time (Step (a) at times $t \geq 1$ in Algorithm 1), one proceeds as follows: (a) generate $u^{1:N} = \text{sort}(v^{1:N})$, where the v^n are independent uniform variates (see p.214 of Devroye, 1986, for a well-known algorithm to generate $u^{1:N}$ directly in $\mathcal{O}(N)$ time, without any sorting); and (b) use the inverse transform method for discrete distributions, recalled in Algorithm 2. We will re-use Algorithm 2 in SQMC.

Algorithm 2 Resampling Algorithm (inverse transform method)

Input: $u^{1:N}$ (such that $0 \leq u^1 \leq \dots \leq u^N \leq 1$), $W^{1:N}$ (normalised weights)

Output: $a^{1:N}$ (labels in $1 : N$)

```

 $s \leftarrow W^1, m \leftarrow 1$ 
for  $n = 1 \rightarrow N$  do
  while  $s < u^n$  do
     $m \leftarrow m + 1$ 
     $s \leftarrow s + W^m$ 
  end while
   $a^n \leftarrow m$ 
end for

```

1.2. Introduction to QMC

QMC (Quasi-Monte Carlo) is generally presented as a way to perform integration with respect to the (semi-closed) hypercube of dimension d :

$$\frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{u}^n) \approx \int_{[0,1]^d} \varphi(\mathbf{u}) \, d\mathbf{u}$$

where the N vectors $\mathbf{u}^n \in [0, 1]^d$ must be chosen so as to have “low discrepancy”, that is, informally, to be spread evenly over $[0, 1]^d$. (We respect the standard convention in the QMC literature to work with space $[0, 1)^d$, rather than $[0, 1]^d$, as it turns out to be technically more convenient.)

Formally, the general notion of discrepancy is defined as

$$D(\mathbf{u}^{1:N}; \mathcal{A}) = \sup_{A \in \mathcal{A}} \left| \frac{1}{N} \sum_{n=1}^N \mathbb{1}(\mathbf{u}^n \in A) - \lambda_d(A) \right|$$

where $\lambda_d(A)$ is the volume (Lebesgue measure on \mathbb{R}^d) of A , and \mathcal{A} is a set of measurable sets. Two discrepancies are particularly useful in this work: the extreme discrepancy,

$$D(\mathbf{u}^{1:N}) = \sup_{[\mathbf{a}, \mathbf{b}]} \left| \frac{1}{N} \sum_{n=1}^N \mathbb{1}(\mathbf{u}^n \in [\mathbf{a}, \mathbf{b}]) - \prod_{i=1}^d (b_i - a_i) \right|$$

which is the discrepancy relative to the set \mathcal{A} of d -dimensional intervals $[\mathbf{a}, \mathbf{b}] :=$

$\prod_{i=1}^d [a_i, b_i]$, $0 \leq a_i < b_i < 1$; and the star discrepancy:

$$D^*(\mathbf{u}^{1:N}) = \sup_{[\mathbf{0}, \mathbf{b}]} \left| \frac{1}{N} \sum_{n=1}^N \mathbf{1}(\mathbf{u}^n \in [\mathbf{0}, \mathbf{b}]) - \prod_{i=1}^d b_i \right|,$$

where again $[\mathbf{0}, \mathbf{b}] = \prod_{i=1}^d [0, b_i]$, $0 < b_i < 1$. When $d = 1$, the star discrepancy is the Kolmogorov–Smirnov statistic for an uniformity test of the points \mathbf{u}^n .

These two discrepancies are related as follows (Niederreiter, 1992, Proposition 2.4):

$$D^*(\mathbf{u}^{1:N}) \leq D(\mathbf{u}^{1:N}) \leq 2^d D^*(\mathbf{u}^{1:N}).$$

The importance of the concept of discrepancy, and in particular of the star discrepancy, is highlighted by the Koksma–Hlawka inequality (see e.g. Kuipers and Niederreiter, 1974, Theorem 5.1):

$$\left| \frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{u}^n) - \int_{[0,1]^d} \varphi(\mathbf{u}) \, d\mathbf{u} \right| \leq V(\varphi) D^*(\mathbf{u}^{1:N})$$

which conveniently separates the effect of the smoothness of φ (as measured by $V(\varphi)$, the total variation in the sense of Hardy and Krause, see Chapter 2 of Niederreiter, 1992 for a definition), and the effect of the discrepancy of the points $\mathbf{u}^{1:N}$. The quantity $V(\varphi)$ is generally too difficult to compute in practice, and the Koksma–Hlawka inequality is used mainly to determine the asymptotic error rate (as $N \rightarrow +\infty$), through the quantity $D^*(\mathbf{u}^{1:N})$.

There are several methods to construct $\mathbf{u}^{1:N}$ so that $D^*(\mathbf{u}^{1:N}) = \mathcal{O}(N^{-1+\epsilon})$ for any $\epsilon > 0$; which is of course better than the Monte Carlo rate $\mathcal{O}_P(N^{-1/2})$. The best known rates are $\mathcal{O}(N^{-1}(\log N)^{d-1})$ for QMC point sets $\mathbf{u}^{N,1:N}$ that are allowed to depend on N (i.e. $\mathbf{u}^{N,1:N}$ are not necessarily the N first elements of $\mathbf{u}^{N+1,1:N+1}$) and $\mathcal{O}(N^{-1}(\log N)^d)$ for QMC sequences (that is $\mathbf{u}^{1:N}$ are the N first elements of a sequence (\mathbf{u}^n) which may be generated iteratively). For simplicity, we will not distinguish further QMC point sets and QMC sequences, and will use the same notation $\mathbf{u}^{1:N}$ in both cases (although our results will apply to both types of construction).

These asymptotic rates seem to indicate that the comparative performance of QMC over Monte Carlo should deteriorate with d : for $d = 10$, $N^{-1}(\log N)^d \leq N^{-1/2}$ only for $N \geq 1.3 \times 10^{39}$. But since these rates correspond to an upper bound for the error size, it is hard to determine beforehand if and when QMC “breaks” with the dimension. For instance, Glasserman (2004, p.327) exhibits a numerical example where QMC remains competitive relative to Monte Carlo for $d \geq 150$ and $N \leq 10^5$.

Describing the different strategies to construct low-discrepancy point sets is beyond the scope of this paper; see again the aforementioned books on QMC. Figure 1 illustrates the greater regularity of a QMC point set over a set of random points.

1.3. Introduction to RQMC

RQMC (randomized QMC) amounts to randomize the points $\mathbf{u}^{1:N}$ in such a way that (a) they still have low discrepancy (with probability one); and (b) each $\mathbf{u}^n \sim \mathcal{U}([0,1]^d)$




Figure 1: QMC versus Monte Carlo: $N = 256$ points sampled independently and uniformly in $[0, 1]^2$ (left); QMC sequence (Sobol') in $[0, 1]^2$ of the same length (right).

marginally. The simplest construction of such RQMC point sets is the random shift method proposed by Cranley and Patterson (1976) in which we take $\mathbf{u}^n = \mathbf{v}^n + \mathbf{w} \pmod{1}$, where $\mathbf{w} \sim \mathcal{U}([0, 1]^d)$ and $\mathbf{v}^{1:N}$ is a low-discrepancy point set.

RQMC has two advantages over QMC. First, one then obtains an unbiased estimator of the integral of interest:

$$\mathbb{E} \left\{ \frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{u}^n) \right\} = \int_{[0,1]^d} \varphi(\mathbf{u}) \, d\mathbf{u},$$

which makes it possible to evaluate the approximation error through independent replications. We will see that, in our context, this unbiasedness property will also be very convenient for another reason: namely to provide an unbiased estimate of the likelihood of the considered state-space model.

Second, Owen (1997a,b, 1998) established that randomization may lead to better rates, in the following sense: under appropriate conditions, and for a certain type of randomization scheme known as nested scrambling, the mean square error of a RQMC estimator is $\mathcal{O}(N^{-3+\epsilon})$. The intuition behind this rather striking result is that randomization may lead to cancellation of certain error terms.

1.4. A note on array-RQMC

Consider the following problem: we have a Markov chain in \mathcal{X} , whose evolution may be formulated as

$$\mathbf{x}_t = \Gamma_t(\mathbf{x}_{t-1}, \mathbf{u}_t), \quad \mathbf{u}_t \sim \mathcal{U}([0, 1]^d), \quad t \geq 1, \quad \mathbf{x}_0 \text{ is fixed,}$$

and we wish to compute the expectation of $\sum_{t=1}^T \varphi_t(\mathbf{x}_t)$, for certain functions φ_t .

From the two previous sections, we see that a simple approach to this problem would be to generate a QMC (or RQMC) point set $\mathbf{u}^{1:N}$ in $[0, 1)^{dT}$, $\mathbf{u}^n = (\mathbf{u}_1^n, \dots, \mathbf{u}_T^n)$, to transform \mathbf{u}_t^n into $\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^n, \mathbf{u}_t^n)$, and finally to return the corresponding empirical average, $N^{-1} \sum_{t=1}^T \varphi_t(\mathbf{x}_t^n)$. The problem with this direct approach is that the dimension dT of $\mathbf{u}^{1:N}$ may be very large, and, as we have seen, equidistribution properties of QMC point sets (as measured by the star discrepancy) deteriorate with the dimension.

An elegant alternative to this approach is the array-RQMC algorithm of L'Ecuyer et al. (2006), see also Lécot and Ogawa (2002), Lécot and Tuffin (2004), and L'Ecuyer et al. (2009). The main idea of this method is to replace the QMC point set in $[0, 1)^{dT}$ by T QMC points sets $\mathbf{u}_t^{1:N}$ in $[0, 1)^d$. Then, \mathbf{x}_t^n is obtained as $\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^{a_t^n}, \mathbf{u}_t^n)$, where the ancestor $\mathbf{x}_{t-1}^{a_t^n}$ of \mathbf{x}_t^n is chosen so as to be the n -th “smallest” point among the \mathbf{x}_{t-1}^n 's. Note that array-RQMC therefore requires to specify a total order for the state space \mathcal{X} ; for instance one may define a certain $\omega : \mathcal{X} \rightarrow \mathbb{R}$ so that $\omega(\mathbf{x}) \leq \omega(\mathbf{x}')$ means that \mathbf{x} is “smaller” than \mathbf{x}' .

Array-RQMC is shown to have excellent empirical performance in the aforementioned papers. On the other hand, it is currently lacking in terms of supporting theory (see however ?, for $d = 1$); in particular, it is not clear how to choose the order ω , beside the obvious case where $\mathcal{X} \subset \mathbb{R}$. The SQMC algorithm we develop in this paper may be seen as an extension of array-RQMC to particle filtering. In particular, it re-uses the essential idea to generate one QMC point set at each step of the simulation process. As an added benefit, the convergence results we obtain for SQMC also apply to array-RQMC, provided the state space is ordered through the Hilbert curve, as explained later.

1.5. Background, plan and notations

QMC is already very popular in Finance for e.g. derivative pricing (Glasserman, 2004), and one may wonder why it has not received more attention in Statistics so far. The main reason seems to be the perceived difficulty to adapt QMC to non-independent simulation such as MCMC (Markov chain Monte Carlo); see however Chen et al. (2011) and references therein, in particular Tribble (2007), for exciting numerical and theoretical results in this direction which ought to change this perception.

Regarding SMC, we are aware of two previous attempts to develop QMC versions of these algorithms: Lemieux et al. (2001) and Fearnhead (2005); see also Guo and Wang (2006) who essentially proposed the same algorithm as Fearnhead (2005). The first paper casts SMC as a Monte Carlo algorithm in $d(T + 1)$ dimensions, where $d = \dim(\mathcal{X})$, and therefore requires to generate a low-discrepancy point set in $[0, 1)^{d(T+1)}$. But, as we have already explained, such an approach may not work well when $d(T + 1)$ is too large.

Our approach is closer to, and partly inspired by, the RPF (regularized particle filter) of Fearnhead (2005), who, in the same spirit as array-RQMC, casts SMC as a sequence of $T + 1$ successive importance sampling steps of dimension d . (The paper focus on the $d = 1$ case.) The main limitation of the RPF is that it has complexity $\mathcal{O}(N^2)$. This

is because the importance sampling steps are defined with respect to a target which is a mixture of N components, hence the evaluation of a single importance weight costs $\mathcal{O}(N)$.

The SQMC algorithm we develop in this paper has complexity $\mathcal{O}(N \log(N))$ per time step. It is also based on a sequence of $T + 1$ importance sampling steps, but of dimension $d + 1$; the first component is used to determine which ancestor \mathbf{x}_{t-1}^m should be assigned to particle \mathbf{x}_t^n . For $d > 1$, this requires us to “project” the set of ancestors $\mathbf{x}_{t-1}^{1:N} \in \mathcal{X}^N$ into $[0, 1)^N$, by means of a space-filling curve known as the Hilbert curve. The choice of this particular space-filling curve is not only for computational convenience, but also because of its nice properties regarding conversion of discrepancy, as we will explain in the paper. (One referee pointed out to us that the use of Hilbert curve in the context of array-RQMC has been suggested by Wächter and Keller (2008), but not implemented.)

The paper is organised as follows. Section 2 derives the general SQMC algorithm, first for $d = 1$, then for any d through the use of the Hilbert curve. Section 3 presents several convergence results; proofs of these results are in the Appendix. Section 4 shows how several standard extensions of SMC, such as forward smoothing, backward smoothing, and PMCMC, may be adapted to SQMC. Section 5 compares numerically SQMC with SMC. Section 6 concludes.

Most random variables in this work will be vectors in \mathbb{R}^d , and will be denoted in bold face, \mathbf{u} or \mathbf{x} . In particular, \mathcal{X} will be an open set of \mathbb{R}^d . The Lebesgue measure in dimension d is denoted by λ_d . Let $\mathcal{P}(\mathcal{X})$ be the set of probability measures defined on \mathcal{X} dominated by λ_d (restricted to \mathcal{X}), and $\pi(\varphi)$ be the expectation of function φ relative to $\pi \in \mathcal{P}(\mathcal{X})$. Let $a : b$ be the set of integers $\{a, \dots, b\}$ for $a \leq b$. We also use this notation for collections of random variables, e.g. $\mathbf{x}_t^{1:N} = (\mathbf{x}_t^1, \dots, \mathbf{x}_t^N)$, $\mathbf{x}_{0:t} = (\mathbf{x}_0, \dots, \mathbf{x}_t)$ and so on.

2. SQMC

The objective of this section is to construct the SQMC algorithm. To this aim, we discuss how to rewrite SMC as a deterministic function of independent uniform variates $\mathbf{u}_t^{1:N}$, $t \in 0:T$, which then may be replaced by low-discrepancy point sets.

2.1. SMC formalisation

A closer inspection of our basic particle filter, Algorithm 1, reveals that this algorithm is entirely determined by (a) the sequence of proposal kernels $(m_t)_{t \geq 0}$ (which determine how particles are simulated) and (b) the sequence of weight functions $(G_t)_{t \geq 0}$ (which determine how particles are weighted). Our introduction to particle filtering focussed on the specific expression (1) for G_t , but useful SMC algorithms may be obtained by considering other weight functions; see e.g. the auxiliary particle filter of Pitt and Shephard (1999), as explained in Johansen and Doucet (2008), or the SMC algorithms for non-sequential problems mentioned in the introduction.

The exact expression and meaning of m_t and G_t will not play a particular role in the rest of the paper, so it is best to think of SMC from now on as a generic algorithm,

again based on a certain sequence (m_t) , $m_0(d\mathbf{x}_0)$ being an initial distribution, and $m_t(\mathbf{x}_{t-1}, d\mathbf{x}_t)$ being a Markov kernel for $t \geq 1$, and a certain sequence of functions, $G_0 : \mathcal{X} \rightarrow \mathbb{R}^+$, $G_t : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$, which produces the following consistent (as $N \rightarrow +\infty$) estimators:

$$\frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{x}_t^n) \rightarrow \overline{\mathbb{Q}}_t(\varphi), \quad \sum_{n=1}^N W_t^n \varphi(\mathbf{x}_t^n) \rightarrow \mathbb{Q}_t(\varphi),$$

where $\varphi : \mathcal{X} \rightarrow \mathbb{R}$, and $\overline{\mathbb{Q}}_t$ and \mathbb{Q}_t are defined as follows:

$$Z_t = \mathbb{E} \left[G_0(\mathbf{x}_0) \prod_{s=1}^t G_s(\mathbf{x}_{s-1}, \mathbf{x}_s) \right], \quad (2)$$

$$\overline{\mathbb{Q}}_t(\varphi) = \frac{1}{Z_{t-1}} \mathbb{E} \left[\varphi(\mathbf{x}_t) G_0(\mathbf{x}_0) \prod_{s=1}^{t-1} G_s(\mathbf{x}_{s-1}, \mathbf{x}_s) \right], \quad (3)$$

$$\mathbb{Q}_t(\varphi) = \frac{1}{Z_t} \mathbb{E} \left[\varphi(\mathbf{x}_t) G_0(\mathbf{x}_0) \prod_{s=1}^t G_s(\mathbf{x}_{s-1}, \mathbf{x}_s) \right], \quad (4)$$

with expectations taken with respect to the law of the non-homogeneous Markov chain (\mathbf{x}_t) , e.g.

$$Z_t = \int_{\mathcal{X}^{t+1}} \left\{ G_0(\mathbf{x}_0) \prod_{s=1}^t G_s(\mathbf{x}_{s-1}, \mathbf{x}_s) \right\} m_0(d\mathbf{x}_0) \prod_{s=1}^t m_s(\mathbf{x}_{s-1}, d\mathbf{x}_s),$$

and with the conventions that $Z_{-1} = 1$ and empty products equal one; e.g. $\overline{\mathbb{Q}}_0(\varphi) = m_0(\varphi)$.

For instance, for the standard filtering problem covered in our introduction, where G_t is set to (1), $\mathbb{Q}_t(\varphi)$ is the filtering expectation of φ , i.e. $\mathbb{E}[\varphi(\mathbf{x}_t) | \mathbf{y}_{0:t}]$, and $\overline{\mathbb{Q}}_t(\varphi)$ is the predictive distribution of φ , i.e. $\mathbb{E}[\varphi(\mathbf{x}_t) | \mathbf{y}_{0:t-1}]$.

2.2. Towards SQMC: SMC as a sequence of importance sampling steps

QMC requires to write any simulation as an explicit function of uniform variates. We therefore make the following assumption for our generic SMC sampler: to generate $\mathbf{x}_0^n \sim m_0(d\mathbf{x}_0)$, one computes $\mathbf{x}_0^n = \Gamma_0(\mathbf{u}_0^n)$, and to generate $\mathbf{x}_t^n | \mathbf{x}_{t-1}^n \sim m_t(\mathbf{x}_{t-1}, d\mathbf{x}_t)$, one computes $\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^n, \mathbf{u}_t^n)$, where $\mathbf{u}_t^n \sim \mathcal{U}([0, 1]^d)$, and the functions Γ_t are easy to evaluate.

Iteration 0 of Algorithm 1 amounts to an importance sampling step, from $m_0(d\mathbf{x}_0)$ to $\mathbb{Q}_0(d\mathbf{x}_0) = m_0(d\mathbf{x}_0)G_0(\mathbf{x}_0)/Z_0$, which produces the following estimator

$$\sum_{n=1}^N W_0^n \varphi(\mathbf{x}_0^n) = \frac{\sum_{n=1}^N G_0(\mathbf{x}_0^n) \varphi(\mathbf{x}_0^n)}{\sum_{m=1}^N G_0(\mathbf{x}_0^m)}$$

of $\mathbb{Q}_0(\varphi)$. To introduce QMC at this stage, we take $\mathbf{x}_0^n = \Gamma_0(\mathbf{u}_0^n)$ where $\mathbf{u}_0^{1:N}$ is a low-discrepancy point set in $[0, 1]^d$.

The key remark that underpins SQMC is that iteration $t \geq 1$ of Algorithm 1 also amounts to an importance sampling step, but this time from

$$\overline{\mathbb{Q}}_t^N(d(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t)) = \sum_{n=1}^N W_{t-1}^n \delta_{\mathbf{x}_{t-1}^n}(d\tilde{\mathbf{x}}_{t-1}) m_t(\mathbf{x}_{t-1}^n, d\mathbf{x}_t) \quad (5)$$

to

$$\mathbb{Q}_t^N(d(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t)) = \frac{1}{\overline{\mathbb{Q}}_t^N(G_t)} \overline{\mathbb{Q}}_t^N(d(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t)) G_t(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t)$$

where \mathbb{Q}_t^N and $\overline{\mathbb{Q}}_t^N$ are two *random* probability measures defined over $\mathcal{X} \times \mathcal{X}$, a set of dimension $2d$. In particular, the generation of random variables $a_{t-1}^{1:N}$ and $\mathbf{x}_t^{1:N}$ in Steps (a) and (b) of Algorithm 1 is equivalent to sampling N times independently random variables $(\tilde{\mathbf{x}}_{t-1}^n, \mathbf{x}_t^n)$ from $\overline{\mathbb{Q}}_t^N(d(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t))$: i.e. $\tilde{\mathbf{x}}_{t-1}^n = \mathbf{x}_{t-1}^{a_{t-1}^n}$ (not to be mistaken with \mathbf{x}_{t-1}^n), and $\mathbf{x}_t^n \sim m_t(\tilde{\mathbf{x}}_{t-1}^n, d\mathbf{x}_t)$.

Based on these remarks, the general idea behind SQMC is to replace at iteration t the N IID random numbers sampled from $\overline{\mathbb{Q}}_t^N(d(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t))$ by a low-discrepancy point set relative to the same distribution.

When $d = 1$, this idea may be implemented as follows: generate a low-discrepancy point set $\mathbf{u}_t^{1:N}$ in $[0, 1]^2$, let $\mathbf{u}_t^n = (u_t^n, v_t^n)$, then set $\tilde{\mathbf{x}}_{t-1}^n = \hat{F}_N^{-1}(u_t^n)$, $\mathbf{x}_t^n = \Gamma_t(\tilde{\mathbf{x}}_{t-1}^n, v_t^n)$, where \hat{F}_N^{-1} is the generalised inverse of the empirical CDF

$$\hat{F}_N(x) = \sum_{n=1}^N W_{t-1}^n \mathbb{1}\{\mathbf{x}_{t-1}^n \leq x\}, \quad x \in \mathcal{X} \subset \mathbb{R}.$$

It is easy to see that the most efficient way to compute $\tilde{\mathbf{x}}_{t-1}^n = \hat{F}_N^{-1}(u_t^n)$ for all $n \in 1:N$ is (a) to sort the \mathbf{x}_{t-1}^n 's, i.e. to find permutation σ such that $\mathbf{x}_{t-1}^{\sigma(1)} \leq \dots \leq \mathbf{x}_{t-1}^{\sigma(N)}$, (b) to sort the $u_t^{1:N}$, call $\text{sort}(u_t^{1:N})$ the corresponding result; (c) to obtain $a_{t-1}^{1:N}$ as the output of Algorithm 2, with inputs $\text{sort}(u_t^{1:N})$ and $W_t^{\sigma(1:N)}$; and finally (d) set $\tilde{\mathbf{x}}_{t-1}^n = \mathbf{x}_{t-1}^{a_{t-1}^n}$.

Algorithm 3 gives a pseudo-code version of SQMC for any d , but note how Step (b) at times $t \geq 1$ simplifies to what we have just described for $d = 1$.

When $d > 1$, the inverse transform method cannot be used to sample from the marginal distribution of $\tilde{\mathbf{x}}_{t-1}$ relative to $\overline{\mathbb{Q}}_t^N(d(\tilde{\mathbf{x}}_{t-1}, \mathbf{x}_t))$, at least unless the \mathbf{x}_{t-1}^n are “projected” to the real line in some sense. This is the point of the Hilbert curve presented in the next section.

2.3. The Hilbert space-filling curve

The Hilbert curve is a continuous fractal map $H : [0, 1] \rightarrow [0, 1]^d$, which “fills” entirely $[0, 1]^d$. H is obtained as the limit of a sequence (H_m) , $m \rightarrow +\infty$, the first terms of which are depicted in Figure 2.

The function H admits a pseudo-inverse $h : [0, 1]^d \rightarrow [0, 1]$, i.e. $H \circ h(\mathbf{x}) = \mathbf{x}$ for all $\mathbf{x} \in [0, 1]^d$. H is not a bijection because certain points $\mathbf{x} \in [0, 1]^d$ have more than one pre-image through H ; however the set of such points is of Lebesgue measure 0.




Figure 2: H_m curve for $d = 2$ and $m = 1$ to $m = 6$ (Source: Wikipedia)

Informally, H transforms $[0, 1]$ into $[0, 1]^d$, while preserving “locality”: if $x, x' \in [0, 1]$ are close, then $H(x)$ and $H(x')$ are close as well. We will establish that h also preserves discrepancy: a low-discrepancy point set in $[0, 1]^d$ remains a low-discrepancy point set in $[0, 1]$ when transformed through h . It is these properties that give to the Hilbert curve its appeal in the SQMC context (as opposed to other space filling curves, such as Z-ordering). We refer to Sagan (1994), Butz (1969) and Hamilton and Rau-Chaplin (2008) for how to compute h in practice for any $d \geq 2$. For $d = 1$, we simply set $H(x) = h(x) = x$ for $x \in [0, 1]$.

The following technical properties of H and H_m will be useful later (but may be skipped on first reading). For $m \geq 0$, let $\mathcal{I}_m^d = \{I_m^d(k)\}_{k=0}^{2^{md}-1}$ be the collection of consecutive closed intervals in $[0, 1]$ of equal size 2^{-md} and such that $\cup \mathcal{I}_m^d = [0, 1]$. For $k \geq 0$, $S_m^d(k) = H_m(I_m^d(k))$ belongs to \mathcal{S}_m^d , the set of the 2^{md} closed hypercubes of volume 2^{-md} that covers $[0, 1]^d$, $\cup \mathcal{S}_m^d = [0, 1]^d$; $S_m^d(k)$ and $S_m^d(k+1)$ are adjacent, i.e. have at least one edge in common (*adjacency property*). If we split $I_m^d(k)$ into the 2^d successive closed intervals $I_{m+1}^d(k_i)$, $k_i = 2^d k + i$ and $i \in 0 : 2^d - 1$, then the $S_{m+1}^d(k_i)$'s are simply the splitting of $S_m^d(k)$ into 2^d closed hypercubes of volume $2^{-d(m+1)}$ (*nesting property*). Finally, the limit H of H_m has the *bi-measure property*: $\lambda_1(A) = \lambda_d(H(A))$ for any measurable set $A \subset [0, 1]$, and satisfies the *Hölder condition* $\|H(x_1) - H(x_2)\|_\infty \leq C_H |x_1 - x_2|^{1/d}$ for any $x_1, x_2 \in [0, 1]$.

2.4. SQMC for $d \geq 2$

Assume now $d \geq 2$, and consider the following change of variables at iteration $t \geq 1$:

$$h_{t-1}^n = h \circ \psi(\mathbf{x}_{t-1}^n) \in [0, 1]$$

where $h : [0, 1]^d \rightarrow [0, 1]$ is the inverse of the Hilbert curve defined in the previous section, and $\psi : \mathcal{X} \rightarrow [0, 1]^d$ is some user-chosen bijection between \mathcal{X} and $\psi(\mathcal{X}) \subset [0, 1]^d$. To preserve the low discrepancy property of $\mathbf{x}_{t-1}^{1:N}$ it is important to choose

Algorithm 3 SQMC algorithm

At time $t = 0$,

- (a) Generate a QMC or a RQMC point set $\mathbf{u}_0^{1:N}$ in $[0, 1]^d$, and compute $\mathbf{x}_0^n = \Gamma_0(\mathbf{u}_0^n)$ for each $n \in 1:N$.
- (b) Compute $w_0^n = G_0(\mathbf{x}_0^n)$ and $W_0^n = w_0^n / \sum_{m=1}^N w_0^m$ for each $n \in 1:N$.

Iteratively, from time $t = 1$ to time $t = T$,

- (a) Generate a QMC or a RQMC point set $\mathbf{u}_t^{1:N}$ in $[0, 1]^{d+1}$; let $\mathbf{u}_t^n = (u_t^n, \mathbf{v}_t^n) \in [0, 1] \times [0, 1]^d$.
 - (b) Hilbert sort: find permutation σ_{t-1} such that $h \circ \psi(\mathbf{x}_{t-1}^{\sigma_{t-1}(1)}) \leq \dots \leq h \circ \psi(\mathbf{x}_{t-1}^{\sigma_{t-1}(N)})$ if $d \geq 2$, or $\mathbf{x}_{t-1}^{\sigma_{t-1}(1)} \leq \dots \leq \mathbf{x}_{t-1}^{\sigma_{t-1}(N)}$ if $d = 1$.
 - (c) Find permutation τ such that $u_t^{\tau(1)} \leq \dots \leq u_t^{\tau(N)}$, generate $a_{t-1}^{1:N}$ using Algorithm 2, with inputs $u_t^{\tau(1:N)}$ and $W_{t-1}^{\sigma_{t-1}(1:N)}$, and compute $\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^{\sigma_{t-1}(a_{t-1}^n)}, \mathbf{v}_t^{\tau(n)})$ for each $n \in 1:N$.
 - (e) Compute $w_t^n = G_t(\mathbf{x}_{t-1}^{\sigma_{t-1}(a_{t-1}^n)}, \mathbf{x}_t^n)$, and $W_t^n = w_t^n / \sum_{m=1}^N w_t^m$ for each $n \in 1:N$.
-

for ψ a mapping which is discrepancy preserving. This requires to select ψ such that $\psi(\mathbf{x}) = (\psi_1(x_1), \dots, \psi_d(x_d))$ where the ψ_i 's are continuous and strictly monotone. But choosing such a ψ is trivial in most applications; e.g. apply the logistic transformation component-wise when $\mathcal{X} = \mathbb{R}^d$ (see Section 5 for more details).

With this change of variables, we obtain particles $h_{t-1}^{1:N}$ that lie in $[0, 1]$, and (5) becomes

$$\overline{\mathbb{Q}}_{t,h}^N(d(\tilde{h}_{t-1}, \mathbf{x}_t)) = \sum_{n=1}^N W_{t-1}^n \delta_{h_{t-1}^n}(d\tilde{h}_{t-1}) m_t(\mathbf{x}_{t-1}^n, d\mathbf{x}_t).$$

Sampling a low-discrepancy sequence from $\overline{\mathbb{Q}}_{t,h}^N(d\tilde{h}_{t-1}, d\mathbf{x}_t)$ may then proceed exactly as for $d = 1$; that is: use the inverse transform method to sample N points $\tilde{h}_{t-1}^{1:N}$ from the marginal distribution $\overline{\mathbb{Q}}_{t,h}^N(d\tilde{h}_{t-1})$, then sample $\mathbf{x}_t^{1:N}$ conditionally on $\tilde{\mathbf{x}}_{t-1}^{1:N}$, with $\tilde{\mathbf{x}}_{t-1}^n = \psi^{-1} \circ H(\tilde{h}_{t-1}^n)$. The exact details of the corresponding operations are the same as for $d = 1$. We therefore obtain the general SQMC algorithm as described in Algorithm 3.

To fully define SQMC, one must choose a particular method to generate point sets $\mathbf{u}_t^{1:N}$ at each iteration. If QMC point sets are generated, one obtains a deterministic algorithm, while if RQMC point sets are generated, one obtains a stochastic algorithm.

2.5. Complexity of SQMC

The complexity of both Steps (b) and (c) (for $t \geq 1$) of the SQMC algorithm is $\mathcal{O}(N \log N)$, because they include a sort operation. The complexity of Step (a) depends on the chosen method for generating the point sets $\mathbf{u}_t^{1:N}$. For instance, Hong and Hickernell (2003) propose a $\mathcal{O}(N \log N)$ method that applies to most constructions of (t, s) -sequences (such as the Faure, the Sobol', the Niederreiter or the Niederreiter-Xing sequences). The cost to randomize a QMC point set is only $\mathcal{O}(N)$ if one chooses the simple random shift approach, while nested scrambling methods for (t, s) -sequences, which are such that all the results below hold, may be implemented at cost $\mathcal{O}(N \log N)$ (Owen, 1995; Hong and Hickernell, 2003).

To summarise, the overall complexity of SQMC is $\mathcal{O}(N \log N)$, provided the method to generate the point sets $\mathbf{u}_t^{1:N}$ is chosen appropriately.

3. Convergence study

We concentrate on two types of asymptotic results (as $N \rightarrow +\infty$): consistency, and stochastic bounds, that is bounds on the mean square error for the randomized SQMC algorithm (i.e. SQMC based on randomized QMC point sets). We leave deterministic bounds of the error (for when deterministic QMC point sets are used) to future work. We find stochastic bounds more interesting, because (a) results from (Owen, 1997a,b, 1998) suggest one might obtain better convergence rates than for deterministic bounds; and (b) the randomized version of SQMC has more applications, as discussed in Section 4.

These results are specialised to the case where the simulation of \mathbf{x}_t^n at time t is based on the inverse transform method, as explained in Section 3.1. Certain of our results require \mathcal{X} to be bounded, so for simplicity we take $\mathcal{X} = [0, 1]^d$, and ψ is set to the identity function. (Recall that, to deal with certain QMC technicalities, we follow the standard practice of taking $\mathcal{X} = [0, 1]^d$ rather than $\mathcal{X} = [0, 1]^d$.) The fact that \mathcal{X} is bounded may not be such a strong restriction, as our results allow for unbounded test functions φ ; thus, one may accommodate for an unbounded state space (and expectations with respect to that space) through appropriate variable transforms.

We introduce the following extreme norm. For any signed measure μ over $\mathcal{X} = [0, 1]^d$,

$$\|\mu\|_{\mathbb{E}} = \sup_{B \in \mathcal{B}_{[0,1]^d}} |\mu(B)|, \quad \mathcal{B}_{[0,1]^d} = \left\{ B : B = \prod_{i=1}^d [a_i, b_i] \subset [0, 1]^d, a_i < b_i \right\}$$

which generalises the extreme discrepancy in the following sense:

$$\|\mathcal{S}(\mathbf{x}^{1:N}) - \lambda_d\|_{\mathbb{E}} = D(\mathbf{x}^{1:N})$$

for any point set $\mathbf{x}^{1:N}$ in \mathcal{X} , where \mathcal{S} is the operator that associates to $\mathbf{x}^{1:N}$ its empirical distribution:

$$\mathbf{x}^{1:N} \in \mathcal{X}^N \mapsto \mathcal{S}(\mathbf{x}^{1:N}) = \frac{1}{N} \sum_{n=1}^N \delta_{\mathbf{x}^n}.$$

Our consistency results will be stated with this norm. Note that $\|\pi^N - \pi\|_{\mathbb{E}} \rightarrow 0$ implies $|\pi^N(\varphi) - \pi(\varphi)| \rightarrow 0$ for any continuous, bounded function φ , by portmanteau lemma (Van der Vaart, 2007, Lemma 2.2).

The next subsection explains how the inverse method may be used to generate \mathbf{x}_t^n given \mathbf{x}_{t-1}^n . The two following subsections state preliminary results that should provide insights on the main ideas that underpin the proofs of our convergence results. Readers interested mostly in the main results may skip these subsections and go directly to Section 3.4 (consistency) and Section 3.5 (stochastic bounds).

This section will use the following standard notations: $\|\varphi\|_{\infty}$ for the supremum norm for functions φ , $L_2(\mathcal{X}, \mu)$ for the set of square integrable functions $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ and $\mathcal{C}_b(\mathcal{X})$ for the set of continuous, bounded functions $\varphi : \mathcal{X} \rightarrow \mathbb{R}$.

3.1. Inverse transform method

We discuss here how to write the simulation of \mathbf{x}_t^n as $\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^n, \mathbf{u}_t^n)$, using the inverse transform method. Our convergence results are specialised to this particular Γ_t .

For a generic distribution $\pi \in \mathcal{P}(\mathcal{X})$, $\mathcal{X} \subset \mathbb{R}^d$, let F_{π} be the Rosenblatt transformation (Rosenblatt, 1952) of π defined through the following chain rule decomposition:

$$F_{\pi}(\mathbf{x}) = (u_1, \dots, u_d)^T, \quad \mathbf{x} = (x_1, \dots, x_d)^T \in \mathcal{X},$$

where, recursively, $u_1 = F_{\pi,1}(x_1)$, $F_{\pi,1}$ being the CDF of the marginal distribution of the first component (relative to π), and for $i \geq 2$, $u_i = F_{\pi,i}(x_i|x_{1:i-1})$, $F_{\pi,i}(\cdot|x_{1:i-1})$ being the CDF of component x_i , conditional on (x_1, \dots, x_{i-1}) , again relative to π . Similarly, we define the multivariate GICDF (generalised inverse CDF) F_{π}^{-1} through the following chain rule decomposition:

$$F_{\pi}^{-1}(\mathbf{u}) = (x_1, \dots, x_d)^T, \quad \mathbf{u} = (u_1, \dots, u_d)^T \in [0, 1]^d,$$

where, recursively, $x_1 = F_{\pi,1}^{-1}(u_1)$, $F_{\pi,1}^{-1}$ being the GICDF of the marginal distribution of the first component (relative to π), and for $i \geq 2$, $x_i = F_{\pi,i}^{-1}(u_i|x_{1:i-1})$, $F_{\pi,i}^{-1}(\cdot|x_{1:i-1})$ being the GICDF of component x_i , conditional on (x_1, \dots, x_{i-1}) , again relative to π . Note that this function depends on the particular order of the components of π . For some probability kernel $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{X})$, define similarly $F_K(\mathbf{x}, \cdot)$ and $F_K^{-1}(\mathbf{x}, \cdot)$ as, respectively, the Rosenblatt transformation and the multivariate GICDF of distribution $K(\mathbf{x}, d\mathbf{x}')$ for a fixed \mathbf{x} .

It is well known that taking $\Gamma_0 = F_{m_0}^{-1}$, and $\Gamma_t = F_{m_t}^{-1}$ lead to valid simulations algorithms, i.e. if $\mathbf{x}_0^n = F_{m_0}^{-1}(\mathbf{u}_0^n)$, resp. $\mathbf{x}_t^n = F_{m_t}^{-1}(\mathbf{x}_{t-1}^n, \mathbf{u}_t^n)$, then $\mathbf{x}_0^n \sim m_0(d\mathbf{x}_0)$, resp. $\mathbf{x}_t^n | \mathbf{x}_{t-1}^n \sim m_t(\mathbf{x}_{t-1}^n, d\mathbf{x}_t)$.

3.2. Preliminary results: importance sampling

Since SQMC is based on importance sampling (e.g. Iteration 0 of Algorithm 3), we need to establish the validity of importance sampling based on low-discrepancy point sets; see Götze (2002); Aistleitner and Dick (2014) for other results on QMC-based importance sampling.

Theorem 1. Let π and q be two probability measures on $[0, 1]^d$ such that the Radon-Nikodym derivative $w(\mathbf{x}) = \pi(d\mathbf{x})/q(d\mathbf{x})$ is continuous and bounded. Let $(\mathbf{x}^{1:N})$ be a sequence of point sets in $[0, 1]^d$ such that $\|\mathcal{S}(\mathbf{x}^{1:N}) - q\|_{\mathbb{E}} \rightarrow 0$ as $N \rightarrow +\infty$, and define

$$\pi^N = \sum_{n=1}^N W^n \delta_{\mathbf{x}^n}, \quad W^n = \frac{w(\mathbf{x}^n)}{\sum_{m=1}^N w(\mathbf{x}^m)}.$$

Then, $\|\pi^N - \pi\|_{\mathbb{E}} \rightarrow 0$ as $N \rightarrow +\infty$.

See Section A.1.2 of the Appendix for a proof.

Recall that in our notations we drop the dependence of point sets on N , i.e. we write $(\mathbf{x}^{1:N})$ rather than $(\mathbf{x}^{N,1:N})$, although in full generality $\mathbf{x}^{1:N}$ may not necessarily be the N first points of a fixed sequence.

The next theorem gives the stochastic error rate when a RQMC point set is used.

Theorem 2. Consider the set-up of Theorem 1. Let $(\mathbf{u}^{1:N})$ be a sequence of random point sets in $[0, 1]^d$ such that $\mathbf{u}^n \sim \mathcal{U}([0, 1]^d)$ marginally and, $\forall \varphi \in L_2([0, 1]^d, \lambda_d)$,

$$\text{Var} \left(\frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{u}^n) \right) = o(r(N)),$$

where $r(N) \rightarrow 0$ as $N \rightarrow +\infty$. Let $\mathbf{x}^{1:N} = F_q^{-1}(\mathbf{u}^{1:N})$ and assume that either one of the following two conditions is verified:

1. F_q^{-1} is continuous and, for any $\epsilon > 0$, there exists a $N_\epsilon \in \mathbb{N}$ such that, almost surely, $D^*(\mathbf{u}^{1:N}) \leq \epsilon$, $\forall N \geq N_\epsilon$;
2. for any $\epsilon > 0$ there exists a $N_\epsilon \in \mathbb{N}$ such that, almost surely,

$$\|\mathcal{S}(\mathbf{x}^{1:N}) - q\|_{\mathbb{E}} \leq \epsilon, \quad \forall N \geq N_\epsilon.$$

Then, for all $\varphi \in L_2(\mathcal{X}, \pi)$,

$$\mathbb{E} |\pi^N(\varphi) - \pi(\varphi)| = o(r(N)^{1/2}), \quad \text{Var}(\pi^N(\varphi)) = o(r(N)).$$

See Section A.1.2 of the Appendix for a proof.

To fix ideas, note that several RQMC strategies reach the Monte Carlo error rate and therefore fulfil the assumptions above with $r(N) = N^{-1+\epsilon}$ for any $\epsilon > 0$ (see e.g. Owen, 1997a, 1998). In addition, nested scrambling methods for (t, s) -sequences in base b (Owen, 1995; Matoúšek, 1998; Hong and Hickernell, 2003) are such that $r(N) = N^{-1}$. This result is established for $N = \lambda b^m$ in Owen (1997a, 1998) and extended for an arbitrary N in Gerber (2014).

3.3. Preliminary results: Hilbert curve and discrepancy

We motivated the use of the Hilbert curve as a way to transform back and forth between $[0, 1]^d$ and $[0, 1]$ while preserving low discrepancy in some sense. This section formalises this idea.

For a probability measure π on $[0, 1]^d$, we write π_h the image by h of π . For a kernel $K : [0, 1]^d \rightarrow \mathcal{P}(\mathcal{X})$, we write $\pi_h \otimes K_h(d(h_1, \mathbf{x}_2))$ the image of $\pi \otimes K$ by the mapping $(\mathbf{x}_1, \mathbf{x}_2) \in [0, 1]^d \times \mathcal{X} \mapsto (h(\mathbf{x}_1), \mathbf{x}_2)$, where $\pi \otimes K$ denotes the joint probability measure $\pi(d\mathbf{x}_1)K(\mathbf{x}_1, d\mathbf{x}_2)$.

The following theorem is a technical result on the conversion of discrepancy through h .

Theorem 3. *Let (π^N) be a sequence of probability measure on $[0, 1]^d$ such that, $\|\pi^N - \pi\|_{\mathbb{E}} \rightarrow 0$, where $\pi(d\mathbf{x}) = \pi(\mathbf{x})\lambda_d(d\mathbf{x})$ admits a bounded probability density $\pi(\mathbf{x})$. Then*

$$\|\pi_h^N - \pi_h\|_{\mathbb{E}} \rightarrow 0, \quad \text{as } N \rightarrow +\infty.$$

See Section A.2.1 of the Appendix for a proof.

The following theorem is an extension of Hlawka and Mück (1972, ‘‘Satz 2’’), which establishes the validity, in the context of QMC, of the multivariate GICDF approach described in Section 3.1. More precisely, for a probability measure π on $[0, 1]^d$, Hlawka and Mück (1972, ‘‘Satz 2’’) show that $\|\mathcal{S}(F_\pi^{-1}(\mathbf{u}^{1:N})) - \pi\|_{\mathbb{E}} \leq cD^*(\mathbf{u}^{1:N})^{1/d}$ (under some conditions on F_π , see below).

Theorem 4. *Let $K : [0, 1]^{d_1} \rightarrow \mathcal{P}([0, 1]^{d_2})$ be a Markov kernel and assume that:*

1. *For a fixed $\mathbf{x}_1 \in [0, 1]^{d_1}$, the i -th coordinate of $F_K(\mathbf{x}_1, \mathbf{x}_2)$ is strictly increasing in $x_{2i} \in [0, 1]$, $i \in 1 : d_2$, and, viewed as a function of \mathbf{x}_1 and \mathbf{x}_2 , $F_K(\mathbf{x}_1, \mathbf{x}_2)$ is Lipschitz;*
2. *$\pi^N(d\mathbf{x}) = \sum_{n=1}^N W_N^n \delta_{\mathbf{x}_1^n}(d\mathbf{x})$, $\mathbf{x}_1^n \neq \mathbf{x}_1^m \forall n \neq m \in 1 : N$, and $\max_{n \in 1:N} W_N^n \rightarrow 0$.*
3. *The sequence (π^N) is such that $\|\pi^N - \pi\|_{\mathbb{E}} \rightarrow 0$ as $N \rightarrow +\infty$, where $\pi(d\mathbf{x}) = \pi(\mathbf{x})\lambda_{d_1}(d\mathbf{x})$ admits a strictly positive bounded density π .*

Let $(\mathbf{u}^{1:N})$, $\mathbf{u}^n = (u^n, \mathbf{v}^n) \in [0, 1]^{1+d_2}$, be a sequence of point sets in $[0, 1]^{1+d_2}$ such that $D^(\mathbf{u}^{1:N}) \rightarrow 0$ as $N \rightarrow +\infty$, and define $P_h^N = (h^{1:N}, \mathbf{x}_2^{1:N})$ where*

$$h^n = F_{\pi_h}^{-1}(u^n), \quad \tilde{\mathbf{x}}_1^n = H(h^n), \quad \mathbf{x}_2^n = F_K^{-1}(\tilde{\mathbf{x}}_1^n, \mathbf{v}^n).$$

Then

$$\|\mathcal{S}(P_h^N) - \pi_h^N \otimes K_h\|_{\mathbb{E}} \rightarrow 0, \quad \text{as } N \rightarrow +\infty.$$

See Section A.2.2 of the Appendix for a proof.

Assumption 1 regarding the regularity of the vector-valued function F_K is the main assumption of the above theorem and comes from Hlawka and Mück (1972, ‘‘Satz 2’’). It is verified as soon as kernel K admits a density that is continuously differentiable on $[0, 1]^d$ (Hlawka and Mück, 1972, p.232). Assumption 2 is a technical condition, which will always hold under the assumptions of our main results.

3.4. Consistency

We are now able to establish the consistency of SQMC; see Appendix A.3 for a proof of the following theorem. For convenience, let $F_{m_t}(\mathbf{x}_{t-1}, \mathbf{x}_t) = F_{m_0}(\mathbf{x}_0)$ when $t = 0$.

Theorem 5. *Consider the set-up of Algorithm 3 where, for all $t \in 0 : T$, $(\mathbf{u}_t^{1:N})$ is a (non random) sequence of point sets in $[0, 1]^{d_t}$, with $d_0 = d$ and $d_t = d + 1$ for $t > 0$, such that $D^*(\mathbf{u}_t^{1:N}) \rightarrow 0$ as $N \rightarrow +\infty$. Assume the following holds for all $t \in 0:T$:*

1. *The components of $\mathbf{x}_t^{1:N}$ are pairwise distinct, $\mathbf{x}_t^n \neq \mathbf{x}_t^m$ for $n \neq m$.*
2. *G_t is continuous and bounded;*
3. *$F_{m_t}(\mathbf{x}_{t-1}, \mathbf{x}_t)$ verifies Assumption 1 of Theorem 4 ;*
4. *$\mathbb{Q}_t(d\mathbf{x}_t) = p_t(\mathbf{x}_t)\lambda_d(d\mathbf{x}_t)$ where $p_t(\mathbf{x}_t)$ is a strictly positive bounded density.*

Let $\widehat{\mathbb{Q}}_t^N(d\mathbf{x}_t) = \sum_{n=1}^N W_t^n \delta_{\mathbf{x}_t^n}(d\mathbf{x}_t)$. Then, under Assumptions 1-4, as $N \rightarrow +\infty$,

$$\|\widehat{\mathbb{Q}}_t^N - \mathbb{Q}_t\|_{\mathbb{E}} \rightarrow 0, \quad \forall t \in 0:T.$$

Assumption 1 is stronger than necessary because for the result to hold it is enough that the number of identical particles does not grow too quickly as $N \rightarrow +\infty$. Note that this is a very weak restriction since Assumption 1 holds almost surely when RQMC point sets are used, since then the particles are generated from a continuous GICDF. The assumption that the weight functions (G_t) are bounded is standard in SMC literature (see e.g. Del Moral, 2004).

3.5. Stochastic bounds

Our second main result concerns stochastic bounds for the randomized version of SQMC, i.e. SQMC based on randomized point sets (\mathbf{u}_t^n) . See Section A.4 of the Appendix for a proof of the next result.

Theorem 6. *Consider the set-up of Algorithm 3 where $(\mathbf{u}_t^{1:N})$, $t \in 0 : T$, are independent sequences of random point sets in $[0, 1]^{d_t}$, with $d_0 = d$ and $d_t = d + 1$ for $t > 0$, such that, for all $t \in 0 : T$, $\mathbf{u}_t^n \sim \mathcal{U}([0, 1]^{d_t})$ marginally and*

1. *For any $\epsilon > 0$, there exists a $N_{\epsilon, t} > 0$ such that, almost surely, $D^*(\mathbf{u}_t^{1:N}) \leq \epsilon$, $\forall N \geq N_{\epsilon, t}$.*
2. *For any function $\varphi \in L_2([0, 1]^{d_t}, \lambda_{d_t})$, $\text{Var}\left(\frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{u}_t^n)\right) \leq C^* \sigma_\varphi^2 r(N)$ where $\sigma_\varphi^2 = \int \{\varphi(\mathbf{u}) - \int \varphi(\mathbf{v}) d\mathbf{v}\}^2 d\mathbf{u}$, and where both C^* and $r(N)$ do not depend on φ .*

In addition, assume that the Assumptions of Theorem 5 are verified and that $F_{m_0}^{-1}$ is continuous. Let $\varphi \in L_2([0, 1]^d, \mathbb{Q}_t)$ for all $t \in 0:T$. Then, $\forall t \in 0:T$,

$$\mathbb{E}\left|\widehat{\mathbb{Q}}_t^N(\varphi) - \mathbb{Q}_t(\varphi)\right| = \mathcal{O}\left(r(N)^{1/2}\right), \quad \text{Var}\left(\widehat{\mathbb{Q}}_t^N(\varphi)\right) = \mathcal{O}\left(r(N)\right).$$

Note that the implicit constants in the line above may depend on φ . Assumptions 1 and 2 are verified for $r(N) = N^{-1}$ if $\mathbf{u}_t^{1:N}$ is the first N points of a nested scrambled (t, s) -sequences in base $b \geq 2$. This result is established for $N = \lambda b^m$ in Owen (1997a, 1998) and can be extended to any pattern of N using Hickernell and Yue (2001, Lemma 1). Consequently, for this construction of RQMC point sets, Theorem 6 shows that the approximation error of SQMC goes to zero at least as fast as for SMC. However, contrary to the $\mathcal{O}(N^{-1})$ convergence rate of SMC, this rate for SQMC based on nested scrambled (t, s) -sequences is not exact but results from a worst case analysis. We can therefore expect to reach faster convergence on a smaller class of functions. The following result shows that it is indeed the case on the class on continuous and bounded functions; see Section A.4.4 of the Appendix for a proof.

Theorem 7. *Consider the set-up of Algorithm 3 where $(\mathbf{u}_t^{1:N})$, $t \in 0:T$, are (t, d_t) -sequences in base $b \geq 2$, with $d_0 = d$ and $d_t = d + 1$ for $t > 0$, independently scrambled such that results in Owen (1997a, 1998) hold. Let $N = \lambda b^m$, $1 \leq \lambda < b$, and assume the following holds:*

1. *Assumptions of Theorem (6) are verified;*
2. *For $t \in 1:T$, $F_{m_t}^{-1}(\mathbf{x}_{t-1}, \mathbf{x}_t)$ is a continuous function of \mathbf{x}_{t-1} .*

Let $\varphi \in \mathcal{C}_b(\mathcal{X})$. Then, $\forall t \in 0:T$,

$$\mathbb{E}|\widehat{\mathbb{Q}}_t^N(\varphi) - \mathbb{Q}_t(\varphi)| = \mathcal{o}(N^{-1/2}), \quad \text{Var}(\widehat{\mathbb{Q}}_t^N(\varphi)) = \mathcal{o}(N^{-1}).$$

Thus, for SQMC based on the first $N = \lambda b^m$ points of nested scrambled (t, s) -sequences in base b , one obtains that the stochastic error of (the random version of) SQMC converges faster than for SMC. Note that we can relax the constraint on N in Theorem 7 using Gerber (2014, Corollary 2).

4. Extensions

4.1. Unbiased estimation of evidence, PMCMC

Like SMC, the randomized version of SQMC (that is SQMC based on RQMC point sets) provides an unbiased estimator of the normalising constant Z_t of the Feynman-Kac model, see (2).

Lemma 8. *Provided that $\mathbf{u}_t^{1:N}$ is a RQMC point set in $[0, 1]^{d_t}$ for $t \in 0 : T$ (i.e. $\mathbf{u}_t^n \sim \mathcal{U}([0, 1]^{d_t})$ marginally), with $d_0 = d$ and $d_t = d + 1$ for $t > 0$, the following quantity*

$$Z_t^N = \left\{ \frac{1}{N} \sum_{n=1}^N G_0(\mathbf{x}_0^n) \right\} \prod_{s=1}^t \left\{ \frac{1}{N} \sum_{n=1}^N G_s(\mathbf{x}_{s-1}^{a_{s-1}^n}, \mathbf{x}_s^n) \right\}$$

is an unbiased estimator of Z_t , $\mathbb{E}[Z_t^N] = Z_t$.

We omit the proof, as it follows the same steps as for SMC (Del Moral, 1996).

In a state-space model parametrised by $\theta \in \Theta$, $Z_t = Z_t(\theta)$ is the marginal likelihood of the data up to time t . One may want to implement a Metropolis-Hastings sampler with respect to posterior density $\pi_T(\theta) \propto p(\theta)Z_T(\theta)$ for the full dataset and for a prior distribution $p(\theta)$, but $Z_T(\theta)$ is typically intractable.

Andrieu et al. (2010) established that, by substituting $Z_T(\theta)$ with an unbiased estimate of $Z_T(\theta)$ in a Metropolis sampler, one obtains an exact MCMC (Markov chain Monte Carlo) algorithm, in the sense that the corresponding MCMC kernel leaves invariant $\pi_T(\theta)$. The so obtained algorithm is called PMMH (Particle marginal Metropolis-Hastings). Andrieu et al. (2010) use SMC to obtain an unbiased estimate of $Z_T(\theta)$, that is, at each iteration a SMC sampler is run to obtain that estimate. We will call PMMH-SQMC the same algorithm, but with SQMC replacing SMC for the evaluation of an unbiased estimate of the likelihood.

The acceptance rate of PMMH depends directly on the variability of the estimates of $Z_T(\theta)$. Since the point of (randomized) SQMC is to provide estimates with a lower variance than SMC (for a given N), one may expect that PMMH-SQMC may require a smaller number of particles than standard PMMH for satisfactory acceptance rates; see Section 5 for a numerical illustration of this.

4.2. Smoothing

Smoothing amounts to compute expectations $\mathbb{Q}_t(\varphi)$ of functions φ of the complete trajectory $\mathbf{x}_{0:t}$; e.g. $\mathbb{Q}_t(\varphi)$ is the expectation of $\varphi(\mathbf{x}_{0:t})$ conditional on data $\mathbf{y}_{0:t}$ for a state-space model with Markov process (\mathbf{x}_t) and observed process (\mathbf{y}_t) . See Briers et al. (2010) for a general overview on SMC smoothing algorithms. This section discusses how to adapt certain of these algorithms to SQMC.

4.2.1. Forward smoothing

Forward smoothing amounts to carry forward the complete trajectories of the particles, rather than simply keeping the last component \mathbf{x}_t^n (as in Algorithm 1). A simple way to formalise forward smoothing is to introduce a path Feynman-Kac model, corresponding to the inhomogeneous Markov process $\mathbf{z}_t = \mathbf{x}_{0:t}$, and weight function (abusing notations) $G_t(\mathbf{z}_t) = G_t(\mathbf{x}_t)$. Then forward smoothing amounts to Algorithm 1 applied to this path Feynman-Kac model (substituting \mathbf{x}_t with $\mathbf{z}_t = \mathbf{x}_{0:t}$).

One may use the same remark to define a SQMC version of forward smoothing: i.e. simply apply SQMC to the same path Feynman-Kac model. The only required modification is that the Hilbert sort of Step (b) at times $t \geq 1$ must now operate on some transformation of the vectors \mathbf{z}_t^n , of dimension $(t+1)d$, rather than vectors \mathbf{x}_t^n of dimension d as in the original version.

Forward smoothing is sometimes used to approximate the smoothing expectation of additive functions, $\varphi(\mathbf{x}_{0:t}) = \sum_{s=0}^t \tilde{\varphi}(\mathbf{x}_s)$, such as the score function of certain models (e.g. Poyiadjis et al., 2011). In that case, one may instead apply SQMC to the Feynman-Kac model corresponding to the inhomogeneous Markov process $\mathbf{z}_t = (\sum_{s=0}^{t-1} \tilde{\varphi}(\mathbf{x}_s), \mathbf{x}_t)$.

Algorithm 4 Backward step of SQMC backward smoothing

Input: $\mathbf{x}_{0:T}^{\sigma_t(1:N)}$, $W_{0:T}^{\sigma_t(1:N)}$ (output of SQMC obtained after the Hilbert sort step, i.e for all $t \in 0:T$, $h \circ \psi(\mathbf{x}_t^{\sigma_t(n)}) \leq h \circ \psi(\mathbf{x}_t^{\sigma_t(m)})$, $n \leq m$) and $\tilde{\mathbf{u}}^{1:N_B}$ a point set in $[0, 1)^{T+1}$; let $\tilde{\mathbf{u}}^n = (\tilde{u}_0^n, \dots, \tilde{u}_T^n)$.

Output: $\tilde{\mathbf{x}}_{0:T}^{1:N_B}$ (N_B trajectories in \mathcal{X}^{T+1})

Find permutation τ such that $\tilde{u}_0^{\tau(1)} \leq \dots \leq \tilde{u}_0^{\tau(N_B)}$, generate $\tilde{a}_T^{1:N_B}$ using Algorithm 2, with inputs $\tilde{u}_0^{\tau(1:N_B)}$ and $W_T^{\sigma_T(1:N)}$, and set $\tilde{\mathbf{x}}_T^n = \mathbf{x}_T^{\tilde{a}_T^n}$ for all $n \in 1:N_B$.

for $t = T - 1 \rightarrow 0$ **do**

For $n \in 1:N_B$, set $\tilde{\mathbf{x}}_t^n = \mathbf{x}_t^{\tilde{a}_t^n}$ where $\tilde{a}_t^n = F_{\pi_t^n}^{-1}(\tilde{u}_{T-t}^{\tau(n)})$, $\pi_t^n = \sum_{m=1}^N \tilde{W}_t^m(\tilde{\mathbf{x}}_{t+1}^n) \delta_m$ and, for $m \in 1:N$,

$$\tilde{W}_t^m(\mathbf{x}_{t+1}) = W_t^{\sigma_t(m)} m_t(\mathbf{x}_{t+1} | \mathbf{x}_t^{\sigma_t(m)}) / \left\{ \sum_{n=1}^N W_t^n m_t(\mathbf{x}_{t+1} | \mathbf{x}_t^n) \right\}.$$

end for

This means that in practice, one may implement the Hilbert sort on a space of much lower dimension (i.e. the dimension of this new \mathbf{z}_t), which is computationally more convenient.

4.2.2. Backward smoothing

Backward smoothing consists of two steps: (a) a forward pass, where SMC is run from time 0 to time T ; and (b) a backward pass, where one constructs a trajectory $\tilde{\mathbf{x}}_{0:T}$ recursively backwards in time, by selecting randomly each component $\tilde{\mathbf{x}}_t$ out of the N particle values \mathbf{x}_t^n generated during the forward pass. An advantage of backward smoothing is that it is less prone to degenerate than forward smoothing. A drawback of backward smoothing is that generating a single trajectory costs $\mathcal{O}(N)$, hence obtaining N of them costs $\mathcal{O}(N^2)$.

Backward smoothing for SQMC may be implemented in a similar way to SMC: see Algorithm 4 for the backward pass that generates N_B trajectories $\tilde{\mathbf{x}}_{0:T}^{1:N_B}$ from the output of the SQMC algorithm. Note that backward smoothing requires that the Markov kernel $m_t(\mathbf{x}_{t-1}, d\mathbf{x}_t)$ admits a closed-form density $m_t(\mathbf{x}_t | \mathbf{x}_{t-1})$ with respect to an appropriate dominating measure. Then one may compute empirical averages over the so obtained N_B trajectories to obtain smoothing estimates in the usual way.

5. Numerical study

The objective of this section is to compare the performance of SMC and SQMC. Our comparisons are either for the same number of particles N , or for the same amount of CPU time to take into account the fact that SQMC has greater complexity than SMC. These comparisons will often summarised through gain factors, which we define as ratios

of mean square errors (for a certain quantity) between SMC and SQMC.

In SQMC, we generate $\mathbf{u}_t^{1:N}$ as a Owen (1995) nested scrambled Sobol' sequence using the C++ library of T. Kollig and A. Keller (<http://www.uni-kl.de/AG-Heinrich/SamplePack.html>). Note that both the generation and the randomization of (t, s) -sequences in base 2 (such as the Sobol' sequence) are very fast since logical operations can be used. In order to sort the particles according to their Hilbert index we use the C++ library of Chris Hamilton (<http://web.cs.dal.ca/~chamilto/hilbert/index.html>) to evaluate $H_m^{-1}(\psi(\mathbf{x}))$, $m \in \mathbb{N}$. Again, Hilbert computations are very fast as they are based on logical operations (see Hamilton and Rau-Chaplin, 2008, for more details). In addition, thanks to the nesting property of the Hilbert curve (see Section 2.3) we only need to take m large enough such that different particles are mapped into different points of $[0, 1)$. Function Γ_t is set to the inverse transform described in Section 3.1, and function ψ to a component-wise (rescaled) logistic transform; that is, $\psi(\mathbf{x}) = (\psi_1(x_1), \dots, \psi_d(x_d))$ with

$$\psi_i(x_i) = \left[1 + \exp\left(-\frac{x_i - \underline{x}_i}{\bar{x}_i - \underline{x}_i}\right) \right]^{-1}, \quad i \in 1:d$$

and where the constants \bar{x}_i and \underline{x}_i are used to solve numerical problems due to high values of $|x_i|$. For instance, when (\mathbf{x}_t) is a stationary process we chose $\bar{x}_i = \mu_i + 2\sigma_i$ and $\underline{x}_i = \mu_i - 2\sigma_i$ where μ_i and σ^2 are respectively the mean and the standard deviation of the stationary distribution of (\mathbf{x}_t) .

SMC is implemented using systematic resampling (Carpenter et al., 1999) and all the random variables are generated using standard methods (i.e. not using the multivariate GICDF). The C/C++ code implementing both SMC and SQMC is available on-line at <https://bitbucket.org/mgerber/sqmc>.

Even if Theorems 6 and 7 are valid for any pattern of N , choosing for N powers of 2 (with 2 the base of the Sobol' sequence) is both natural and optimal for QMC methods based on (scrambled) (t, s) -sequences (see e.g. Owen, 1997b; Hickernell and Yue, 2001; and Chapter 5 of Dick and Pillichshammer, 2010). Comparing the performance of SQMC for different patterns of N is beyond the scope of this paper (see Gerber, 2014, for a discussion of this point) and therefore we follow in this numerical study the standard approach in the QMC literature by considering values of N that are powers of 2. We nevertheless do one exception to this rule for the PMMH estimation on real data (Section 5.3) because doubling the number of particles to reduce the variance of the likelihood estimate used in the Metropolis-Hastings ratio may be very inefficient from a computational point of view. As we will see, allowing N to differ from powers of the Sobol' base does not seem to alter the performance of SQMC.

One may expect the two following situations to be challenging for SQMC: (a) small N (because our results are asymptotic); and (b) large d (because of the usual deterioration of QMC with respect of the dimension, and also because of the Hilbert sort step). Thus we consider examples of varying dimensions (from 1 to 10), and we will also make N vary within a large range (between 2^4 and 2^{17}).

5.1. Example 1: A non linear and non stationary univariate model

We consider the following popular toy example (Gordon et al., 1993; Kitagawa, 1996):

$$\begin{cases} y_t = \frac{x_t^2}{a} + \epsilon_t, & \epsilon_t \sim \mathcal{N}_1(0, 1), \quad t \geq 0 \\ x_t = b_1 x_{t-1} + b_2 \frac{x_{t-1}}{1+x_{t-1}^2} + b_3 \cos(b_4 t) + \sigma \nu_t, & \nu_t \sim \mathcal{N}_1(0, 1), \quad t > 0 \end{cases} \quad (6)$$

and $x_0 \sim \mathcal{N}_1(0, 2)$, where $\mathcal{N}_d(\boldsymbol{\mu}, \Sigma)$ denotes the d -dimensional Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix Σ . We generate observations from 100 time steps of the model, with the parameters set as in Gordon et al. (1993): $a = 20$, $\mathbf{b} = (0.5, 25, 8, 1.2)$, $\sigma^2 = 10$, $x_0 = 0.1$. Note that inference in this model is non trivial because the observation y_t does not allow to identify the sign of x_t , and because the weight function $G_t(x_t)$ is bimodal if $y_t > 0$ (with modes at $\pm(20y_t)^{1/2}$). In addition, we expect this model to be challenging for SQMC due to the high non linearity of the Markov transition $m_t(\mathbf{x}_{t-1}, d\mathbf{x}_t)$.

All the results presented below are based on 500 independent runs of SMC and SQMC. Figure 4 presents results concerning the estimation of the log-likelihood functions evaluated at the true value of the parameters. The two top graphs show that, compared to SMC, SQMC yields faster convergence of both the mean and the variance of the estimates.

These better consistency properties of SQMC are also illustrated on the bottom left graph of Figure 4 where we have reported for each N the range in which lies the 500 estimates of the log-likelihood. From this plot we see that quickly the SQMC estimates stay in a very tiny interval while, on the contrary, the SMC estimates are much more dispersed, even for large values of N .

The bottom right panel of Figure 4 shows the MSE of SQMC and SMC as a function of CPU time. One sees that the gain of SQMC over SMC does not only increase with N , as predicted by the theory, but also with the CPU time which is of more practical interest. On the other hand, in this particular case (log-likelihood evaluation for this univariate model), when N is small the reduction in MSE brought by SQMC does not compensate its greater running time. Nevertheless, we observe that SQMC outperforms SMC very quickly, that is, as soon as the CPU time is larger or equal to $10^{-1.5} \approx 0.03$ seconds.

In the left graph of Figure 3 we have reported the gain factor for the estimation of $\mathbb{E}[x_t|y_{0:t}]$ as a function of t and for different values of N . From this plot we observe both significant and increasing gain of SQMC over SMC.

The right panel of Figure 3 compares SQMC and SMC backward smoothing for the estimation of $\mathbb{E}[x_t|y_{0:T}]$ as a function of t and for $N \in \{2^7, 2^9\}$. As for the filtering problem, SQMC significantly outperforms SMC with gain factors that increase with the number of particles.




Figure 3: Filtering (left graph) and backward smoothing (right graph) for the toy example (6): gain factor as a function of t for the estimation of $\mathbb{E}[x_t|y_{0:t}]$ and for the estimation of $\mathbb{E}[x_t|y_{0:T}]$, obtained from 500 independent runs of SQMC and SMC.

5.2. Example 2: Multivariate stochastic volatility model

We consider the following multivariate stochastic volatility model (SV) proposed by Chan et al. (2006):

$$\begin{cases} \mathbf{y}_t = S_t^{1/2} \boldsymbol{\epsilon}_t, & t \geq 0 \\ \mathbf{x}_t = \boldsymbol{\mu} + \Phi(\mathbf{x}_{t-1} - \boldsymbol{\mu}) + \Psi^{1/2} \boldsymbol{\nu}_t, & t > 0 \end{cases} \quad (7)$$

where $S_t = \text{diag}(\exp(x_{t1}), \dots, \exp(x_{td}))$, Φ and Ψ are diagonal matrices and $(\boldsymbol{\epsilon}_t, \boldsymbol{\nu}_t) \sim \mathcal{N}_{2d}(\mathbf{0}_{2d}, C)$, with C a correlation matrix and $\mathbf{0}_{2d} = (0, \dots, 0) \in \mathbb{R}^{2d}$.

In order to study the relative performance of SQMC over SMC as the dimension d of the hidden process increases we perform simulations for $d \in \{1, 2, 4, 10\}$. The parameters we use for the simulations are the same as in Chan et al. (2006): $\phi_{ii} = 0.9$, $\mu_i = -9$, $\psi_{ii}^2 = 0.1$ for all $i = 1, \dots, d$ and

$$C = \begin{pmatrix} 0.6\mathbf{1}_d + 0.4\mathcal{I}_d & -0.1\mathbf{1}_d - 0.2\mathcal{I}_d \\ -0.1\mathbf{1}_d - 0.2\mathcal{I}_d & 0.8\mathbf{1}_d + 0.2\mathcal{I}_d \end{pmatrix}$$

where \mathcal{I}_d is the d -dimensional identity matrix, and $\mathbf{1}_d$ is the $d \times d$ matrix having one in all its entries. Note that the errors terms $\boldsymbol{\epsilon}_t$ and $\boldsymbol{\nu}_t$ are correlated so that the weight function G_t depends now both on \mathbf{x}_{t-1} and on \mathbf{x}_t . The prior distribution for \mathbf{x}_0 is the stationary distribution of the process (\mathbf{x}_t) and we take $T = 399$.

The three first panels of Figure 5 present results for the estimation of the log-likelihood (evaluated at the true value of the parameters and for the complete dataset $y_{0:T}$), for $d \in \{1, 2, 4\}$. One sees that the gain factor increases quickly with N , and, more importantly,

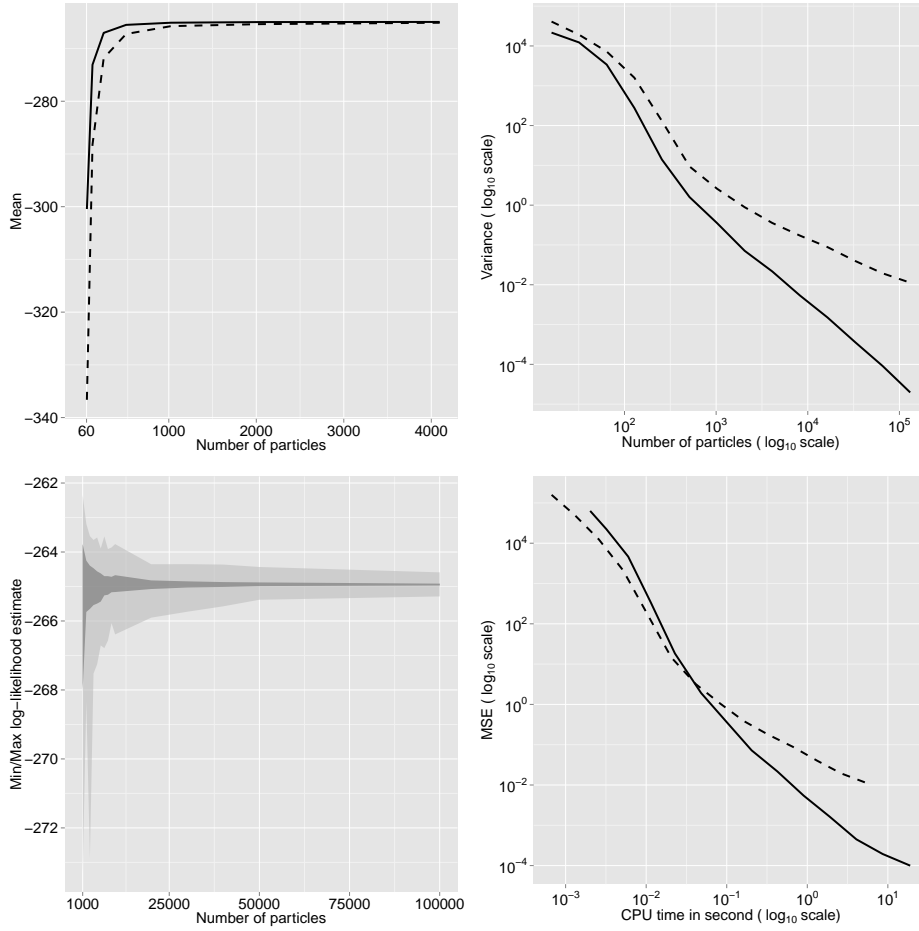


Figure 4: Log-likelihood estimation for the toy example (6). The solid lines are for SQMC while the dashed lines are for SMC. In the bottom-left graph, the dark (light) area shows the range in which lies the SQMC (SMC) estimates of the log-likelihood function. All the results are obtained from 500 independent runs of SQMC and SMC.

the MSE of SQMC converges faster than SMC even as a function of CPU time. In fact, except for a very small interval for the univariate model, SQMC always outperforms SMC in terms of MSE for the same CPU effort. We note the particularly impressive values of the gain factor we obtain for $d = 1$ when N is large: around 4.2×10^4 for $N = 2^{17}$. The last panel of Figure 5 plots the gain factors as a function of N , for same values of d , plus $d = 10$. The improvement brought by SQMC decreases with the dimension, and in fact, for $d = 10$, the gain factor is essentially one for the considered values of N ; yet for $d = 4$ we still observe some notable improvement; e.g. a gain factor of 10 for $N \approx 10^5$. We now focus on $d = 1, 2$ and 4.

Figure 6 represents the evolution with respect to t of the MSE for the partial log-

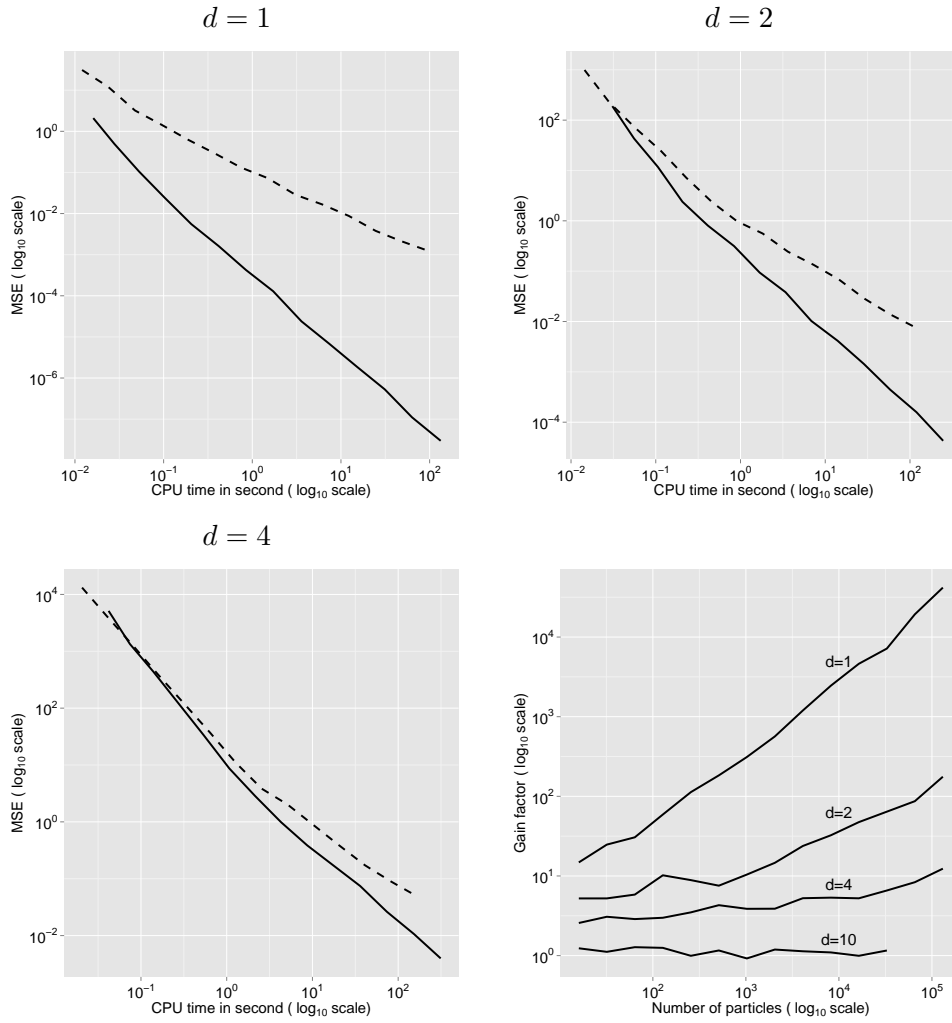


Figure 5: Log-likelihood estimation of SV model (7): MSE as a function of CPU time, for $d = 1, 2, 4$; gain factor as a function of N for $d = 1, 2, 4$ and 10. The solid lines are for SQMC while the dashed lines are for SMC. The graphs are obtained from 200 independent runs of SQMC and SMC.

likelihood of data $y_{0:t}$ up to time t ; gain factors are reported for different values of N . As we can see from these graphs, the performance of SQMC does not seem to depreciate with t .

Finally, Figure 7 shows that SQMC also give impressive gain when $d > 1$ concerning the estimation of the filtering expectation $\mathbb{E}[x_{1t}|\mathbf{y}_{0:t}]$ of the first component of \mathbf{x}_t .

5.3. Application: Bayesian estimation of MSV using PMMH on real data

To compare SMC to SQMC when used as a way to approximate the likelihood within a PMMH algorithm, as described in Section 4.1, we turn our attention to the Bayesian estimation of the multivariate SV model (7), for $d = 2$. As in Chan et al. (2006), we take the following prior:


$$\phi_{ii} \sim \mathcal{U}((0, 1)), \quad 1/\psi_{ii}^2 \sim \text{Gamma}(10 \exp(-10), 10 \exp(-3)) \quad i = 1, \dots, d,$$

where ϕ_{ii} and ψ_{ii}^2 denotes respectively the diagonal elements of Φ and Ψ , and a flat prior for $\boldsymbol{\mu}$. In addition, we assume that C is uniformly distributed on the space of correlation matrices which are such that the errors terms $\boldsymbol{\epsilon}_t$ and $\boldsymbol{\nu}_t$ are independents (no leverage effects). To sample from the posterior distribution of the parameters we use a Gaussian random walk Metropolis-Hastings algorithm with covariance matrix Σ calibrated so that the acceptance probability of the algorithm becomes, as $N \rightarrow +\infty$, close to 25%. The matrix Σ , as well as the starting point of the Markov chain, are calibrated using a pilot run of the algorithm with $\Sigma = 0.011^2 \mathcal{I}_8$ and starting at the value of the parameters we used above for the simulations. To compare PMMH-SQMC with PMMH-SMC, we run the two algorithms during 10^5 iterations and for values of N ranging from 10 to 200, where N increases from 10 to 100 by increment of 10 and then by increment of 50.


We consider the following dataset: the two series are the mean-corrected daily return on the Nasdaq and S&P 500 indices for the period ranging from the 3rd January 2012 to the 21th October 2013 so that the data set contains 452 observations.

Figure 8 shows the Metropolis-Hastings acceptance rate and the effective sample sizes (see Robert and Casella, 2004, Section 12.3.5, for a definition) for the PMMH-SQMC algorithm and for the standard PMMH algorithm. We first observe that the acceptance rate of PMMH-SQMC increases very quickly with N . Indeed, it is already of 20% for only 30 particles while for the same number of particles the acceptance rate for the standard PMMH is approximatively 6.5%. As far as the acceptance rate is concerned, there is no significant gain to take $N > 60$ for the PMMH-SQMC algorithm while for the plain Monte Carlo algorithm the acceptance rate is only about 20% for $N = 200$ and therefore much smaller than the target of 25%. Looking at the results for the effective sample sizes (ESSs), we see that the same conclusions hold. More precisely, for the PMMH-SQMC algorithm, the ESSs increase with N much faster than for PMMH-SMC. Indeed, for $N \in 10 : 50$, the ESSs for the former is between 2.18 and 14.94 times larger than for PMMH-SMC.

Univariate SV model



Bivariate SV model



Four dimensional SV model




Figure 6: Log-likelihood estimation of the SV model (7): gain factor as a function of t , obtained from 200 independent runs of SQMC and SMC.

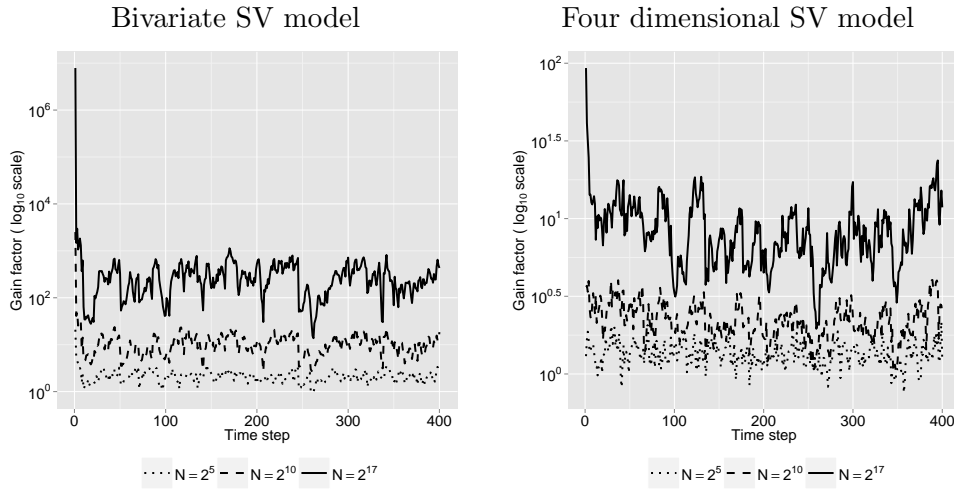


Figure 7: Filtering of the multivariate SV model (7): gain factor as a function of t for the estimation of $\mathbb{E}[x_{1t}|\mathbf{y}_{0:t}]$, obtained from 200 independent runs of SQMC and SMC.




Figure 8: Metropolis-Hastings acceptance rate and effective sample sizes (one per parameter) for the multivariate SV model (7), $d = 2$, and real data: the solid lines are for PMMH-SQMC while the dashed ones are for PMMH-SMC. The results are obtained from a Markov Chain of length 10^5 .

5.4. Example 3: Neural decoding

Neural decoding models are used for brain-machine interface in order to make inference about an organism’s environment from its neural activity. More precisely, we consider the problem of decoding a set of environment variables $\mathbf{p}_t \in \mathbb{R}^2$, from the firing ensemble of d_y neurons. The latent vector \mathbf{p}_t may be interpreted as two-dimensional hand kinetics for motor cortical decoding (see Koyama et al., 2010, and references therein for more details about neural decoding models). Noting $\dot{\mathbf{p}}_t$ the vector of velocities, the neural decoding model we consider is given by (Koyama et al., 2010)

$$\begin{cases} y_{ti} | \mathbf{x}_{0:t} \sim \mathcal{P}(\Delta \exp(\alpha_i + \beta_i^T \mathbf{x}_t)), & i \in 1 : d_y, & t \geq 0 \\ \mathbf{x}_t = \Phi \mathbf{x}_{t-1} + \Psi \boldsymbol{\epsilon}_t, & \boldsymbol{\epsilon}_t \sim \mathcal{N}_2(\mathbf{0}_2, \sigma^2 \mathcal{I}_2), & t > 0 \end{cases} \quad (8)$$

and $\mathbf{x}_0 \sim \mathcal{N}_4(\mathbf{0}_4, \mathcal{I}_4)$, where $\mathbf{x}_t = (\mathbf{p}_t, \dot{\mathbf{p}}_t)$, the y_{ti} ’s are conditionally independent, $\mathcal{P}(\lambda)$ denotes the Poisson distribution with parameter λ , Δ is the duration of the interval over which spikes are counted at each time step, and

$$\Phi = \begin{pmatrix} \mathcal{I}_2 & \Delta \mathcal{I}_2 \\ \mathbf{0}_2 & \mathcal{I}_2 \end{pmatrix}, \quad \Psi^T = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Realistic values for the parameters, see Koyama et al. (2010), that we will take in our simulations, are $d_y = 10$, $T = 23$, $\Delta = 0.03$, $\sigma^2 = 0.019$, $\alpha_i \stackrel{i.i.d.}{\sim} \mathcal{N}_1(2.5, 1)$, $\beta_i \sim \mathcal{U}([0, 1]^d)$.

One important aspect of this model is that the dimension of the noise term $\boldsymbol{\epsilon}_t$ is lower than the dimension of \mathbf{x}_t . As a result, two components of \mathbf{x}_t are deterministic functions of \mathbf{x}_{t-1} . Many tracking problems have a similar structure.

This requires us to slightly adapt SQMC as follows: one samples jointly the ancestor variables $a_{t-1}^{1:N}$ and the new velocities $\dot{\mathbf{p}}_t^n$ as in Steps (b) and (c) of Algorithm 3, then one obtains the new \mathbf{p}_t^n as $\mathbf{p}_t^n = \mathbf{p}_{t-1}^n + \dot{\mathbf{p}}_t^n$, i.e. the deterministic linear transformation of $\mathbf{p}_{t-1}^{a_{t-1}^n}$ and $\dot{\mathbf{p}}_{t-1}^n$ defined by the model. Note that in this case the dimension of the point set $\mathbf{u}_t^{1:N}$ is 3 for $t > 0$; we could say that $d = 2$ in this case, even if the dimension of \mathbf{x}_t itself is 4.

Figures 9 and 10 present, respectively, results for the estimation of the log-likelihood (evaluated at the true value of the parameters) and for the estimation of the filtering expectation $\mathbb{E}[x_{ti} | \mathbf{y}_{0:t}]$ for $i \in 1 : d$. Concerning the log-likelihood estimation we observe fast increase of the gain factor after about 2^{11} particles with a maximum close to 21 when N is very large. The gain of SQMC compensates its longer running time after only about 0.17 seconds. Important and increasing (in N) gains are also observed for the estimation of the filtering expectations.

6. Conclusion and future work

The main message of the paper is that SMC users should be strongly encouraged to switch to SQMC, as SQMC is “typically” much more accurate (produces estimates with smaller errors) than SMC. We add the word “typically” to recall that our asymptotic




Figure 9: Log-likelihood estimation of the neural decoding model (8). The left graph gives the ratio of the SMC and the SQMC MSE. In the right graph, the solid line is for SQMC while the dashed line is for SMC. The graphs are obtained from 200 independent runs of SQMC and SMC.

analysis, by construction, proves only that the SQMC error is smaller than the SMC error *for N large enough*. But our range of numerical examples, which are representative of real-world filtering problems, makes us optimistic that in most practical cases SQMC should outperform SMC even for moderate values of N .

The main price to pay to switch to SQMC is that users should spend some time thinking on how to write the simulation of \mathbf{x}_t^n given \mathbf{x}_{t-1}^n as $\mathbf{x}_t^n = \Gamma_t(\mathbf{x}_{t-1}^n, \mathbf{u}_t^n)$, where $\mathbf{u}_t^n \sim \mathcal{U}([0, 1]^d)$ and Γ_t is a deterministic function that is easy to evaluate. Fortunately, this is often straightforward. In fact, there are many models of interest where \mathbf{x}_t^n given \mathbf{x}_{t-1}^n is linear and Gaussian. Since this case is already implemented in our program, adapting it to such a model should be just a matter of changing a few lines of code (to evaluate the probability density of \mathbf{y}_t given \mathbf{x}_t).

Regarding future work, the most pressing tasks seem (a) to refine the convergence rate of the SQMC error; and (b) to establish that it does not degenerate over time (in the spirit of time-uniform estimates for SMC, see p. 244 of Del Moral, 2004). Regarding the former, we make the interesting conjecture that the mean square error of SQMC converges at rate $\mathcal{O}(N^{-1-2/d})$. This would explain why the relative performance of SQMC decreases with the dimension. Fortunately, a majority of the state space models of interest in signal processing, finance, or other fields are such that $d \leq 6$. A notable exception is geophysical data assimilation (in e.g. meteorology or oceanography) for which d can be very large, but for such large-dimensional problems SMC seems to perform too poorly for practical use anyway (Bocquet et al., 2010).

Finally, it is also our hope that this paper will help QMC garner wider recognition in Bayesian computation and related fields. Granted, QMC is more technical than standard Monte Carlo, and there is perhaps something specific about particle filtering that makes




Figure 10: Filtering of the Neuro decoding model (8). From left to right and from top to bottom, the graphs give the ratio of the SMC and SQMC MSE for the estimation of $\mathbb{E}[x_{kt}|\mathbf{y}_{0:t}]$ as a function of t , $k = 1, \dots, 4$, and are obtained from 200 independent runs of SQMC and SMC.

the introduction of QMC so effective. Yet we cannot help but think that the full potential of QMC in Statistics remains under-explored.

Acknowledgements

We thank the referees, Christophe Andrieu, Simon Barthelmé, Arnaud Doucet, Paul Fearnhead, Simon Lacoste-Julien, and Art Owen for excellent remarks that helped us to greatly improve the paper. The second author is partially supported by a grant from the French National Research Agency (ANR) as part of the “Investissements d’Avenir” program (ANR-11-LABEX-0047).

References

- Aistleitner, C. and Dick, J. (2014). Functions of bounded variation, signed measures, and a general Koksma-Hlawka inequality. *arXiv:1406.0230*.
- Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *J. R. Statist. Soc. B*, 72(3):269–342.
- Barvínek, E., Daler, I., and Francu, J. (1991). Convergence of sequences of inverse functions. *Archivum Mathematicum*, 27(3-4):2001–204.
- Bocquet, M., Pires, C. A., and Wu, L. (2010). Beyond Gaussian statistical modeling in geophysical data assimilation. *Monthly Weather Review*, 138(8):2997–3023.
- Briers, M., Doucet, A., and Maskell, S. (2010). Smoothing algorithms for state-space models. *Ann. of the Inst. of Stat. Math.*, 62(1):61–89.
- Butz, A. R. (1969). Convergence with Hilbert’s space filling curve. *Journal of Computer and System Science*, 3(2):128–146.
- Cappé, O., Moulines, E., and Rydén, T. (2005). *Inference in Hidden Markov Models*. Springer-Verlag, New York.
- Carpenter, J., Clifford, P., and Fearnhead, P. (1999). Improved particle filter for non-linear problems. *IEE Proc. Radar, Sonar Navigation*, 146(1):2–7.
- Chan, D., Kohn, R., and Kirby, C. (2006). Multivariate stochastic volatility models with correlated errors. *Econometric reviews*, 25(2-3):245–274.
- Chen, S., Dick, J., and Owen, A. B. (2011). Consistency of Markov chain quasi-Monte Carlo on continuous state spaces. *Ann. Stat.*, 39(2):673–701.
- Chopin, N. (2002). A sequential particle filter for static models. *Biometrika*, 89:539–552.
- Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Stat.*, 32(6):2385–2411.

- Cranley, R. and Patterson, T. (1976). Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, 13(6):904–914.
- Del Moral, P. (1996). Non-linear filtering: interacting particle resolution. *Markov processes and related fields*, 2(4):555–581.
- Del Moral, P. (2004). *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems with Applications*. Probability and its Applications. Springer Verlag, New York.
- Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. *J. R. Statist. Soc. B*, 68(3):411–436.
- Del Moral, P. and Guionnet, A. (1999). Central limit theorem for nonlinear filtering and interacting particle systems. *Ann. Appl. Prob.*, 9:275–297.
- Devroye, L. (1986). *Non-Uniform Random Variate Generation*. Springer-Verlag, New York.
- Dick, J. and Pillichshammer, F. (2010). *Digital nets and sequences: discrepancy theory and quasi-Monte Carlo integration*. Cambridge University Press.
- Doucet, A., de Freitas, N., and Gordon, N. J. (2001). *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York.
- Fearnhead, P. (2005). Using random quasi-Monte Carlo within particle filters, with application to financial time series. *J. Comput. Graph. Statist.*, 14(4):751–769.
- Gerber, M. (2014). On integration methods based on scrambled nets of arbitrary size. *ArXiv preprint*.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer Verlag.
- Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. F, Comm., Radar, Signal Proc.*, 140(2):107–113.
- Götz, M. (2002). Discrepancy and the error in integration. *Monatsh. Math.*, 136(2):99–121.
- Guo, D. and Wang, X. (2006). Quasi-monte carlo filtering in nonlinear dynamic systems. *Signal Processing, IEEE Transactions on*, 54(6):2087–2098.
- Hamilton, C. H. and Rau-Chaplin, A. (2008). Compact Hilbert indices for multi-dimensional data. In *Proceedings of the First International Conference on Complex, Intelligent and Software Intensive Systems*.
- Hickernell, F. J. and Yue, R.-X. (2001). The mean square discrepancy of scrambled (t, s) -sequences. *SIAM Journal on Numerical Analysis*, 38:1089–1112.

- Hlawka, E. and Mück, R. (1972). Über eine transformation von gleichverteilten folgen II. *Computing*, 9:127–138.
- Hong, H. S. and Hickernell, F. J. (2003). Algorithm 823: Implementing scrambled digital sequences. *ACM Trans. Math. Softw.*, 29(2):95–109.
- Johansen, A. M. and Doucet, A. (2008). A note on auxiliary particle filters. *Statist. Prob. Letters*, 78(12):1498–1504.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Statist.*, 5:1–25.
- Koyama, S., Castellanos Pérez-Bolde, L., Shalizi, C. R., and Kass, R. E. (2010). Approximate methods for state-space models. *J. Am. Statist. Assoc.*, 105(489):170–180.
- Kuipers, L. and Niederreiter, H. (1974). *Uniform distribution of sequences*. Wiley-Interscience.
- Künsch, H. R. (2005). Recursive Monte Carlo filters: Algorithms and theoretical analysis. *Ann. Stat.*, 33:1983–2021.
- Lécot, C. and Ogawa, S. (2002). Quasirandom walk methods. In *Monte Carlo and Quasi-Monte Carlo Methods 2000*, pages 63–85. Springer.
- Lécot, C. and Tuffin, B. (2004). Quasi-monte carlo methods for estimating transient measures of discrete time markov chains. In *Monte Carlo and Quasi-Monte Carlo Methods 2002*, pages 329–343. Springer.
- L’Ecuyer, P., Lécot, C., and L’Archevêque-Gaudet, A. (2009). On array-RQMC for Markov Chains: Mapping alternatives and convergence rates. In *Monte Carlo and quasi-Monte Carlo methods 2008*, pages 485–500. Springer Berlin Heidelberg.
- L’Ecuyer, P., Lécot, C., and Tuffin, B. (2006). A randomized quasi-Monte Carlo simulation method for Markov chains. In *Monte Carlo and quasi-Monte Carlo methods 2004*, pages 331–342. Springer Berlin Heidelberg.
- Lemieux, C. (2009). *Monte Carlo and Quasi-Monte Carlo Sampling (Springer Series in Statistics)*. Springer.
- Lemieux, C., Ormoneit, D., and Fleet, D. J. (2001). Lattice particle filters. *Proceeding UAI’01 Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 395–402.
- Matoůšek, J. (1998). On the L2 -discrepancy for anchored boxes. *Journal of Complexity*, 14:527–556.
- Neal, R. M. (2001). Annealed importance sampling. *Statist. Comput.*, 11:125–139.
- Niederreiter, H. (1992). *Random Number Generation and Quasi-Monte Carlo Methods*. CBMS-NSF Regional conference series in applied mathematics.

- Owen, A. B. (1995). Randomly permuted (t, m, s) -nets and (t, s) -sequences. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing. Lecture Notes in Statistics*, volume 106, pages 299–317. Springer, New York.
- Owen, A. B. (1997a). Monte Carlo variance of scrambled net quadrature. *SIAM Journal on Numerical Analysis*, 34(5):1884–1910.
- Owen, A. B. (1997b). Scramble net variance for integrals of smooth functions. *Ann. Stat.*, 25(4):1541–1562.
- Owen, A. B. (1998). Scrambling Sobol’ and Niederreiter-Xing points. *Journal of complexity*, 14(4):466–489.
- Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: auxiliary particle filters. *J. Am. Statist. Assoc.*, 94:590–599.
- Poyiadjis, G., Doucet, A., and Singh, S. S. (2011). Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, 98:65–80.
- Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods, 2nd ed.* Springer-Verlag, New York.
- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Ann. Math. Stat.*, 23(3):470–472.
- Sagan, H. (1994). *Space-Filling curves.* Springer-Verlag.
- Tribble, S. D. (2007). *Markov chain Monte Carlo algorithms using completely uniformly distributed driving sequences.* PhD thesis, Stanford Univ. MR2710331.
- Van der Vaart, A. W. (2007). *Asymptotic Statistics.* Cambridge series in statistical and probabilistic mathematics.
- Wächter, C. and Keller, A. (2008). Efficient simultaneous simulation of Markov chains. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*, pages 669–684. Springer.

A. Proofs

A.1. Importance sampling: Theorems 1 and 2

A.1.1. Preliminary calculation

Let $\hat{q}(d\mathbf{x}) = \mathcal{S}(\mathbf{x}^{1:N})(d\mathbf{x}) = N^{-1} \sum_{n=1}^N \delta_{\mathbf{x}^n}(d\mathbf{x})$, and, as a preliminary calculation, take $\varphi \in L_2([0, 1]^d, \lambda_d)$ and

$$\begin{aligned}
|\pi^N(\varphi) - \pi(\varphi)| &= \left| \frac{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) \varphi(\mathbf{x}^n)}{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n)} - \pi(\varphi) \right| \\
&\leq \left| \frac{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) \varphi(\mathbf{x}^n)}{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n)} - N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) \varphi(\mathbf{x}^n) \right| \\
&\quad + \left| N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) \varphi(\mathbf{x}^n) - q(w\varphi) \right| \\
&\leq \frac{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) |\varphi(\mathbf{x}^n)|}{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n)} |q(w) - \hat{q}(w)| \\
&\quad + |\hat{q}(w\varphi) - q(w\varphi)|. \tag{9}
\end{aligned}$$

We will use this inequality in the two following proofs.

A.1.2. Proof of Theorem 1

Take $\varphi = \mathbb{1}_B$ for $B \in \mathcal{B}_{[0,1]^d}$ in (9). Consider the first term above. The ratio is bounded by 1, and (since w is bounded) $|q(w) - \hat{q}(w)| \rightarrow 0$ by portmanteau lemma (Van der Vaart, 2007, Lemma 2.2). Now consider the second term.

We follow essentially the same steps as in Van der Vaart (2007, Lemma 2.2). Without loss of generality we assume that $q(d\mathbf{x})$ is a continuous probability measure (the same argument as in Van der Vaart, 2007, is used for the general case).

Let $\epsilon > 0$ and take $J \in \mathcal{B}_{[0,1]^d}$ such that $q(J^c) \leq \epsilon$. Since J is compact, $w(\cdot)$ is uniformly continuous on J . Let $\eta > 0$ be such that $\|\mathbf{x} - \mathbf{y}\| \leq \eta \implies |w(\mathbf{x}) - w(\mathbf{y})| \leq \epsilon$, $\forall (\mathbf{x}, \mathbf{y}) \in J^2$. Let $\{J_k\}_{k=1}^m$ be a split of J into a finite collection of m closed hyperrectangles with radius (at most) η . Let $g(\mathbf{x}) = \sum_{k=1}^m w(\mathbf{x}_k) \mathbb{I}_{J_k}(\mathbf{x})$ and note that $|w(\mathbf{x}) - g(\mathbf{x})| \leq 2^d \epsilon$, $\forall \mathbf{x} \in J$. Thus

$$\begin{aligned}
\left| \int_B w(\mathbf{x}) \{\hat{q}(d\mathbf{x}) - q(d\mathbf{x})\} \right| &\leq \left| \int_B \{w(\mathbf{x}) - g(\mathbf{x})\} \hat{q}(d\mathbf{x}) \right| + \left| \int_B g(\mathbf{x}) \{\hat{q}(d\mathbf{x}) - q(d\mathbf{x})\} \right| \\
&\quad + \left| \int_B \{w(\mathbf{x}) - g(\mathbf{x})\} q(d\mathbf{x}) \right|
\end{aligned}$$

where for the first term we have

$$\begin{aligned} \left| \int_B (w(\mathbf{x}) - g(\mathbf{x})) \hat{q}(\mathrm{d}\mathbf{x}) \right| &\leq \left| \int_{B \cap J} (w(\mathbf{x}) - g(\mathbf{x})) \hat{q}(\mathrm{d}\mathbf{x}) \right| + \left| \int_{B \cap J^c} w(\mathbf{x}) \hat{q}(\mathrm{d}\mathbf{x}) \right| \\ &\leq 2^d \epsilon + \|w\|_\infty \hat{q}(J^c) \\ &\leq \epsilon(2^d + 2\|w\|_\infty) \end{aligned} \quad (10)$$

as $\hat{q}(J^c)$ converges to $q(J^c)$, and thus $\hat{q}(J^c) \leq 2\epsilon$ for N large enough; and for the second term

$$\begin{aligned} \left| \int_B g(\mathbf{x}) \{\hat{q}(\mathrm{d}\mathbf{x}) - q(\mathrm{d}\mathbf{x})\} \right| &\leq \sum_{k=1}^m w(\mathbf{x}_k) \left| \int_{\bar{J}_k \cap B} \{\hat{q}(\mathrm{d}\mathbf{x}) - q(\mathrm{d}\mathbf{x})\} \right| \\ &\leq \|\hat{q}(\mathrm{d}\mathbf{x}) - q(\mathrm{d}\mathbf{x})\|_{\mathbb{E}} \sum_{k=1}^m w(\mathbf{x}_k). \end{aligned} \quad (11)$$

Finally, for the third term:

$$\begin{aligned} \left| \int_B \{w(\mathbf{x}) - g(\mathbf{x})\} q(\mathrm{d}\mathbf{x}) \right| &\leq \left| \int_{B \cap J} \{w(\mathbf{x}) - g(\mathbf{x})\} q(\mathrm{d}\mathbf{x}) \right| + \left| \int_{B \cap J^c} w(\mathbf{x}) q(\mathrm{d}\mathbf{x}) \right| \\ &\leq \epsilon(2^d + \|w\|_\infty). \end{aligned} \quad (12)$$

Putting (10)-(12) together shows that, for all $B \in \mathcal{B}_{[0,1]^d}$

$$\begin{aligned} \left| \int_B w(\mathbf{x}) \{\hat{q}(\mathrm{d}\mathbf{x}) - q(\mathrm{d}\mathbf{x})\} \right| &\leq \epsilon(2^{d+1} + 3\|w\|_\infty) + \|\hat{q}(\mathrm{d}\mathbf{x}) - q(\mathrm{d}\mathbf{x})\|_{\mathbb{E}} \sum_{k=1}^m w(\mathbf{x}_k) \\ &\leq \epsilon(2^{d+2} + 3\|w\|_\infty) \end{aligned} \quad (13)$$

for N large enough (as $\|\hat{q}(\mathrm{d}\mathbf{x}) - q(\mathrm{d}\mathbf{x})\|_{\mathbb{E}} \rightarrow 0$) which concludes the proof of Theorem 1.

A.1.3. Proof of Theorem 2

We prove first L_1 convergence (first part of Theorem 2). We start again from (9), but for any $\varphi \in L_2([0,1]^d, \lambda_d)$. For the second term, by Jensen's inequality

$$\mathbb{E} |\hat{q}(w\varphi) - q(w\varphi)| \leq [\mathrm{Var} \{\hat{q}(w\varphi)\}]^{1/2} = \mathcal{O}(r(N)^{1/2})$$

by assumption. For the first term, using Cauchy-Schwartz, $\mathbb{E}(|CD|) \leq \{\mathbb{E}(C^2)\mathbb{E}(D^2)\}^{1/2}$ with

$$C = \frac{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) |\varphi(\mathbf{x}^n)|}{N^{-1} \sum_{n=1}^N w(\mathbf{x}^n)}, \quad D = q(w) - \hat{q}(w),$$

we have $\{\mathbb{E}(D^2)\}^{1/2} = \mathcal{O}(r(N)^{1/2})$, and what remains to prove is that $\mathbb{E}(C^2) = \mathcal{O}(1)$.

From (13), and under Assumption 2, one sees that there exists N_ϵ such that with probability one $N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) \geq 1/2$ as soon as $N \geq N_\epsilon$. Under Assumption 1, a

bound similar to (13) is easily obtained by replacing $\mathbf{x}^{1:N}$ with $\mathbf{u}^{1:N}$ and observing that $w \circ F_q^{-1}$ is continuous and bounded. Thus, for N large enough

$$\mathbb{E}(C^2) \leq 4\mathbb{E} \left\{ \left[N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) |\varphi(\mathbf{x}^n)| \right]^2 \right\} \leq \mathcal{O}(r(N)) + \pi(|\varphi|)^2 = \mathcal{O}(1).$$

We now prove L_2 convergence (second part of Theorem 2):

$$\text{Var} \{ \pi^N(\varphi) \} \leq \left[\text{Var} \{ \pi^N(\varphi) - \hat{q}(w\varphi) \}^{1/2} + \text{Var} \{ \hat{q}(w\varphi) \}^{1/2} \right]^2,$$

with $\text{Var} \{ \hat{q}(w\varphi) \} = \mathcal{O}(r(N))$ by assumption, and for the first term:

$$\begin{aligned} \mathbb{E} \left[\{ \pi^N(\varphi) - \hat{q}(w\varphi) \}^2 \right] &= \mathbb{E} \left[\left\{ \sum_{n=1}^N \{ W^n - N^{-1} w(\mathbf{x}^n) \} \varphi(\mathbf{x}^n) \right\}^2 \right] \\ &= \mathbb{E} \left[\left\{ 1 - N^{-1} \sum_{n=1}^N w(\mathbf{x}^n) \right\}^2 \left\{ \sum_{n=1}^N W^n \varphi(\mathbf{x}^n) \right\}^2 \right] \\ &= \mathbb{E} \left[\frac{\{1 - \hat{q}(w)\}^2 \hat{q}(w\varphi)^2}{\hat{q}(w)^2} \right] \\ &\leq 4\mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \hat{q}(w\varphi)^2 \right] \end{aligned}$$

for N large enough, using the same argument as above (as $\hat{q}(w) \rightarrow 1$). Then

$$\begin{aligned} \mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \hat{q}(w\varphi)^2 \right] &\leq \mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \{ \hat{q}(w\varphi) - \pi(\varphi) \}^2 \right] - \pi(\varphi)^2 \mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \right] \\ &\quad + 2|\pi(\varphi)| \mathbb{E} \left[|\hat{q}(w\varphi)| \{1 - \hat{q}(w)\}^2 \right] \end{aligned}$$

where for the second term, $\mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \right] = \text{Var} [\hat{q}(w)] = \mathcal{O}(r(N))$, for the first term

$$\begin{aligned} \mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \{ \hat{q}(w\varphi) - \pi(\varphi) \}^2 \right] &= \mathbb{E} \left[\{1 - \hat{q}(w)\}^2 \{ \hat{q}(w\varphi) - q(w\varphi) \}^2 \right] \\ &\leq (1 + \|w\|_\infty)^2 \text{Var} [\hat{q}(w\varphi)] \\ &= \mathcal{O}(r(N)) \end{aligned}$$

and finally for the third term

$$\mathbb{E} \left[|\hat{q}(w\varphi)| \{1 - \hat{q}(w)\}^2 \right] \leq \mathbb{E} \left[|\hat{q}(w\varphi) - q(w\varphi)| \{1 - \hat{q}(w)\}^2 \right] + |q(w\varphi)| \text{Var} [\hat{q}(w)]$$

with

$$\begin{aligned} \mathbb{E} \left[|\hat{q}(w\varphi) - q(w\varphi)| \{1 - \hat{q}(w)\}^2 \right] &\leq (1 + \|w\|_\infty) \text{Var} [\hat{q}(w\varphi)]^{1/2} \text{Var} [\hat{q}(w)]^{1/2} \\ &= \mathcal{O}(r(N)) \end{aligned}$$

which concludes the proof.

For subsequent uses (see the proof of Theorem 6), we note that these computations imply, for N large enough,

$$\text{Var}\{\pi^N(\varphi)\} \leq \left\{ 2(1 + \|w\|_\infty)\text{Var}[\hat{q}(w\varphi)]^{1/2} + (1 - 2|\pi(\varphi)|)\text{Var}[\hat{q}(w)]^{1/2} \right\}^2 \quad (14)$$

$$|\pi^N(\varphi) - \pi(\varphi)| \leq [\text{Var}\{\hat{q}(w\varphi)\}]^{1/2} + 2[\text{Var}\{\hat{q}(w)\}]^{1/2} [\text{Var}\{\hat{q}(w\varphi)\} + \pi(|\varphi|)^2]^{1/2}. \quad (15)$$

A.2. Hilbert curve and discrepancy: Theorems 3 and 4

The proofs in this section rely on the properties of the Hilbert curve laid out in Section 2.3 and the corresponding notations.

A.2.1. Theorem 3

We first show that $\|\pi_h^N - \pi_h\|_{\mathbb{E}} = \sup_{0 \leq a < b \leq 1} |\pi_h^N([a, b]) - \pi_h([a, b])|$. Because π_h is a continuous probability measure on $[0, 1)$, the result is obvious if π_h^N is continuous as well. Let $0 \leq a < b < 1$ be such that b is a discontinuity point of $F_{\pi_h^N}$ and let $\delta > 0$ be small enough so that $\pi_h^N([a, b]) = \pi_h^N([a, b + \delta])$ and $b + \delta \leq 1$. Then,

$$\left| |\pi_h^N([a, b]) - \pi_h([a, b])| - |\pi_h^N([a, b + \delta]) - \pi_h([a, b + \delta])| \right| \leq \pi_h([b, b + \delta]).$$

By the bi-measure property of the Hilbert curve, the set $H([b, b + \delta])$ has Lebesgue measure δ in $[0, 1)^d$ and therefore, $\pi_h([b, b + \delta]) = \pi(H[b, b + \delta]) \leq \|\pi\|_\infty \delta$ where $\|\pi\|_\infty < +\infty$ by assumption. Hence, for all $\epsilon > 0$ small enough,

$$\left| \|\pi_h^N - \pi_h\|_{\mathbb{E}} - \sup_{0 \leq a < b \leq 1} |\pi_h^N([a, b]) - \pi_h([a, b])| \right| \leq \epsilon.$$

To prove the theorem note that the above computations imply that

$$\|\pi_h^N - \pi_h\|_{\mathbb{E}} \leq 2 \sup_{b \in (0, 1)} |\pi_h^N([0, b]) - \pi_h([0, b])|.$$

To bound the right-hand side, let $I = [0, b]$, $b \in (0, 1)$, and $m \in \mathbb{N}$ (which may depend on N) and assume first that $b \geq 2^{-dm}$, so that $I_m^d(0) \subseteq I$. Take $\tilde{I} = [0, k^* 2^{-dm}]$, where $k^* \leq (2^{dm} - 1)$ is the largest integer such that $k^* 2^{-dm} \leq b$. Then

$$\begin{aligned} |\pi_h^N(I) - \pi_h(I)| &\leq \left| F_{\pi_h^N}(k^* 2^{-dm}) - F_{\pi_h}(k^* 2^{-dm}) \right| \\ &\quad + \left| \pi_h^N(I) - F_{\pi_h^N}(k^* 2^{-dm}) - \left\{ \pi_h(I) - F_{\pi_h}(k^* 2^{-dm}) \right\} \right| \\ &= |\pi^N(J) - \pi(J)| + \left| \pi_h^N((k^* 2^{-dm}, b]) - \pi_h((k^* 2^{-dm}, b]) \right| \end{aligned} \quad (16)$$

with $J = H(\tilde{I})$. Since \tilde{I} is the union of k^* intervals in \mathcal{I}_m^d , J is the union of k^* hypercubes in \mathcal{S}_m^d , and therefore (using a similar argument as above and Niederreiter, 1992, Proposition 2.4),

$$|\pi^N(J) - \pi(J)| \leq c \|\pi^N - \pi\|_{\mathbb{E}} \leq 2^{dm} r(N)$$

for a constant c and where $r(N) = \|\pi^N - \pi\|_{\mathbb{E}}$.

For the second term of (16), by the properties of the Hilbert curve,

$$\begin{aligned} \left| \pi_h^N \left((k2^{-dm}, b] \right) - \pi_h \left((k2^{-dm}, b] \right) \right| &\leq \pi_h^N \left(I_m^d(k) \right) + \pi_h \left(I_m^d(k) \right) \\ &= \pi^N \left(S_m^d(k) \right) + \pi \left(S_m^d(k) \right) \\ &\leq 2\pi \left(S_m^d(k) \right) + r(N) \\ &= \mathcal{O} \left(2^{-dm} \vee r(N) \right) \end{aligned}$$

where the last inequality comes from the fact that $\pi(\mathbf{x})$ is a bounded density.

In case $b < 2^{-dm}$, similar computations show that

$$\left| \pi_h^N(I) - \pi_h(I) \right| \leq \pi_h^N(I_m^d(0)) + \pi_h(I_m^d(0)) = \mathcal{O} \left(2^{-dm} \vee r(N) \right).$$

To conclude, we choose m so that $2^{-dm} = \mathcal{O}(r(N)^{1/2})$, which gives

$$\sup_{b \in (0,1)} \left| \pi_h^N([0, b]) - \pi_h([0, b]) \right| = \mathcal{O} \left(r(N)^{1/2} \right).$$

Finally, since replacing $[0, b]$ by $[0, a]$ changes nothing to the proof of the result above, one may conclude that $\sup_{I \in \mathcal{B}_{[0,1]}} \left| \pi_h^N(I) - \pi_h(I) \right| = \mathcal{O}(r(N)^{1/2})$.

A.2.2. Proof of Theorem 4

Preliminary computations

The proof of this result is based on Hlawka and Mück (1972, ‘‘Satz 2’’). Compared to this latter, the main technical difficulty comes from the fact that the Rosenblatt transformation $F_{\pi_h^N \otimes K_h}$ is not continuous because π_h^N is a weighted sum of Dirac measures. To control the ‘‘jumps’’ of the inverse Rosenblatt transformation $F_{\pi_h^N \otimes K_h}^{-1}$ introduced by the discontinuity of π_h^N , we first prove the following Lemma.

Lemma 9. *Consider the set-up of Theorem 4. For $n \in 1:N$, let $h_1^n = H(\mathbf{x}_1^n)$ and assume that the points $h_1^{1:N}$ are labelled so that $n < m \implies h_1^n < h_1^m$. (Note that the inequality is strict because, by Assumption 2 of Theorem 4, the points $\mathbf{x}^{1:N}$ are distinct.) Without loss of generality, assume that $h_1^1 > 0$ and let $h_1^0 = 0$. Then, as $N \rightarrow +\infty$,*

$$\max_{n \in 1:N} |h_1^n - h_1^{n-1}| \rightarrow 0.$$

To prove this Lemma, let $J_N = [h_1^{n^*-1}, h_1^{n^*}]$ where $|h_1^{n^*} - h_1^{n^*-1}| = \max_{n \in 1:N} |h_1^n - h_1^{n-1}|$. Since J_N contains at most two points, we have

$$\pi_h(J_N) \leq \pi_h^N(J_N) + r_2(N) \leq 2r_1(N) + r_2(N)$$

where $r_1(N) = \max_{n \in 1:N} W_N^n$ and $r_2(N) = \|\pi_h^N - \pi_h\|_{\mathbb{E}}$; note $r_1(N) \rightarrow 0$ by Assumption 2 of Theorem 4 while $r_2(N) \rightarrow 0$ by Assumption 3 of Theorem 4 and by Theorem 3. Therefore, $\pi_h(J_N) \rightarrow 0$ as $N \rightarrow +\infty$.

Assume now that $\max_{n \in 1:N} |h_1^n - h_1^{n-1}| \not\rightarrow 0$. Then, this means that there exists a $\epsilon \in (0, 1)$ such that, for all $N > 1$ there exists a $N^* \geq N$ for which $\lambda_1(J_{N^*}) \geq \epsilon$. Assume first that $J_{N^*} \subset [0, 1 - \frac{\epsilon}{2}]$. In that case, we have $\pi_h(J_{N^*}) \geq c_\epsilon$ for a constant $c_\epsilon > 0$. Indeed, by the continuity of the Hilbert curve, the set $H([0, 1 - \frac{\epsilon}{2}])$ is compact and therefore, $\forall \mathbf{x} \in H([0, 1 - \frac{\epsilon}{2}])$, $\pi(\mathbf{x}) \geq \underline{\pi}^{(\epsilon)}$ for a constant $\underline{\pi}^{(\epsilon)} > 0$ because the density $\pi(\mathbf{x})$ is continuous and strictly positive. Therefore, if $J_{N^*} \subset [0, 1 - \frac{\epsilon}{2}]$, we have

$$\pi_h(J_{N^*}) = \pi(H(J_{N^*})) \geq \underline{\pi}^{(\epsilon)} \lambda_d(H(J_{N^*})) = \underline{\pi}^{(\epsilon)} \lambda_1(J_{N^*}) \geq \epsilon \underline{\pi}^{(\epsilon)}$$

where the second equality uses the bi-measure property of the Hilbert curve.

Assume now that $J_{N^*} \not\subset [0, 1 - \frac{\epsilon}{2}]$. Write $J_{N^*} = [a_{N^*}, b_{N^*}]$ and note that, since $\lambda_1(J_{N^*}) \geq \epsilon$, we have $a_{N^*}^* < 1 - \epsilon$ and therefore

$$\pi_h(J_{N^*}) = \pi_h\left(\left[a_{N^*}, 1 - \frac{\epsilon}{2}\right]\right) + \pi_h\left(\left[1 - \frac{\epsilon}{2}, b_{N^*}\right]\right) \geq \left(1 - \frac{\epsilon}{2} - a_{N^*}\right) \underline{\pi}^{(\epsilon)} \geq \frac{\epsilon}{2} \underline{\pi}^{(\epsilon)}.$$

Thus, this shows that if $\max_{n \in 1:N} |h_1^n - h_1^{n-1}| \not\rightarrow 0$, then there exists a $\epsilon \in (0, 1)$ such that $\limsup_{N \rightarrow +\infty} \pi_h(J_N) \geq (\epsilon \underline{\pi}^{(\epsilon)})/2 > 0$. This contradicts the fact that $\pi_h(J_N) \rightarrow 0$ as $N \rightarrow +\infty$ and the proof is complete.

Proof of Theorem 4

We use the shorthand $\alpha_N(B) = \mathcal{S}(\mathbf{u}^{1:N})(B)$ for any set $B \subset [0, 1]^{1+d_2}$. One has

$$\|\mathcal{S}(P_h^N) - \pi_h^N \otimes K_h\|_{\mathbb{E}} = \sup_{B \in \mathcal{B}_{[0,1]^{1+d_2}}^N} |\alpha_N(E^N(B)) - \lambda_{1+d_2}(E^N(B))|$$

where

$$\mathcal{B}_{[0,1]^{1+d_2}}^N = \left\{ B = [\mathbf{a}, \mathbf{b}] \in \mathcal{B}_{[0,1]^{1+d_2}} : \min_{n \in 1:N} h(\mathbf{x}_1^n) \leq F_{\pi_h^N}(b_1) \leq \max_{n \in 1:N} h(\mathbf{x}_1^n) \right\},$$

and where, for an arbitrary set $\tilde{B} = [a_1, b_1] \times [\mathbf{a}', \mathbf{b}']$ with $0 \leq a_1 \leq b_1 < 1$ and with $0 \leq a'_i \leq b'_i < 1$ for all $i \in 1 : d_2$, we use the shorthand $E^N(\tilde{B})$ for the set

$$\left\{ (u_1, \mathbf{u}_2) \in [0, 1]^{1+d_2} : F_{\pi_h^N}(a_1) \leq u_1 \leq F_{\pi_h^N}(b_1), \mathbf{u}_2 \in F_{K_h}\left(F_{\pi_h^N}^{-1}(u_1), [\mathbf{a}', \mathbf{b}']\right) \right\}.$$

Let \mathcal{P} be a partition of $[0, 1]^{1+d_2}$ in $L^{d_1+d_2}$ congruent hyperrectangles W of size $L^{-d_1} \times L^{-1} \times \dots \times L^{-1}$ where $L \geq 1$ is an arbitrary integer. Let $B = [a_1, b_1] \times [\mathbf{a}', \mathbf{b}'] \in \mathcal{B}_{[0,1]^{1+d_2}}^N$, \mathcal{U}_1 the set of the elements of \mathcal{P} that are strictly in $E^N(B)$, \mathcal{U}_2 the set of elements $W \in \mathcal{P}$ such that $W \cap \partial(E^N(B)) \neq \emptyset$, $U_1 = \cup \mathcal{U}_1$, $U_2 = \cup \mathcal{U}_2$, and $U'_1 = E^N(B) \setminus U_1$ so that

$$\alpha_N(E^N(B)) - \lambda_{1+d_2}(E^N(B)) = \alpha_N(U_1) - \lambda_{1+d_2}(U_1) + \alpha_N(U'_1) - \lambda_{1+d_2}(U'_1).$$

To bound $\alpha_N(U'_1) - \lambda_{1+d_2}(U'_1)$, note that we can cover U'_1 with sets in \mathcal{U}_2 , hence

$$\alpha_N(U'_1) - \lambda_{1+d_2}(U'_1) \leq \alpha_N(U_2), \quad \text{and} \quad \alpha_N(U'_1) - \lambda_{1+d_2}(U'_1) \geq -\lambda_{1+d_2}(U_2)$$

so that, by the definition of $D(\mathbf{u}^{1:N})$,

$$\begin{aligned} |\alpha_N(U'_1) - \lambda_{1+d_2}(U'_1)| &\leq |\alpha_N(U_2) - \lambda_{1+d_2}(U_2)| + \lambda_{1+d_2}(U_2) \\ &\leq \#\mathcal{U}_2 \left\{ D(\mathbf{u}^{1:N}) + L^{-(d_1+d_2)} \right\}. \end{aligned}$$

We therefore have

$$\begin{aligned} |\alpha_N(E^N(B)) - \lambda_{1+d_2}(E^N(B))| &\leq |\alpha_N(U_1) - \lambda_{1+d_2}(U_1)| + \#\mathcal{U}_2 \left\{ D(\mathbf{u}^{1:N}) + L^{-(d_1+d_2)} \right\} \\ &\leq L^{d_1+d_2} D(\mathbf{u}^{1:N}) + \#\mathcal{U}_2 \left\{ D(\mathbf{u}^{1:N}) + L^{-(d_1+d_2)} \right\}. \end{aligned}$$

The rest of the proof is dedicated to bounding $\#\mathcal{U}_2$, the number of hyperrectangles in \mathcal{P} required to cover $\partial(E^N(B))$. To that effect, first note that, using the continuity of F_{K_h} and the fact that B and $E^N(B)$ are closed sets, we can easily show that $E^N(\partial(B)) \subset \partial(E^N(B))$. Let $\#\mathcal{U}_2^{(1)}$ and $\#\mathcal{U}_2^{(2)}$ be, respectively, the number of hyperrectangles in \mathcal{P} we need to cover $E^N(\partial(B))$ and to cover $P(B) := \partial(E^N(B)) \setminus E^N(\partial(B))$. Hence, $\#\mathcal{U}_2 \leq \#\mathcal{U}_2^{(1)} + \#\mathcal{U}_2^{(2)}$ and we now bound $\#\mathcal{U}_2^{(i)}$, $i \in 1 : 2$.

To bound $\#\mathcal{U}_2^{(1)}$ we first cover $\partial(B)$ with hyperrectangles belonging to a partition \mathcal{P}' of the set $[0, 1]^{1+d_2}$. We construct \mathcal{P}' as a partition of the set $[0, 1]^{1+d_2}$ into hyperrectangles W' of size $L'^{-d_1} \times L'^{-1} \times \dots \times L'^{-1}$ such that, for all points (h_1, \mathbf{x}_2) and (h'_1, \mathbf{x}'_2) in W' , we have

$$\|F_{K_h}(h_1, \mathbf{x}_2) - F_{K_h}(h'_1, \mathbf{x}'_2)\|_\infty = \|F_K(H(h_1), \mathbf{x}_2) - F_K(H(h'_1), \mathbf{x}'_2)\|_\infty \leq L^{-1} \quad (17)$$

and

$$|F_{\pi_h^N}(h_1) - F_{\pi_h^N}(h'_1)| \leq L^{-d_1}. \quad (18)$$

Let $L' = 2^m$ for an integer $m \geq 0$, so that h_1 and h'_1 are in the same interval $I_m^{d_1}(k) \in \mathcal{I}_m^{d_1}$, and $H(h_1)$ and $H(h'_1)$ belong to the same hypercube in $\mathcal{S}_m^{d_1}$. Let C_K be the Lipschitz constant of F_K , then

$$\begin{aligned} \|F_K(H(h_1), \mathbf{x}_2) - F_K(H(h'_1), \mathbf{x}'_2)\|_\infty &\leq C_K \left\{ \|\mathbf{x}_2 - \mathbf{x}'_2\|_\infty \vee \|H(h_1) - H(h'_1)\|_\infty \right\} \\ &\leq C_K L'^{-1} \end{aligned}$$

and Condition (17) is verified as soon as $L' \geq C_K L$. Let us now look at Condition (18). We have:

$$\begin{aligned} |F_{\pi_h^N}(h_1) - F_{\pi_h^N}(h'_1)| &\leq 2\|F_{\pi_h^N} - F_{\pi_h}\|_\infty + |F_{\pi_h}(h_1) - F_{\pi_h}(h'_1)| \\ &\leq 2r_2(N) + |F_{\pi_h}(h_1) - F_{\pi_h}(h'_1)| \end{aligned}$$

where, as in the proof of Lemma 9, $r_2(N) = \|\pi_h^N - \pi_h\|_E$. Since h_1 and h'_1 are in the same interval $I_m^{d_1}(k) \in \mathcal{I}_m^{d_1}$,

$$|F_{\pi_h}(h_1) - F_{\pi_h}(h'_1)| \leq \pi_h \left(I_m^{d_1}(k) \right) = \pi \left(S_m^{d_1}(k) \right) \leq \frac{\|\pi\|_\infty}{(L')^{d_1}}$$

as π is bounded. To obtain both (17) and (18), we can take $L' = 2^m$ to be the smallest power of 2 such that $L' \geq k_N L$ where

$$k_N = C_K + \left(\frac{\|\pi\|_\infty}{(1 - L^{d_1} 2r_2(N))} \right)^{1/d_1}$$

which implies that we assume from now on that $L^{-d_1} \geq 4r_2(N)$ for N large enough.

Let $R \in \partial B$ be a d_2 -dimensional face of B and let \mathcal{R} be the set of hyperrectangles $W' \in \mathcal{P}'$ such that $R \cap W' \neq \emptyset$. Note that $\#\mathcal{R} \leq L'^{d_1+d_2-1} \leq (2k_N L)^{d_1+d_2-1}$. For each $W' \in \mathcal{R}$, take a point $\mathbf{r}^{W'} = (r_1^{W'}, \mathbf{r}_2^{W'}) \in R \cap W'$ and define

$$\tilde{\mathbf{r}}^{W'} = (\tilde{r}_1^{W'}, \tilde{\mathbf{r}}_2^{W'}) = F_{\pi_h^N \otimes K_h}(\mathbf{r}^{W'}) \in E^N(R).$$

Let $\tilde{\mathcal{R}}$ be the collection of hyperrectangles \tilde{W} of size $4L^{-d_1} \times 2L^{-1} \times \dots \times 2L^{-1}$ and having point $\tilde{\mathbf{r}}^{W'}$, $W' \in \mathcal{R}$, as middle point.

For an arbitrary $\mathbf{u} = (u_1, \mathbf{u}_2) \in E^N(R)$, let $h_1 = a_1 \vee F_{\pi_h^N}^{-1}(u_1)$ and $\mathbf{x}_2 = F_{K_h}^{-1}(h_1, \mathbf{u}_2)$. Since $\mathbf{x} = (h_1, \mathbf{x}_2) \in R$, \mathbf{x} is in one hyperrectangle $W' \in \mathcal{R}$. Hence, using (17) and (18),

$$|u_1 - \tilde{r}_1^{W'}| \leq |F_{\pi_h^N}(h_1) - F_{\pi_h^N}(r_1^{W'})| + |u_1 - F_{\pi_h^N}(h_1)| \leq L^{-d_1} + r_1(N),$$

where, as in the proof of Lemma 9, $r_1(N) = \max_{n \in 1:N} W_N^n$, and

$$\|\mathbf{u}_2 - \tilde{\mathbf{r}}_2^{W'}\|_\infty = \|F_{K_h}(h_1, \mathbf{x}_2) - F_{K_h}(r_1^{W'}, \mathbf{r}_2^{W'})\|_\infty \leq L^{-1}.$$

Assume from now on that $L^{-d_1} \geq r_1(N) + 4r_2(N)$. Then, this shows that \mathbf{u} belongs to the hyperrectangle $\tilde{W} \in \tilde{\mathcal{R}}$ with center $\tilde{\mathbf{r}}^{W'}$ so that $E^N(R)$ is covered by at most $\#\tilde{\mathcal{R}} = \#\mathcal{R} \leq (2k_N L)^{d_1+d_2-1}$ hyperrectangles $\tilde{W} \in \tilde{\mathcal{R}}$. To go back to the initial partition of $[0, 1]^{1+d_2}$ with hyperrectangles in \mathcal{P} , remark that every hyperrectangles in $\tilde{\mathcal{R}}$ is covered by at most c^* hyperrectangles in \mathcal{P} for a constant c^* . Finally, since the set ∂B is made of the union of $2(d_2 + 1)$ d_2 -dimensional faces of B , we have

$$\#\mathcal{U}_2^{(1)} \leq c_N L^{d_1+d_2-1} \tag{19}$$

where $c_N = c^* 2(d_2 + 1)(2k_N)^{d_1+d_2-1}$.

We now consider the problem of bounding $\#\mathcal{U}_2^{(2)}$, the number of hyperrectangles in \mathcal{P} we need to cover the set $P(B) = \partial(E^N(B)) \setminus E^N(\partial(B))$. Note that $P(B)$ contains the boundaries of the set $E^N(B)$ that are due to the discontinuities of $F_{\pi_h^N \otimes K_h}$.

To that effect, we show that there exists a finite collection $\{D_m^N\}_{m=1}^k$ of sets in $\mathcal{B}_{[0,1]^{1+d_2}}^N$ such that, for any $\mathbf{u} = (u_1, \mathbf{u}_2) \in P(B)$, there exists a $m^* \in 1 : k$ and a point $\tilde{\mathbf{u}} =$

$(\tilde{u}_1, \tilde{\mathbf{u}}_2) \in E^N(\partial(D_m^N))$ which verifies $\tilde{u}_1 = u_1$ and $\|\mathbf{u}_2 - \tilde{\mathbf{u}}_2\|_\infty \leq Cr_3(N)^{1/d_1}$ for a constant C and where $r_3(N) = \max_{n \in 1:N} |h_1^n - h_1^{n-1}|$; note that $r_3(N) \rightarrow 0$ as $N \rightarrow +\infty$ by Lemma 9. Hence, by taking L small enough (i.e. such that $L^{-1} \geq 2Cr_3(N)^{1/d_1}$), we have $\#\mathcal{U}_2^{(2)} \leq \sum_{m=1}^k \#\mathcal{U}_2^{(D_m^N)}$ where $\#\mathcal{U}_2^{(D_m^N)}$ is the number of hyperrectangles in \mathcal{P} we need to cover $E^N(\partial(D_m^N))$. Then, because the bound we derived above for the number of these hyperrectangles required to cover $E^N(\partial(B))$ is uniform in $B \in \mathcal{B}_{[0,1]^{1+d_2}}^N$, one can conclude using (19) that $\mathcal{U}_2^{(2)} \leq kc_N L^{d_1+d_2-1}$.

To construct the collection $\{D_m^N\}_{m=1}^k$, let $\mathbf{u} = (u_1, \mathbf{u}_2) \in P(B)$, that is, $u_1 = F_{\pi_h^N}(h_1^{n^*})$ for a $n^* \in 1 : N$ and $\mathbf{u}_2 = F_{K_h}(h_1^{n^*}, \mathbf{x}^*)$ with $\mathbf{x}^* \in (\mathbf{a}', \mathbf{b}')$. By the definition of the boundary of a set, for any $\epsilon > 0$ there exists a $\mathbf{v} = (v_1, \mathbf{v}_2) \notin E^N(B)$ such that $\|\mathbf{u} - \mathbf{v}\|_\infty \leq \epsilon$. Let $\epsilon > 0$ and assume that the point $\mathbf{v} = (u_1 - \epsilon, \mathbf{u}_2)$ verifies this condition, that is, $\mathbf{u}_2 \notin F_{K_h}(h_1^{n^*-1}, [\mathbf{a}', \mathbf{b}'])$, $n^* > 1$. (The case $v_1 = (u_1 + \epsilon, \mathbf{u}_2)$ is treated in a similar way, just replace $n^* - 1$ by $n^* + 1$ in what follows.)

We now show that there exists a set $B^N \in \mathcal{B}_{[0,1]^{1+d_2}}^N$ and a point $\tilde{\mathbf{u}} = (u_1, \tilde{\mathbf{u}}_2) \in E^N(\partial(B^N))$ such that $\|\mathbf{u}_2 - \tilde{\mathbf{u}}_2\|_\infty \leq Cr_3(N)^{1/d_1}$ for a constant C . We consider the set $B^N = [a_1, b_1] \times [\mathbf{a}^N, \mathbf{b}^N]$ where $\mathbf{a}^N < \mathbf{b}^N \in [0, 1]^{d_2}$. In order to construct $[\mathbf{a}^N, \mathbf{b}^N]$, we write $F_i(h_1, x_{1:i-1}, x_i)$ the i -th coordinate of $F_{K_h}(h_1, \mathbf{x})$ (with the natural convention $F_i(h_1, x_{1:i-1}, x_i) = F_1(h_1, x_1)$ when $i = 1$).

Let i^* be smallest index $i \in 1 : d_2$ such that $u_{2i} \neq F_i(h_1^{n^*-1}, x_{1:i-1}, x_i)$, $\forall \mathbf{x} \in [\mathbf{a}', \mathbf{b}']$. Then, for $i \in 1 : (i^* - 1)$, set $\tilde{u}_{2i} = u_{2i}$ and $\tilde{x}_i = x_i^*$, while, for $i \in 1 : i^*$, we set $a_i^N = a_i'$ and $b_i^N = b_i'$.

To choose \tilde{u}_{2i^*} and \tilde{x}_{i^*} we proceed as follows: if $F_{i^*}(h_1^{n^*-1}, x_{1:i^*-1}^*, b_{i^*}') < u_{2i^*}$, we take $\tilde{u}_{2i^*} = F_{i^*}(h_1^{n^*}, x_{1:i^*-1}^*, b_{i^*}')$ and $\tilde{x}_{i^*} = b_{i^*}'$ so that, noting C_H the Hölder constant of H ,

$$0 \leq \tilde{u}_{2i^*} - u_{2i^*} \leq F_{i^*}(h_1^{n^*}, x_{1:i^*-1}^*, b_{i^*}') - F_{i^*}(h_1^{n^*-1}, x_{1:i^*-1}^*, b_{i^*}') \leq C_K C_H r_3(N)^{1/d_1}$$

as required; if $F_{i^*}(h_1^{n^*-1}, x_{1:i^*-1}^*, a_{i^*}') > u_{2i^*}$, we take $\tilde{u}_{2i^*} = F_{i^*}(h_1^{n^*}, x_{1:i^*-1}^*, a_{i^*}')$ and $\tilde{x}_{i^*} = a_{i^*}'$ so that

$$0 \leq u_{2i^*} - \tilde{u}_{2i^*} \leq F_{i^*}(h_1^{n^*-1}, x_{1:i^*-1}^*, a_{i^*}') - F_{i^*}(h_1^{n^*}, x_{1:i^*-1}^*, a_{i^*}') \leq C_K C_H r_3(N)^{1/d_1}$$

as required.

Then, for $i \in (i^* + 1) : d_2$, take $\tilde{u}_{2i} = u_{2i}$ and $a_i^N = 0$. Finally, to construct the right boundaries b_i^N , $i \in (i^* + 1) : d_2$, we define

$$u_{2i}^* = \max_{n \in 1:N} \left\{ \sup \{v \in F_i(h_1^n, [0, b_{1:i}'])\} \right\}, \quad i = 1, \dots, d_2.$$

Note that $u_{2i}^* \in (0, 1)$ for all $i \in 1 : d_2$. Indeed, the continuity of F_i and the fact that $[0, b_{1:i}']$ is compact imply that

$$v_i^n := \sup \{v \in F_i(h_1^n, [0, b_{1:i}'])\} \in F_i(h_1^n, [0, b_{1:i}']).$$

Then, since $b_i' \in (0, 1)$ and F_i is strictly increasing with respect to its i -th coordinate on $[0, 1]$, we indeed have $v_i^n \in (0, 1)$ for all $n \in 1 : N$.

The right boundaries b_i^N , $i \in (i^* + 1) : d_2$ are then defined recursively as follows:

$$b_i^N = \inf \{c \in [0, 1], g_i(c) \geq u_{2i}^*\}, \quad i = i^* + 1, \dots, d_2$$

where

$$g_i(c) = \min_{(h_1, x_{1:i-1}) \in [a_1, b_1] \times [a_{1:i-1}^N, b_{1:i-1}^N]} \tilde{F}_i(h_1, x_{1:i-1}, c),$$

with $\tilde{F}_i(\cdot)$ the continuous extension of $F_i(\cdot)$ on $[0, 1]^{i+1}$. (Note that such an extension exists because F_i is Lipschitz.) Because $\tilde{F}_i(h_1, x_{1:i-1}, c)$ is continuous in $(h_1, x_{1:i-1}, c)$ and $[a_1, b_1] \times [a_{1:i-1}^N, b_{1:i-1}^N] \times [0, 1]$ is compact, the function g_i is continuous on $[0, 1]$ with $g_i(0) = 0$ and $g_i(1) = 1$. Therefore, as $u_{2i}^* \in (0, 1)$, we indeed have $b_i^N \in (0, 1)$ for all $i \in (i^* + 1) : d_2$, as required.

To show that $\tilde{\mathbf{u}} = (u_1, \tilde{\mathbf{u}}_2) \in E^N(\partial(B^N))$, note that, by the construction of \mathbf{b}^N we have, for all $i \in (i^* + 1) : d_2$,

$$F_i(h_1, x_{1:i-1}, b_i^N) \geq u_{2i}^* \geq u_{2i}, \quad \forall (h_1, x_{1:i-1}) \in [a_1, b_1] \times [a_{1:i-1}^N, b_{1:i-1}^N].$$

Therefore, by the continuity of F_i , for any $(h_1, x_{1:i-1}) \in [a_1, b_1] \times [a_{1:i-1}^N, b_{1:i-1}^N]$ there exists a $x_i \leq b_i^N$ such that $F_i(h_1, x_{1:i-1}, x_i) = u_{2i}$. Hence, for $i \in (i^* + 1) : d_2$, \tilde{x}_i is selected recursively as the unique solution of $F_i(h_1^{n^*}, \tilde{x}_{1:i-1}, \tilde{x}_i) = u_{2i}$. This concludes to show that there exists a $\tilde{\mathbf{x}} \in B^N$ such that $\tilde{\mathbf{u}}_2 = F_{K_h}(h_1^{n^*}, \tilde{\mathbf{x}})$ and $\|\mathbf{u}_2 - \tilde{\mathbf{u}}_2\|_\infty \leq C_K C_H r_3(N)^{1/d_1}$. Moreover, since $\tilde{x}_{i^*} = b_{i^*}^N$, we have $\tilde{\mathbf{x}} \in \partial(B^N)$ and therefore $\tilde{\mathbf{u}} \in \partial(E^N(B^N))$.

Finally, note that the set B^N depends only on i^* , the smallest index $i \in 1 : d_2$ such that $u_{2i} \neq F_i(h_1^{n^*}, x_{1:i-1}, x_i)$, $\forall \mathbf{x} \in [\mathbf{a}', \mathbf{b}']$. Defining $D_{i^*}^N = B^N$, this shows that the collection $\{D_i^N\}_{i=1}^{d_2}$ of sets in $\mathcal{B}_{[0,1]^{d+1}}^N$ satisfies the desired properties.

Finally, we may conclude the proof as follows:

$$\|\mathcal{S}(P_h^N) - \pi_h^N \otimes K_h\|_E \leq L^{d_1+d_2} D(\mathbf{u}^{1:N}) + (d_2 + 1)c_N L^{d_1+d_2-1} \left(D(\mathbf{u}^{1:N}) + L^{-(d_1+d_2)} \right)$$

where the optimal value of L is such that $L = \mathcal{O} \left(D(\mathbf{u}^{1:N})^{-\frac{1}{1+d_1+d_2}} \right)$. Let $r(N) = r_1(N) + 2r_2(N) + (2C_K C_H)^{d_1} r_3(N)$. Then, if $r(N) D(\mathbf{u}^{1:N})^{-\frac{d_1}{1+d_1+d_2}} = \mathcal{O}(1)$, L verifies all the conditions above and we have $c_N = \mathcal{O}(1)$. Thus

$$\|\mathcal{S}(P_h^N) - \pi_h^N \otimes K_h\|_E = \mathcal{O} \left(D(\mathbf{u}^{1:N})^{\frac{1}{1+d_1+d_2}} \right).$$

Otherwise, if $r(N) D(\mathbf{u}^{1:N})^{-\frac{d_1}{1+d_1+d_2}} \rightarrow +\infty$, let $L = \mathcal{O}(r(N)^{-\frac{1}{d_1}})$. Then $c_N = \mathcal{O}(1)$ and

$$\begin{aligned} L^{d_1+d_2} D(\mathbf{u}^{1:N}) &= \mathcal{O}(r(N))^{\frac{1}{d_1} - \frac{1+d_1+d_2}{d_1}} D(\mathbf{u}^{1:N}) \\ &= \mathcal{O}(r(N)^{1/d_1}) \left(\mathcal{O}(r(N))^{-1} D(\mathbf{u}^{1:N})^{\frac{d_1}{1+d_1+d_2}} \right)^{\frac{1+d_1+d_2}{d_1}} \\ &= \mathcal{O} \left(r(N)^{1/d_1} \right). \end{aligned}$$

Therefore $\|\mathcal{S}(P_h^N) - \pi_h^N \otimes K_h\|_E = \mathcal{O}(1)$, which concludes the proof.

A.3. Consistency: proof of Theorem 5

We first prove the following Lemma:

Lemma 10. *Let $(\pi^N \otimes K)$ be a sequence of probability measures on $[0, 1]^{d_1+d_2}$. Assume that $\|\pi^N - \pi\|_E = o(1)$, $\pi \in \mathcal{P}([0, 1]^{d_1})$ and that $F_K(\mathbf{x}_1, \mathbf{x}_2)$ is Hölder continuous with its i -th component strictly increasing in x_{2i} , $i \in 1 : d_2$. Then, as $N \rightarrow +\infty$,*

$$\|\pi^N \otimes K - \pi \otimes K\|_E \rightarrow 0.$$

To prove this result, let $B_1 \times B_2 \in \mathcal{B}_{[0,1]^{d_1+d_2}}$, $B_2 = [\mathbf{a}_2, \mathbf{b}_2]$,

$$\begin{aligned} \left| \int_{B_1 \times B_2} (\pi^N \otimes K - \pi \otimes K)(d\mathbf{x}_1, d\mathbf{x}_2) \right| &= \left| \int_{B_1} K(\mathbf{x}_1, B_2)(\pi^N - \pi)(d\mathbf{x}_1) \right| \\ &= \left| \int_{B_1} \lambda_{d_2}(F_K(\mathbf{x}_1, B_2))(\pi^N - \pi)(d\mathbf{x}_1) \right|. \end{aligned}$$

The function $\mathbf{x}_1 \rightarrow \lambda_{d_2}(F_K(\mathbf{x}_1, B_2))$ is continuous and bounded and therefore we proceed as in the proof of Theorem 1. But since $\lambda_{d_2}(F_K(\mathbf{x}_1, B_2))$ depends on $(\mathbf{a}_2, \mathbf{b}_2)$ and we want to take the supremum over $\mathbf{a}_2, \mathbf{b}_2 \in (0, 1)^{d_2}$, we need to make sure that, on a compact set J , for any $\epsilon > 0$ we can find $\eta > 0$ which does not depend on $(\mathbf{a}_2, \mathbf{b}_2)$ such that, for $\mathbf{x}_1, \mathbf{x}'_1 \in J$,

$$\|\mathbf{x}_1 - \mathbf{x}'_1\|_\infty \leq \eta \implies |\lambda_{d_2}(F_K(\mathbf{x}_1, B_2)) - \lambda_{d_2}(F_K(\mathbf{x}'_1, B_2))| \leq \epsilon.$$

To see that this is true, note that $\partial(F_K(\mathbf{x}_1, B_2)) = F_K(\mathbf{x}_1, \partial B_2)$. Hence, for any point $\mathbf{c} \in \partial F_K(\mathbf{x}_1, B_2)$ there exists a $\mathbf{p} \in \partial B_2$ such that $\mathbf{c} = F_K(\mathbf{x}_1, \mathbf{p})$ and therefore, by the Hölder property of F_K , we have

$$\|\mathbf{x}_1 - \mathbf{x}'_1\|_\infty \leq \eta \implies \|\mathbf{c} - \mathbf{c}'\|_\infty \leq C_K \eta^\kappa, \quad \mathbf{c}' = F_K(\mathbf{x}'_1, \mathbf{p}) \in \partial F_K(\mathbf{x}'_1, B_2)$$

where C_K and κ are respectively the Hölder constant and the Hölder exponent of F_K . Let \tilde{F}_K be the continuous extension of F_K on $[0, 1]^{d_1+d_2}$ (which exists because F_K is Hölder continuous on $[0, 1]^{d_1+d_2}$). Let $w > 0$, $\mathbf{x} \in [0, 1]^{d_1}$ and $\mathbf{a} \leq \mathbf{b}$, $(\mathbf{a}, \mathbf{b}) \in [0, 1]^{2d_2}$. Then, define

$$A^+(w, \mathbf{x}, \mathbf{a}, \mathbf{b}) = \left\{ \mathbf{u} \in [0, 1]^{d_2} : \exists \mathbf{p} \in \partial[\mathbf{a}, \mathbf{b}] \text{ such that } \|\mathbf{u} - \tilde{F}_K(\mathbf{x}, \mathbf{p})\|_\infty \leq C_K w^\kappa \right\}$$

and, noting $\tilde{F}_i(\mathbf{x}_1, \mathbf{x}_2)$ the i -th component of $\tilde{F}_K(\mathbf{x}_1, \mathbf{x}_2)$, $i \in 1 : d_2$,

$$\begin{aligned} A^-(w, \mathbf{x}, \mathbf{a}, \mathbf{b}) &= \left\{ \mathbf{u} \in \tilde{F}_K(\mathbf{x}, [\mathbf{a}, \mathbf{b}]) : \exists \mathbf{p} \in \partial[\mathbf{a}, \mathbf{b}] \right. \\ &\quad \left. \text{such that } |u_i - \tilde{F}_i(\mathbf{x}, \mathbf{p})| \geq C_K w^\kappa, \forall i \in 1 : d_2 \right\}. \end{aligned}$$

Let $B^* = \{(\mathbf{a}, \mathbf{b}) \in [0, 1]^{2d_2} : a_i \leq b_i, i \in 1 : d_2\}$ and $f : \mathbb{R}^+ \times [0, 1]^{d_1} \times B^* \rightarrow [0, 1]$ be the mapping

$$(w, \mathbf{x}, \mathbf{a}, \mathbf{b}) \in \mathbb{R}^+ \times [0, 1]^{d_1} \times B^* \mapsto f(w, \mathbf{x}, \mathbf{a}, \mathbf{b}) = \lambda_{d_2}(A^+(w, \mathbf{x}, \mathbf{a}, \mathbf{b})) - \lambda_{d_2}(A^-(w, \mathbf{x}, \mathbf{a}, \mathbf{b})).$$

Note that for a fix w the function $f(w, \cdot)$ is continuous on $[0, 1]^d \times B^*$ (as \tilde{F}_K is continuous). Therefore, for all \mathbf{x}_1 and \mathbf{x}'_1 in J such that $\|\mathbf{x}_1 - \mathbf{x}'_1\| \leq \eta$, we have

$$|\lambda_{d_2}(F_K(\mathbf{x}_1, B_2)) - \lambda_{d_2}(F_K(\mathbf{x}'_1, B_2))| \leq f(\eta, \mathbf{x}_1, \mathbf{a}_2, \mathbf{b}_2) \leq m(\eta)$$

with

$$m(\eta) := \max_{(\mathbf{x}, \mathbf{a}, \mathbf{b}) \in J \times B^*} f(\eta, \mathbf{x}, \mathbf{a}, \mathbf{b}).$$

Because f is continuous and $J \times B^*$ is compact, $m(\eta)$ is continuous so that, for any $\epsilon > 0$, there exists a $\eta > 0$ (that depends only on $m(\cdot)$ and therefore independent of B_2) such that $m(\eta) \leq \epsilon$. This concludes the proof of the Lemma.

We now prove Theorem 5. By the result of Hlawka and Mück (1972, ‘‘Satz 2’’) and Assumption 3, $(\mathbf{x}_0^{1:N})$ is such that $\|\mathcal{S}(\mathbf{x}_0^{1:N}) - m_0\|_{\mathbb{E}} = o(1)$. In addition, the importance weight function $\mathbb{Q}_0(d\mathbf{x}_0)/m_0(d\mathbf{x}_0) = G_0(\mathbf{x}_0)/m_0(G_0)$ is continuous and bounded by Assumption 2. Therefore, $\|\hat{\mathbb{Q}}_0^N - \mathbb{Q}_0\|_{\mathbb{E}} = o(1)$ by Theorem 1.

Assume that the result is true at time $t \geq 0$ and let $P_{t+1,h}^N = (h_t^{1:N}, \mathbf{x}_{t+1}^{1:N})$ where $h_t^n = h(\mathbf{x}_t^{\sigma_t(a_t^n)})$. Then, the result is true at time $t + 1$ if

$$\|\mathcal{S}(P_{t+1,h}^N) - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}} = o(1). \quad (20)$$

To see that, let $G_{t,h}(h_{t-1}, \mathbf{x}_t) = G_t(H(h_{t-1}), \mathbf{x}_t)$ and Ψ_{t+1} be the Boltzmann-Gibbs transformation associated to $G_{t+1,h}$ (see Del Moral, 2004, Definition 2.3.3). Then, the importance weight function

$$\frac{\Psi_{t+1}(\mathbb{Q}_{t,h} \otimes m_{t+1})}{\mathbb{Q}_{t,h} \otimes m_{t+1}}(d(h_t, \mathbf{x}_{t+1})) = \frac{G_{t+1,h}(h_t, \mathbf{x}_{t+1})}{\mathbb{Q}_t \otimes m_{t+1}(G_{t+1})}$$

is continuous and bounded (by Assumption 2 and the continuity of the Hilbert curve) and therefore Theorem 1 implies that $\|\hat{\mathbb{Q}}_{t+1}^N - \mathbb{Q}_{t+1}\|_{\mathbb{E}} = o(1)$ if (20) is verified.

To show (20), note that

$$\|\mathcal{S}(P_{t+1,h}^N) - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}} \leq \|\mathcal{S}(P_{t+1,h}^N) - \bar{\mathbb{Q}}_{t+1,h}^N\|_{\mathbb{E}} + \|\bar{\mathbb{Q}}_{t+1,h}^N - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}}.$$

By the inductive hypothesis, $\|\hat{\mathbb{Q}}_t^N - \mathbb{Q}_t\|_{\mathbb{E}} = o(1)$ so that, by Theorem 3, Assumption 3, the Hölder property of the Hilbert curve and Lemma 10,

$$\|\bar{\mathbb{Q}}_{t+1,h}^N - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}} = \|\hat{\mathbb{Q}}_t^N \otimes m_{t+1,h} - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}} = o(1).$$

Finally, note that

$$W_t^n \leq \frac{\|G_t\|_{\infty}}{\mathcal{S}(P_{t,h}^N)(G_{t,h})} = o(1)$$

because $\mathcal{S}(P_{t,h}^N)(G_{t,h}) = \mathcal{O}(N^{-1})$ by the inductive hypothesis and the fact that $G_{t,h}$ is continuous and bounded (by Assumption 2 and the continuity of the Hilbert curve). Together with the inductive hypothesis and Assumptions 1, 3-4, this implies that all the assumptions of Theorem 4 are verified and therefore $\|\mathcal{S}(P_{t+1,h}^N) - \bar{\mathbb{Q}}_{t+1,h}^N\|_{\mathbb{E}} = o(1)$ as required.

A.4. Stochastic bounds

A.4.1. Setup of the proof of Theorem 6

The result is proved by induction. By Assumption 2 of Theorem 5, the weight function $\mathbb{Q}_0(d\mathbf{x}_0)/m_0(d\mathbf{x}_0) = G_0(\mathbf{x}_0)/m_0(G_0)$ is continuous and bounded. Therefore, the continuity of $F_{m_0}^{-1}$, the assumptions on $(\mathbf{u}_0^{1:N})$ (Assumptions 1 and 2) and Theorem 2 give the result at time $t = 0$.

Assume that the result is true at time $t \geq 0$ and let $\hat{I}_{t+1}^N = \hat{\mathbb{Q}}_{t+1}^N(\varphi)$ where $\varphi : [0, 1]^d \rightarrow \mathbb{R}$ verifies the conditions of the theorem. As mentioned previously, iteration $t+1$ of SQMC is a QMC importance sampling step from the proposal distribution $\overline{\mathbb{Q}}_{t+1,h}^N$ to the target $w_{t+1,h}^h(h_t, \mathbf{x}_{t+1})\overline{\mathbb{Q}}_{t+1,h}^N(d(h_t, \mathbf{x}_{t+1}))$ where

$$w_{t+1,h}^N(h_t, \mathbf{x}_{t+1}) := \frac{G_{t+1,h}(h_t, \mathbf{x}_{t+1})}{C_{t+1}^N}$$

with $C_{t+1}^N = \overline{\mathbb{Q}}_{t+1,h}^N(G_{t+1,h})$ and $G_{t+1,h}$ as in the proof of Theorem 5. To bound $\text{Var}\{\hat{I}_{t+1}^N\}$ and $\mathbb{E}|\hat{I}_{t+1}^N - \mathbb{Q}_{t+1}(\varphi)|$ we therefore naturally want to use expression (14) and (15) derived in the proof of Theorem 2. To that effect, we need to show that, for N large enough and almost surely, the assumptions given in Theorem 2 on the weight function and on the point set at hand (Assumption 2 of Theorem 2) are satisfied.

To see that the conditions on the weight function are fulfilled, note first that $w_{t+1,h}^N$ is continuous by Assumption 2 of Theorem 5 and by the continuity of the Hilbert curve. To show that $w_{t+1,h}^N$ is almost surely bounded for N large enough, first note that, by Assumption 1, it is clear from the proofs of Theorem 3 and of Theorem 5 that, for all $\epsilon > 0$ and for all $t \geq 0$, there exists a $N_{\epsilon,t}^*$ such that, almost surely,

$$\|\hat{\mathbb{Q}}_{t,h}^N - \mathbb{Q}_{t,h}\|_{\mathbb{E}} \leq \epsilon, \quad \forall N \geq N_{\epsilon,t}^*.$$

In addition, under the assumptions of the theorem, $(C_{t+1}^N)^{-1}$ is almost surely bounded above and below away from 0, for N large enough. Indeed, by Lemma 10 (and using the Hölder property of the Hilbert curve), $\|\overline{\mathbb{Q}}_{t+1,h}^N - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}} = o(1)$ and, in particular, under the conditions of the theorem, for any $\delta > 0$, we have, almost surely,

$$\|\overline{\mathbb{Q}}_{t+1,h}^N - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\| \leq \delta \tag{21}$$

for N large enough (see the proof of Lemma 10 and the proof of Theorem 1). Writing $C_{t+1} = \mathbb{Q}_{t,h} \otimes m_{t+1,h}(G_{t+1})$, this observation, together with the fact that

$$|C_{t+1}^N - C_{t+1}| = |\overline{\mathbb{Q}}_{t+1,h}^N(G_{t+1,h}) - \mathbb{Q}_{t,h} \otimes m_{t+1,h}(G_{t+1})|$$

where $G_{t+1,h}$ is continuous and bounded (by Assumption 2 of Theorem 5 and the continuity of the Hilbert curve), implies that, almost surely, $C_{t+1} + \delta \geq C_{t+1}^N \geq C_{t+1} - \delta := c_\delta > 0$ for N large enough (computations as in the proof of Theorem 1). Hence, almost surely, $\|w_{t+1,h}^N\|_\infty \leq c_\delta^{-1} \|G_{t+1}\|_\infty$, for N large enough.

Finally, to show that the point set $P_{t+1,h}^N$ (defined as in the proof of Theorem 5) verifies Assumption 2 of Theorem 2, note that, from Theorem 5 and under the assumptions of the theorem, for any $\epsilon > 0$ there exists a N_ϵ such that, almost surely, $\|\mathcal{S}(P_{t+1,h}^N) - \mathbb{Q}_{t,h} \otimes m_{t+1,h}\|_{\mathbb{E}} \leq \epsilon$ for all $N \geq N_\epsilon$. Together with (21), this shows that, as required, for any $\epsilon > 0$ we have, almost surely and for N large enough, $\|\mathcal{S}(P_{t+1,h}^N) - \overline{\mathbb{Q}}_{t+1,h}^N\|_{\mathbb{E}} \leq \epsilon$.

A.4.2. Proof of Theorem 6: L_2 -convergence

Using expression (14) given in the proof of Theorem 2, we have for N large enough

$$\begin{aligned} \text{Var}\{\hat{I}_{t+1}^N\} &\leq \left[2(1 + c_\delta^{-1}\|G_{t+1}\|_\infty)\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(w_{t+1,h}^N)\}^{1/2} + \right. \\ &\quad \left. \{1 - 2\mathbb{E}[\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N)]\} \text{Var}\{\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N)\}^{1/2}\right]^2. \end{aligned} \quad (22)$$

We first bound $\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(w_{t+1,h}^N)\}$. Let \mathcal{F}_t^N be the σ -algebra generated by the point set $(h_{1:t-1}^{1:N}, \mathbf{x}_{1:t}^{1:N})$. Then, by Assumption 2,

$$\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(w_{t+1,h}^N)|\mathcal{F}_t^N\} \leq C^* r(N) \sigma_N^2$$

with C^* as in the statement of the theorem and $\sigma_N^2 \leq \|w_{t+1,h}^N\|_\infty \leq c_\delta^{-1}\|G_{t+1}\|_\infty$ almost surely and for N large enough. Therefore, since $\mathbb{E}[\mathcal{S}(P_{t+1,h}^N)(w_{t+1,h}^N)|\mathcal{F}_t^N] = 1$, we have

$$\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(w_{t+1,h}^N)\} = \mathcal{O}(r(N)). \quad (23)$$

Next, we need to bound $\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N)\}$. Note that

$$\overline{\mathbb{Q}}_{t+1}^N \left((C_{t+1}^N)^{-2} \varphi^2 G_{t+1}^2 \right) \leq \frac{1}{(C_{t+1}^N)^2} \|G_{t+1}\|_\infty^2 \widehat{\mathbb{Q}}_t^N(m_{t+1}(\varphi^2)),$$

where the last factor is almost surely finite for all N . Indeed, since $\varphi \in L_2(\mathcal{X}, \mathbb{Q}_{t+1})$, $m_{t+1}(\varphi^2)(\mathbf{x}_t)$ is finite for almost all $\mathbf{x}_t \in \mathcal{X}$ and the integral with respect to $\widehat{\mathbb{Q}}_t^N$ is a finite sum. Hence, for all N , $\varphi \in L_2(\mathcal{X}^2, \overline{\mathbb{Q}}_{t+1,h}^N)$ almost surely so that, by Assumption 2, we have almost surely

$$\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N)|\mathcal{F}_t^N\} \leq C^* r(N) \sigma_{N,\varphi}^2$$

where, with probability one and for N large enough, $\sigma_{N,\varphi}^2 \leq c_\delta^{-2} \|G_{t+1}\|_\infty^2 \overline{\mathbb{Q}}_{t+1}^N(\varphi^2)$. We now need to show that $\mathbb{E}[\overline{\mathbb{Q}}_{t+1}^N(\varphi^2)]$ is bounded.

In order to establish this, we prove that for all $t \geq 0$ and for N large enough, we have, $\forall f \in L_1(\mathcal{X}^2, \mathbb{Q}_t \otimes m_{t+1})$,

$$\mathbb{E}[\overline{\mathbb{Q}}_{t+1}^N(f)] \leq c_{t+1} \mathbb{Q}_t \otimes m_{t+1}(|f|) \quad (24)$$

for constant c_{t+1} .

Equation (24) is true for $t = 0$. Indeed, let $f \in L_1(\mathcal{X}^2, \mathbb{Q}_0 \otimes m_1)$ and note that, under the conditions of the theorem, almost surely and for N large enough, $\{\mathcal{S}(\mathbf{x}_0^{1:N})(G_0)\}^{-1} \leq \tilde{c}_0 < \infty$ for a constant \tilde{c}_0 . Hence, for N large enough, we have

$$\begin{aligned} \mathbb{E}[\bar{\mathbb{Q}}_1^N(f)] &= \mathbb{E} \left[\{\mathcal{S}(\mathbf{x}_0^n)(G_0)\}^{-1} \frac{1}{N} \sum_{n=1}^N G_0(\mathbf{x}_0^n) \int_{\mathcal{X}} f(\mathbf{x}_0^n, \mathbf{x}_1) m_1(\mathbf{x}_0^n, d\mathbf{x}_1) \right] \\ &\leq c_0 \mathbb{Q}_0 \otimes m_1(|f|) \end{aligned}$$

with $c_0 = \tilde{c}_0 m_0(G_0)$. Assume that (24) is true for $t \geq 0$ and note that, under the conditions of the theorem, almost surely and for N large enough, $\{\mathcal{S}(P_t^N)(G_t)\}^{-1} \leq \tilde{c}_t < \infty$ for a constant \tilde{c}_t . Then, for N large enough (with the convention $G_t(\mathbf{x}_{t-1}, \mathbf{x}_t) = G_0(\mathbf{x}_0)$ if $t = 0$),

$$\begin{aligned} \mathbb{E}[\bar{\mathbb{Q}}_{t+1}^N(f)] &= \mathbb{E} \left[\{\mathcal{S}(P_t^N)(G_t)\}^{-1} \frac{1}{N} \sum_{n=1}^N G_t(\mathbf{x}_{t-1}^{\sigma_{t-1}(a_{t-1}^n)}, \mathbf{x}_t^n) \int_{\mathcal{X}} f(\mathbf{x}_t^n, \mathbf{x}_{t+1}) m_{t+1}(\mathbf{x}_t^n, d\mathbf{x}_{t+1}) \right] \\ &\leq \tilde{c}_t \frac{1}{N} \sum_{n=1}^N \mathbb{E} \left\{ \mathbb{E} \left[\int_{\mathcal{X}} G_t(\mathbf{x}_{t-1}^{\sigma_{t-1}(a_{t-1}^n)}, \mathbf{x}_t^n) |f(\mathbf{x}_t^n, \mathbf{x}_{t+1})| m_{t+1}(\mathbf{x}_t^n, d\mathbf{x}_{t+1}) | \mathcal{F}_t^N \right] \right\} \\ &= \tilde{c}_t \mathbb{E} \left[\int_{\mathcal{X}^3} G_t(\mathbf{x}_{t-1}, \mathbf{x}_t) |f(\mathbf{x}_t, \mathbf{x}_{t+1})| \bar{\mathbb{Q}}_t^N \otimes m_{t+1}(d\mathbf{x}_{t-1:t+1}) \right] \\ &\leq \tilde{c}_t c_{t-1} \int_{\mathcal{X}^3} G_t(\mathbf{x}_{t-1}, \mathbf{x}_t) |f(\mathbf{x}_t, \mathbf{x}_{t+1})| \mathbb{Q}_{t-1} \otimes m_t \otimes m_{t+1}(d\mathbf{x}_{t-1:t+1}) \\ &= c_t \int_{\mathcal{X}^3} |f(\mathbf{x}_t, \mathbf{x}_{t+1})| \Psi_t(\mathbb{Q}_{t-1} \otimes m_t) \otimes m_{t+1}(d\mathbf{x}_{t-1:t+1}) \\ &= c_t \int_{\mathcal{X}^2} |f(\mathbf{x}_t, \mathbf{x}_{t+1})| \mathbb{Q}_t \otimes m_{t+1}(d\mathbf{x}_{t:t+1}) \\ &= c_t \mathbb{Q}_t \otimes m_{t+1}(|f|) \end{aligned}$$

with $c_t = c_{t-1} \tilde{c}_t [\mathbb{Q}_{t-1} \otimes m_t(G_t)]$, Ψ_t be the Boltzmann-Gibbs transformation associated to G_t (see Del Moral, 2004, Definition 2.3.4) and where the second inequality uses the inductive hypothesis and the fact that the mapping

$$(\mathbf{x}_{t-1}, \mathbf{x}_t) \mapsto G_t(\mathbf{x}_{t-1}, \mathbf{x}_t) m_{t+1}(|f|)(\mathbf{x}_t)$$

belongs to $L_1(\mathcal{X}^2, \mathbb{Q}_{t-1} \otimes m_t)$. This shows (24) and therefore, for N large enough, $\mathbb{E}[\sigma_{N,\varphi}^2] \leq c$ for a constant c so that $\mathbb{E}[\text{Var}\{\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N) | \mathcal{F}_t^N\}] = \mathcal{O}(r(N))$. In addition

$$\begin{aligned} \mathbb{E}[\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N) | \mathcal{F}_t^N] &= \frac{\hat{\mathbb{Q}}_t^N(m_{t+1}(\varphi G_{t+1}))}{C_{t+1}^N} \\ &= \frac{\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1}))}{C_{t+1}^N} + \frac{(\hat{\mathbb{Q}}_t^N - \mathbb{Q}_t)(m_{t+1}(\varphi G_{t+1}))}{C_{t+1}^N} \end{aligned}$$

where $\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1})) < +\infty$ because $\varphi \in L_2(\mathcal{X}, \mathbb{Q}_{t+1})$. Since

$$\frac{\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1}))}{C_{t+1}^N} = C_{t+1}^{-1} \mathbb{Q}_t(m_{t+1}(\varphi G_{t+1})) + \frac{C_{t+1} - C_{t+1}^N}{C_{t+1}^N C_{t+1}} \mathbb{Q}_t(m_{t+1}(\varphi G_{t+1})),$$

we therefore have, for N large enough,

$$\begin{aligned} \text{Var} \left\{ \frac{\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1}))}{C_{t+1}^N} \right\} &= [\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1}))]^2 \text{Var} \left\{ \frac{C_{t+1} - C_{t+1}^N}{C_{t+1}^N C_{t+1}} \right\} \\ &\leq \frac{[\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1}))]^2}{(c_\delta C_{t+1})^2} \mathbb{E} \left[\left\{ (\widehat{\mathbb{Q}}_t^N - \mathbb{Q}_t)(m_{t+1}(G_{t+1})) \right\}^2 \right]. \end{aligned}$$

Since $\|G_{t+1}\|_\infty < +\infty$, $m_{t+1}(G_{t+1})$ is bounded and the inductive hypothesis implies that the term on the right of the inequality sign is $\mathcal{O}(r(N))$. In addition, for all N large enough,

$$\text{Var} \left\{ \frac{1}{C_{t+1}^N} \left(\widehat{\mathbb{Q}}_t^N - \mathbb{Q}_t \right) (m_{t+1}(\varphi G_{t+1})) \right\} \leq c_\delta^{-2} \mathbb{E} \left[\left\{ (\widehat{\mathbb{Q}}_t^N - \mathbb{Q}_t)(m_{t+1}(\varphi G_{t+1})) \right\}^2 \right].$$

Since $[m_{t+1}(\varphi G_{t+1})(\mathbf{x}_t)]^2 \leq \|G_{t+1}\|_\infty m_{t+1}(\varphi^2 G_{t+1})(\mathbf{x}_t)$, we have

$$\mathbb{Q}_t(\{m_{t+1}(\varphi G_{t+1})\}^2) \leq \|G_{t+1}\|_\infty C_{t+1} \mathbb{Q}_t(\varphi^2) < +\infty$$

by assumption. Therefore, $m_{t+1}(\varphi G_{t+1}) \in L_2(\mathcal{X}, \mathbb{Q}_t)$ so that, by the inductive hypothesis, $\text{Var}\{\mathbb{E}[\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N) | \mathcal{F}_t^N]\} = \mathcal{O}(r(N))$. Hence,

$$\text{Var} \{ \mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N) \} = \mathcal{O}(r(N)). \quad (25)$$

The last term of (22) we need to control is

$$\mathbb{E} [\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N)] = \mathbb{E} \left[\frac{1}{C_{t+1}^N} \widehat{\mathbb{Q}}_t^N(m_{t+1}(\varphi G_{t+1})) \right].$$

Since we saw that $m_{t+1}(\varphi G_{t+1}) \in L_2(\mathcal{X}, \mathbb{Q}_t)$, we have, for N large enough,

$$\begin{aligned} \mathbb{E} [\mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N)] &= \mathbb{E} \left[\frac{1}{C_{t+1}^N} \right] \mathbb{Q}_t(m_{t+1}(\varphi G_{t+1})) \\ &\quad + \mathbb{E} \left[\frac{1}{C_{t+1}^N} (\widehat{\mathbb{Q}}_t^N - \mathbb{Q}_t)(m_{t+1}(\varphi G_{t+1})) \right] \\ &\leq c_\delta^{-1} \left[\mathbb{Q}_t(m_{t+1}(\varphi G_{t+1})) + \mathcal{O}(r(N)^{1/2}) \right] \end{aligned} \quad (26)$$

using previous computations.

Combining (22), (23), (25) and (26), one obtains $\text{Var}\{\hat{I}_{t+1}^N\} = \mathcal{O}(r(N))$.

A.4.3. Proof of Theorem 6: L_1 -convergence

Let $I_{t+1} = \mathbb{Q}_{t+1}(\varphi)$ and $I_{t+1}^N = \mathbb{Q}_{t+1}^N(\varphi)$ so that

$$\mathbb{E}[|\hat{I}_{t+1}^N - I_{t+1}|] \leq \mathbb{E}[|\hat{I}_{t+1}^N - I_{t+1}^N|] + \mathbb{E}[|I_{t+1}^N - I_{t+1}|].$$

Then, using expression (15) in the proof of Theorem 2, we have, for N large enough,

$$\begin{aligned} \mathbb{E} \left[|\hat{I}_{t+1}^N - I_{t+1}^N| \right] &\leq \text{Var} \left\{ \mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N) \right\}^{1/2} + 2 \left(\text{Var} \left\{ \mathcal{S}(P_{t+1,h}^N)(\varphi w_{t+1,h}^N) \right\} \right. \\ &\quad \left. + \left\{ \mathbb{E} \left[\mathcal{S}(P_{t+1,N})(\varphi w_{t+1,h}^N) \right]^2 \right\}^{1/2} \text{Var} \left\{ \mathcal{S}(P_{t+1,h}^N)(w_{t+1,h}^N) \right\} \right)^{1/2} \\ &= \mathcal{O}(r(N)^{1/2}) \end{aligned}$$

from above computations. In addition,

$$\begin{aligned} \mathbb{E} [|I_{t+1}^N - I_{t+1}|] &= \mathbb{E} \left[\left| \left(\frac{\hat{\mathbb{Q}}_t^N}{C_{t+1}^N} - \frac{\mathbb{Q}_t}{C_{t+1}} \right) (m_{t+1}(\varphi G_{t+1})) \right| \right] \\ &\leq \mathbb{E} \left[\frac{1}{C_{t+1}} \left| (\hat{\mathbb{Q}}_t^N - \mathbb{Q}_t) (m_{t+1}(\varphi G_{t+1})) \right| \right] \\ &\quad + \mathbb{E} \left[\frac{|C_{t+1} - C_{t+1}^N|}{C_{t+1}^N C_{t+1}} \left| \hat{\mathbb{Q}}_t^N (m_{t+1}(\varphi G_{t+1})) \right| \right]. \end{aligned}$$

By the inductive hypothesis and the above computations, the first term after the inequality sign is $\mathcal{O}(r(N)^{1/2})$. In addition, for N large enough, the second term after the inequality sign is bounded by

$$\begin{aligned} \mathbb{E} \left[\frac{|C_{t+1} - C_{t+1}^N|}{C_{t+1}^N C_{t+1}} \left| \hat{\mathbb{Q}}_t^N (m_{t+1}(\varphi G_{t+1})) \right| \right] &\leq \frac{\delta}{c_\delta C_{t+1}} \mathbb{E} \left[\left| (\hat{\mathbb{Q}}_t^N - \mathbb{Q}_t) (m_{t+1}(\varphi G_{t+1})) \right| \right] \\ &\quad + \frac{|\mathbb{Q}_t (m_{t+1}(\varphi G_{t+1}))|}{c_\delta C_{t+1}} \mathbb{E} \left[\left| (\hat{\mathbb{Q}}_t^N - \mathbb{Q}_t) (m_{t+1}(G_{t+1})) \right| \right]. \end{aligned}$$

We saw above that the first term on the right-hand side is $\mathcal{O}(r(N)^{1/2})$. In addition, $m_{t+1}(G_{t+1})$ belongs to $L_2(\mathcal{X}, \mathbb{Q}_t)$ because $\|G_{t+1}\|_\infty < +\infty$. Hence, by the inductive hypothesis, the second term after the inequality sign is also $\mathcal{O}(r(N)^{1/2})$ and the proof is complete.

A.4.4. Proof of Theorem 7

To avoid confusion between the t of the time index and the t of the (t, s) -sequence we replace the latter by \tilde{t} in what follows.

The result is true at time $t = 0$ by Theorem 2. To obtain the result for $t \geq 1$ we need to modify the steps in the proof of Theorem 6 where we do not use the inductive hypothesis. Inspection of this proof shows that we only need to establish that, for any function $\varphi \in \mathcal{C}_b([0, 1]^{1+d})$, we have

$$\mathbb{E} \left[\text{Var} \left(\mathcal{S}(P_{t+1,h}^N)(\varphi) | \mathcal{F}_t^N \right) \right] = o(N^{-1}).$$

Let $N = \lambda b^m$. Then, from the proof of Owen (1998, Theorem 1), and using the same notations as in that paper (note in particular the new meaning for symbol u), we have

$$\text{Var}(\mathcal{S}(P_{t+1,h}^N)(\varphi)|\mathcal{F}_t^N) \leq \frac{c}{N} \sum_{|u|>0} \sum_{|\kappa|>m-\bar{t}-|u|} \sigma_{N,u,\kappa}^2$$

for a constant c , where $|u|$ is the cardinal of $u \subseteq \{1, \dots, d+1\}$, κ is a vector of $|u|$ nonnegative integers k_j , $j \in u$, and $|\kappa| = \sum_{j \in u} k_j$. Note that κ depends implicitly on u . The $\sigma_{N,u,\kappa}^2$'s are such that

$$\sigma_N^2 = \overline{\mathbb{Q}}_{t+1,h}^N(\varphi^2) - \overline{\mathbb{Q}}_{t+1,h}^N(\varphi)^2 = \sum_{|u|>0} \sum_{\kappa} \sigma_{N,u,\kappa}^2,$$

with $\sigma_{N,u,\kappa}^2 = \int_{[0,1]^{1+d}} \nu_{N,u,\kappa}(\mathbf{x})^2 d\mathbf{x}$ and

$$\nu_{N,u,\kappa}(\mathbf{x}) = \sum_{\tau(u,\kappa)} \sum_{\gamma(u)} \langle \varphi \circ F_{\overline{\mathbb{Q}}_{t+1,h}^N}^{-1}, \psi_{u,\kappa,\tau,\gamma} \rangle \psi_{u,\kappa,\tau,\gamma}(\mathbf{x})$$

where $\langle f_1, f_2 \rangle = \int f_1(\mathbf{x})f_2(\mathbf{x})d\mathbf{x}$, $\psi_{u,\kappa,\tau,\gamma}$ is bounded and all the sums in the definition of $\nu_{N,u,\kappa}(\mathbf{x})$ are finite (see Owen, 1997a, for more details).

Similarly, let

$$\sigma^2 = \mathbb{Q}_{t,h} \otimes m_{t+1,h}(\varphi^2) - \mathbb{Q}_{t,h} \otimes m_{t+1,h}(\varphi)^2 = \sum_{|u|>0} \sum_{\kappa} \sigma_{u,\kappa}^2,$$

where $\sigma_{u,\kappa}^2 = \int_{[0,1]^{1+d}} \nu_{u,\kappa}(\mathbf{x})^2 d\mathbf{x}$ and with

$$\nu_{u,\kappa}(\mathbf{x}) = \sum_{\tau(u,\kappa)} \sum_{\gamma(u)} \langle \varphi \circ F_{\mathbb{Q}_{t,h} \otimes m_{t+1,h}}^{-1}, \psi_{u,\kappa,\tau,\gamma} \rangle \psi_{u,\kappa,\tau,\gamma}(\mathbf{x}).$$

We first want to establish that $|\sigma_{N,u,\kappa}^2 - \sigma_{u,\kappa}^2| = o(1)$ almost surely. Note that

$$\|\nu_{N,u,\kappa} - \nu_{u,\kappa}\|_\infty \leq c \sum_{\tau(u,\kappa)} \sum_{\gamma(u)} \left| \langle \varphi \circ F_{\overline{\mathbb{Q}}_{t+1,h}^N}^{-1}, \psi_{u,\kappa,\tau,\gamma} \rangle - \langle \varphi \circ F_{\mathbb{Q}_{t,h} \otimes m_{t+1,h}}^{-1}, \psi_{u,\kappa,\tau,\gamma} \rangle \right|$$

for a constant $c > 0$. To show that the term inside the absolute value sign is almost surely $o(1)$, assume that for all $\tilde{u} \in [0, 1)$, $|F_{\overline{\mathbb{Q}}_{t,h}^N}^{-1}(\tilde{u}) - F_{\mathbb{Q}_{t,h}}^{-1}(\tilde{u})| = o(1)$ almost surely.

Using the continuity of φ and the continuity of the Hilbert curve H , and the fact that that $F_{m_{t+1}}^{-1}(\mathbf{x}_t, \mathbf{x}_{t+1})$ is a continuous function of \mathbf{x}_t (Assumption 2), we have for any $(h_t, \mathbf{x}_{t+1}) \in [0, 1)^{d+1}$

$$\left| \varphi \circ F_{\overline{\mathbb{Q}}_{t+1,h}^N}^{-1}(h_t, \mathbf{x}_{t+1}) - \varphi \circ F_{\mathbb{Q}_{t,h} \otimes m_{t+1,h}}^{-1}(h_t, \mathbf{x}_{t+1}) \right| = o(1), \quad \text{a.s.}$$

and therefore, since φ and $\psi_{u,\kappa,\tau,\gamma}$ are bounded, we have, by the dominated convergence Theorem,

$$\left| \langle \varphi \circ F_{\overline{\mathbb{Q}}_{t+1,h}^N}^{-1}, \psi_{u,\kappa,\tau,\gamma} \rangle - \langle \varphi \circ F_{\mathbb{Q}_{t,h} \otimes m_{t+1,h}}^{-1}, \psi_{u,\kappa,\tau,\gamma} \rangle \right| \rightarrow 0, \quad \text{a.s.}$$

We now establish that, for all $\tilde{u} \in [0, 1)$, $|F_{\widehat{\mathbb{Q}}_{t,h}^N}^{-1}(\tilde{u}) - F_{\mathbb{Q}_{t,h}}^{-1}(\tilde{u})| \rightarrow 0$ almost surely. The proof of this result is inspired from Barvínek et al. (1991, Theorem 2).

First, note that because $p_t(\mathbf{x}_t) > 0$ for all $\mathbf{x}_t \in [0, 1)^d$ (Assumption 4 of Theorem 5) the function $F_{\mathbb{Q}_{t,h}}$ is continuous and strictly increasing on $[0, 1)$ (see the proof of Lemma 9). Let $\epsilon > 0$ and $\tilde{u}_1 \in [0, 1)$. Then, by the continuity of $F_{\mathbb{Q}_{t,h}}^{-1}$, there exists a $\delta_{\tilde{u}_1, \epsilon} > 0$ such that,

$$|\tilde{u}_1 - \tilde{u}| \leq \delta_{\tilde{u}_1, \epsilon}, \implies |F_{\mathbb{Q}_{t,h}}^{-1}(\tilde{u}_1) - F_{\mathbb{Q}_{t,h}}^{-1}(\tilde{u})| \leq \epsilon. \quad (27)$$

In the proof of Theorem 6 we saw that, for any $\delta_0 > 0$, there exists a N_{δ_0} such that, for all $N \geq N_{\delta_0}$,

$$\|F_{\widehat{\mathbb{Q}}_{t,h}^N} - F_{\mathbb{Q}_{t,h}}\|_{\infty} \leq \delta_0, \quad \text{a.s.} \quad (28)$$

Let $x_N = F_{\widehat{\mathbb{Q}}_{t,h}^N}^{-1}(\tilde{u}_1)$ and $u_N = F_{\mathbb{Q}_{t,h}}(x_N)$. Then, by (28),

$$|F_{\widehat{\mathbb{Q}}_{t,h}^N}(x_N) - F_{\mathbb{Q}_{t,h}}(x_N)| \leq \delta_0, \quad \forall N \geq N_{\delta_0}, \quad \text{a.s.}$$

Let $r_N(\tilde{u}_1) = F_{\widehat{\mathbb{Q}}_{t,h}^N}(F_{\widehat{\mathbb{Q}}_{t,h}^N}^{-1}(\tilde{u}_1)) - \tilde{u}_1$ so that

$$|F_{\widehat{\mathbb{Q}}_{t,h}^N}(x_N) - F_{\mathbb{Q}_{t,h}}(x_N)| = |\tilde{u}_1 + r_N(\tilde{u}_1) - u_N| \leq \delta_0, \quad \forall N \geq N_{\delta_0}, \quad \text{a.s.}$$

Now note that $|r_N(\tilde{u}_1)| \leq \frac{\|G_t\|_{\infty}}{N\mathcal{S}(P_t^N)(G_t)}$ a.s.. Then, it is easy to see that, for all $\delta' > 0$, there exists a $N_{\delta'}$ such that, a.s., $|r_N(\tilde{u}_1)| \leq \delta'$ for all $N \geq N_{\delta'}$. Let $\delta = \delta_0 + \delta'$ and set $N_{\delta} := N_{\delta_0} \vee N_{\delta'}$. Then, for $N \geq N_{\delta}$, we have almost surely $|\tilde{u}_1 - u_N| \leq \delta$. By taking δ_0 and δ' such that $\delta = \delta_{\tilde{u}_1, \epsilon}$, (27) implies that

$$|F_{\mathbb{Q}_{t,h}}^{-1}(\tilde{u}_1) - F_{\mathbb{Q}_{t,h}}^{-1}(u_N)| \leq \epsilon, \quad \forall N \geq N_{\delta_{\tilde{u}_1, \epsilon}}, \quad \text{a.s.}$$

In addition, $F_{\mathbb{Q}_{t,h}}^{-1}(u_N) = x_N = F_{\widehat{\mathbb{Q}}_{t,h}^N}^{-1}(\tilde{u}_1)$ and therefore, $\forall N \geq N_{\delta_{\tilde{u}_1, \epsilon}}$,

$$|F_{\mathbb{Q}_{t,h}}^{-1}(\tilde{u}_1) - F_{\widehat{\mathbb{Q}}_{t,h}^N}^{-1}(\tilde{u}_1)| \leq \epsilon, \quad \text{a.s.}$$

Consequently, $\|\nu_{N,u,\kappa} - \nu_{u,\kappa}\|_{\infty} = o(1)$ almost surely so that, by the dominated convergence Theorem, $\sigma_{N,u,\kappa}^2 \rightarrow \sigma_{u,\kappa}^2$ almost surely. Also, because φ is continuous and bounded, $\sigma_N^2 \rightarrow \sigma^2$ almost surely by Theorem 5 and portmanteau lemma (Van der Vaart, 2007, Lemma 2.2). To simplify the notations in what follows, let $\tilde{\sigma}_{u,l}^2 = \sum_{\kappa:|\kappa|=l} \sigma_{u,\kappa}^2$ and $\tilde{\sigma}_{N,u,l}^2 = \sum_{\kappa:|\kappa|=l} \sigma_{N,u,\kappa}^2$ (remark that l depends implicitly on u), and note that, for $m \geq \tilde{t} + d + 1$,

$$\sum_{|u|>0} \sum_{l>m-\tilde{t}-|u|} \tilde{\sigma}_{N,u,l}^2 = \sigma_N^2 - \sum_{|u|>0} \sum_{l=0}^{\infty} \mathbb{I}(l \leq m - \tilde{t} - |u|) \tilde{\sigma}_{N,u,\kappa}^2.$$

By Fubini's Theorem,

$$\mathbb{E} \left[\sum_{|u|>0} \sum_{l>m-\tilde{t}-|u|} \tilde{\sigma}_{N,u,l}^2 \right] = \mathbb{E}[\sigma_N^2] - \sum_{|u|>0} \sum_{l=0}^{\infty} \mathbb{I}(l \leq m - \tilde{t} - |u|) \mathbb{E}[\tilde{\sigma}_{N,u,l}^2]$$

where, by the dominated convergence Theorem, $\mathbb{E}[\sigma_N^2] \rightarrow \sigma^2$. In addition, since in the definition of $\tilde{\sigma}_{N,u,l}^2$ and $\tilde{\sigma}_{u,l}^2$ the notation $\sum_{\kappa:|\kappa|=l}$ denotes a finite sum, we have, for any u and l , $\tilde{\sigma}_{N,u,l}^2 \rightarrow \tilde{\sigma}_{u,l}^2$ almost surely and therefore, by the dominated convergence Theorem, $\mathbb{E}[\tilde{\sigma}_{N,u,\kappa}^2] \rightarrow \tilde{\sigma}_{u,\kappa}^2$ (because $\tilde{\sigma}_N^2$ is bounded by $\|\varphi\|_{\infty}^2$). Hence, using Fatou's lemma,

$$\begin{aligned} 0 &\leq \limsup_{m \rightarrow +\infty} \mathbb{E} \left[\sigma_N^2 - \sum_{|u|>0} \sum_{l=0}^{\infty} \mathbb{I}(l \leq m - \tilde{t} - |u|) \tilde{\sigma}_{N,u,l}^2 \right] \\ &\leq \limsup_{m \rightarrow +\infty} \mathbb{E}[\sigma_N^2] + \limsup_{m \rightarrow +\infty} \left\{ - \sum_{|u|>0} \sum_{l=0}^{\infty} \mathbb{I}(l \leq m - \tilde{t} - |u|) \mathbb{E}[\tilde{\sigma}_{N,u,l}^2] \right\} \\ &\leq \sigma^2 - \sum_{|u|>0} \sum_{l=0}^{\infty} \liminf_{m \rightarrow +\infty} \mathbb{I}(l \leq m - \tilde{t} - |u|) \mathbb{E}[\tilde{\sigma}_{N,u,l}^2] \\ &= \sigma^2 - \sum_{|u|>0} \sum_{l=0}^{\infty} \tilde{\sigma}_{u,l}^2 = 0, \end{aligned}$$

since the indicator functions converge to one.