University of Pennsylvania
ScholarlyCommons

Statistics Papers

Wharton Faculty Research

2016

Beardwood-Halton-Hammersly Theorem for Stationary Ergodic Sequences: A Counterexample

Alessandro Arlotto
University of Pennsylvania

J. Michael Steele
University of Pennsylvania

Follow this and additional works at: http://repository.upenn.edu/statistics_papers

Part of the Physical Sciences and Mathematics Commons

Beardwood-Halton-Hammersly Theorem for Stationary Ergodic Sequences: A Counterexample

Abstract

We construct a stationary ergodic process $X_1, X_2, \ldots$ such that each X_t has the uniform distribution on the unit square and the length L_n of the shortest path through the points $X_1, X_2, \ldots, X_n$ is *not* asymptotic to a constant times the square root of n. In other words, we show that the Beardwood, Halton, and Hammersley theorem does not extend from the case of independent uniformly distributed random variables to the case of stationary ergodic sequences with uniform marginal distributions.

Keywords

traveling salesman problem, Beardwood-Halton-Hammersley theorem, subadditive Euclidean functional, stationary ergodic processes, equidistribution, construction of stationary processes

Disciplines

Physical Sciences and Mathematics

BEARDWOOD-HALTON-HAMMERSLEY THEOREM
FOR STATIONARY ERGODIC SEQUENCES:
A COUNTEREXAMPLE

ALESSANDRO ARLOTTO AND J. MICHAEL STEELE

ABSTRACT. We construct a stationary ergodic process $X_1, X_2, \ldots$ such that each X_t has the uniform distribution on the unit square and the length L_n of the shortest path through the points $X_1, X_2, \ldots, X_n$ is *not* asymptotic to a constant times the square root of n. In other words, we show that the Beardwood, Halton and Hammersley theorem does not extend from the case of independent uniformly distributed random variables to the case of stationary ergodic sequences with uniform marginal distributions.

MATHEMATICS SUBJECT CLASSIFICATION (2000): Primary 60D05, 90B15; Secondary 60F15, 60G10, 60G55, 90C27.

KEY WORDS: traveling salesman problem, Beardwood-Halton-Hammersley theorem, subadditive Euclidean functional, stationary ergodic processes, equidistribution, construction of stationary processes.

1. INTRODUCTION

Given a sequence $x_1, x_2, \ldots, x_n, \ldots$ of points in the unit square $[0,1]^2$, we let $L(x_1, x_2, \ldots, x_n)$ denote the length of the shortest path through the n points $x_1, x_2, \ldots, x_n$; that is, we let

$$(1) \qquad L(x_1, x_2, \ldots, x_n) = \min_\sigma \sum_{t=1}^{n-1} |x_{\sigma(t)} - x_{\sigma(t+1)}|,$$

where the minimum is over all permutations $\sigma : \{1, 2, \ldots, n\} \to \{1, 2, \ldots, n\}$ and where $|x - y|$ denotes the usual Euclidean distance between elements x and y of $[0,1]^2$. The classical theorem of Beardwood, Halton and Hammersley (1959) tells us in the leading case that if $X_1, X_2, \ldots$ is a sequence of independent random variables with the uniform distribution on $[0,1]^2$, then there is a constant $\beta > 0$ such that

$$(2) \qquad \lim_{n \to \infty} n^{-1/2} L(X_1, X_2, \ldots, X_n) = \beta \quad \text{with probability one.}$$

The Beardwood, Halton and Hammersley (BHH) constant β has been studied extensively, and, although its exact value is still unknown, sophisticated numerical computations (Applegate et al., 2006) suggest that $\beta \approx 0.714\ldots$. The best available analytical bounds are much rougher; we only know with certainty that $0.62499 \leq \beta \leq 0.91996$. (See, e.g., Finch, 2003, pp. 497–498.)

The BHH theorem is a strong law for independent identically distributed random variables, and, for reasons detailed below, it is natural to ask if there is an *analogous ergodic theorem* where one relaxes the hypotheses to those of the classic ergodic theorem for partial sums. Our main goal is to answer this question; specifically, we construct a stationary ergodic process with uniform marginals on $[0,1]^2$ for which the length of the shortest path through n points is *not* asymptotic to a constant times $n^{1/2}$.

Theorem 1 (No Ergodic BHH). *There are constants $c_1 < c_2$ and a stationary and ergodic process $X_1, X_2, \ldots$ such that X_1 is uniformly distributed on $[0,1]^2$ and such that with probability one*

$$(3) \qquad \liminf_{n \to \infty} \frac{L(X_1, X_2, \ldots, X_n)}{n^{1/2}} \leq c_1 < c_2 \leq \limsup_{n \to \infty} \frac{L(X_1, X_2, \ldots, X_n)}{n^{1/2}}.$$

This theorem is obtained as a corollary of the next theorem where the condition of ergodicity is dropped. In this case, one can construct processes for which there is a more explicit control of the expected minimal path length.

Theorem 2 (Asymptotics of Expected Path Lengths). *There is a stationary process $X_1^*, X_2^*, \ldots$ such that X_1^* is uniformly distributed on $[0,1]^2$ and such that*

$$(4) \quad \liminf_{n \to \infty} \frac{\mathbb{E}[L(X_1^*, X_2^*, \ldots, X_n^*)]}{n^{1/2}} \leq 2^{-1/2}\beta < \beta \leq \limsup_{n \to \infty} \frac{\mathbb{E}[L(X_1^*, X_2^*, \ldots, X_n^*)]}{n^{1/2}};$$

where β is the BHH constant.

Our construction proves more broadly that there are no ergodic analogs for many of the other subadditive Euclidean functionals such as the Steiner Tree Problem (cf. Hwang et al., 1992; Yukich, 1998). We will return to these and other general considerations in Section 10 where we also describe some open problems, but first we explain more fully the motivation for Theorems 1 and 2.

BHH Theorems for Dependent Sequences: Motivation and Evolution

The traveling salesman problem (or, TSP) has a remarkably extensive literature; the monographs of Lawler et al. (1985), Gutin and Punnen (2002) and Applegate et al. (2006) note that it is among the of the most studied of all problems in combinatorial optimization. Moreover, the TSP has had a role in many practical and computational developments. For example, the TSP provided important motivation for the theory of NP-completeness, the design of polynomial time approximations, the methods of constraint generation in the theory of linear programming, and — most relevant here — the design of probabilistic algorithms.

Probabilistic algorithms are of two fundamentally different kinds. One class of algorithms uses *internal randomization*; for example, one may make a preliminary randomization of a list before running QuickSort. The other class of algorithms takes the view that the *problem input follows a probabilistic model*. Algorithms of this second kind are especially common in application areas such as vehicle routing and the layout of very large-scale integrated circuits, or VLSI.

The partitioning algorithm proposed by Karp (1977) is an algorithm of the second kind that was directly motivated by the BHH theorem. The partitioning algorithm was further analyzed in Steele (1981) and Karp and Steele (1985), and now it is well understood that for any $\epsilon > 0$ the partitioning algorithm produces a path in

time $O(n \log n)$ that has length that is asymptotically almost surely within a factor of $1 + \epsilon$ of the length of the optimal path.

Arora (1998) subsequently discovered a polynomial time algorithm that will determine a $(1+\epsilon)$-optimal path for *any set* of n points in Euclidean space, and, while the Arora algorithm is of great theoretical interest, the degree of the polynomial time bound depends on $1/\epsilon$ in a way that limits its practicality. On the other hand, Karp's partitioning algorithm is immanently practical, and it has been widely used, especially in vehicle routing problems (see e.g. Bertsimas et al., 1990; Bertsimas and van Ryzin, 1993; Pavone et al., 2007)

Still, at the core of Karp's partitioning algorithm is the assumption that the problem instance $X_1, X_2, \ldots, X_n$ can be viewed as a sequence of independent and identically distributed points in $\mathbb{R}^2$. While this assumption is feasible in some circumstances, there are certainly many more real-world problems where one would want to accommodate dependent sequences. In addition to the direct benefits, any time one establishes an analog of the BHH theorem for dependent sequences, one simultaneously robustifies a substantial body of algorithmic work.

Some extensions of the BHH theorem are easily done. For example, one can quickly confirm that the analog of the BHH theorem is valid for an infinite sequence $Y_1, Y_2, \ldots$ of exchangeable random variables with compact support. To see why this is so, one can appeal to the full theorem of Beardwood, Halton and Hammersley (1959) which asserts that for a sequence $X_1, X_2, \ldots$ of independent, identically distributed random variables with values in a compact subset of $\mathbb{R}^d$ one has the almost sure limit

$$\text{(5)} \qquad \lim_{n \to \infty} n^{-(d-1)/d} L(X_1, X_2, \ldots, X_n) = c_d \int_{\mathbb{R}^d} f(x)^{(d-1)/d} \, dx,$$

where f is the density of the absolutely continuous part of the distribution of X_1 and $0 < c_d < \infty$ is a constant that depends only on the dimension $d \geq 2$. Now, given an infinite sequences $Y_1, Y_2, \ldots$ of exchangeable random variables with values with compact support in $\mathbb{R}^d$, the de Finetti representation theorem and the limit (5) tell us that there is a random variable $0 \leq C(\omega) < \infty$ such that

$$\lim_{n \to \infty} n^{-(d-1)/d} L(Y_1, Y_2, ..., Y_n) = C \quad \text{almost surely.}$$

This particular extension of the BHH theorem is just a simple corollary of the original BHH theorem. Later, in Proposition 8 of Section 3, we give another extension of the BHH theorem to what we call *locally uniform* processes. This extension requires some delicate preparation.

There are also more wide-ranging possibilities. If one views the BHH theorem as a reflection of the evenness of the asymptotic placement of the observations, then there is a much different way to seek to generalize the BHH theorem. One can even obtain analogs of the BHH theorem for certain non-random sequences. These results are rather crude, but in some contexts, such as VLSI planning, these analogs may be more relevant than the traditional BHH theorem.

Here one considers the sequence of points $x_1, x_2, \ldots, x_n$ in $[0,1]^2$ for which one has strong control of the rectangle discrepancy that is defined by

$$D_n = D(x_1, x_2, \ldots, x_n) = \sup_{Q \in \mathcal{Q}} \left| \frac{1}{n} \sum_{t=1}^{n} \mathbb{1}_Q(x_t) - \lambda(Q) \right|,$$

where $\mathcal{Q}$ is the set of all axis-aligned rectangles $Q \subseteq [0,1]^2$, $\mathbb{1}_Q(x)$ is the indicator function that equals one if $x \in Q$, and $\lambda(Q)$ is the Lebesgue measure of Q. Inequalities of Steele (1980), more recently refined by Steinerberger (2010), then suffice to show that if one has

$$(6) \qquad\qquad nD_n = o(n^\delta) \quad \text{as } n \to \infty \quad \text{for all } \delta \in (0,1),$$

then one also has the pointwise limit

$$(7) \qquad\qquad \lim_{n\to\infty} \frac{\log L(x_1, x_2, \ldots, x_n)}{\log n} = \frac{1}{2}.$$

The leading example of such a sequence is $x_n = (n\phi_1 \bmod 1, n\phi_2 \bmod 1)$ where ϕ_1 and ϕ_2 are algebraic irrationals that are linearly independent over the rational numbers. A deep theorem of Schmidt (1964) tells us that the discrepancy of this sequence satisfies the remarkable estimate $nD_n = O((\log n)^{3+\gamma})$ for all $\gamma > 0$, and this is more than one needs for the discrepancy criterion (6).

In contrast, for an independent uniformly distributed sequence on $[0,1]^2$, one only has $nD_n = O_p(n^{1/2})$, and this is vastly weaker than the discrepancy condition (6). Moreover, even when one does have (6), the conclusion (7) falls a long distance from what the BHH theorem gives us.

The point of this review is that there are good reasons to want a theorem of BHH type for dependent sequences. Some spotty progress has been made; nevertheless, we are still far away from being able to give a simple, easily checked, criterion for a dependent compactly supported sequence to satisfy a limit that parallels the BHH theorem.

In the theory of the law of large numbers there are two natural bookends. On one side one has the Kolmogorov law of large numbers for independent identically distributed random variables with a finite mean, and on the other side one has the Birkhoff ergodic theorem. The BHH theorem initiates a strong law theory for the traveling salesman functional, and, in the fullness of time, there will surely be analogs of that theorem for various classes of dependent random variables. The perfect bookend for the theory would be a theorem that asserts that stationarity, uniformity, and ergodicity suffice. Theorems 1 and 2 show that the limit theory of the TSP cannot be bookended so nicely.

Observations on the Proofs and Methods

The main idea of the proof of Theorem 2 is that one can construct a stationary process such that, along a subsequence of $\tau_1 < \tau_2 < \cdots$ of successive (and ever larger) times, the ensemble of observations up to time τ_j will alternately either look very much like an independent uniformly distributed sample or else look like a sample that has far too many "twin cities," i.e. pairs of points that are excessively close together on a scale that depends on τ_j. To make this idea precise, we use a sequence of parameterized transformation of stationary processes where each transformation adds a new epoch with too many twin cities — *at an appropriate scale*. Finally, we show that one can build a single stationary process with infinitely many such epochs.

This limit process provides the desired example of a stationary, uniform process for which the minimal length paths have expectations that behave much differently from those of Beardwood, Halton, and Hammersley. The ergodic process required by Theorem 1 is then obtained by an extreme point argument that uses Choquet's

representation of a stationary measure with uniform marginals as a mixture of stationary ergodic measures with uniform marginals.

Finally, there are a two housekeeping observations. First, the classical BHH theorem (5) requires the distribution of the observations to have support in a compact set, so the difference between the length of the shortest *path* and the length of shortest *tour* through $X_1, X_2, \ldots, X_n$ is bounded by a constant that does not depend on n. Consequently, in limit a theorem such as Theorem 1 the distinction between tours and paths is immaterial. For specificity, all of the analyses done here are for the shortest path functional.

Also, in Section 2, we further note that it is immaterial here whether one takes the square $[0,1]^2$ with its natural metric, or if one takes the metric on $[0,1]^2$ to be the metric of the flat torus $\mathcal{T}$ that is obtained from the unit square with the natural identification of the opposing boundary edges. Theorems 1 and 2 are stated above for the traditional Euclidean metric, but, as we later make explicit, most of our analysis will be done with respect to the torus metric.

2. Two Transformations of a Stationary Sequence: the $H_{\epsilon,N}$ and $T_{\epsilon,N}$ Transformations

Our construction depends on the iteration of procedures that transform a given stationary process into another stationary process with additional properties. For integers $a \leq b$, we denote by $[a:b]$ the set $\{a, a+1, \ldots, b\}$, and, given a doubly infinite sequence of random variables $\mathcal{X} = (\ldots, X_{-1}, X_0, X_1, \ldots)$, we define the $[a:b]$ segment of $\mathcal{X}$ to be the subsequence

$$(8) \qquad \mathcal{X}[a:b] = (X_a, X_{a+1}, \ldots, X_b).$$

Now, given an integer p, we say that the process $\mathcal{X}$ is *periodic in distribution with period p* if we have

$$\mathcal{X}[a:a+k] \stackrel{d}{=} \mathcal{X}[b:b+k] \quad \text{for all } k \geq 0 \text{ and all } b \text{ such that } b = a \mod p.$$

This is certainly a weaker condition than stationarity, but, by an old randomization trick, one can transform a process that is periodic in distribution to a closely related process that is stationary. We will eventually apply this construction infinitely many times, so to fix ideas and notation, we first recall how it works in the simplest setting.

Lemma 3 (Passage from Periodicity in Distribution to Stationarity). *If the $\mathbb{R}^d$-valued doubly infinite sequence $\widehat{\mathcal{X}} = (\ldots, \widehat{X}_{-1}, \widehat{X}_0, \widehat{X}_1, \ldots)$ is periodic in distribution with period p, and, if I is chosen independently and uniformly from $[0:p-1]$, then the doubly infinite sequence $\widetilde{\mathcal{X}} = (\ldots, \widetilde{X}_{-1}, \widetilde{X}_0, \widetilde{X}_1, \ldots)$ defined by setting*

$$\widetilde{X}_t = \widehat{X}_{t+I} \quad \text{for all } t \in \mathbb{Z}$$

is a stationary process.

Proof. Fix $0 \leq j < \infty$ and take Borel subsets A_0, A_1, …, A_j of $\mathbb{R}^d$. By the definition of $\widetilde{\mathcal{X}}$ and by conditioning on I, one then has

$$\mathbb{P}(\widetilde{X}_t \in A_0, \widetilde{X}_{t+1} \in A_1, \ldots, \widetilde{X}_{t+j} \in A_j)$$

$$= \frac{1}{p} \sum_{i=0}^{p-1} \mathbb{P}(\widehat{X}_{t+i} \in A_1, \widehat{X}_{t+1+i} \in A_2, \ldots, \widehat{X}_{t+j+i} \in A_j)$$

(9)
$$= \frac{1}{p} \sum_{i=0}^{p-1} \mathbb{P}(\widehat{X}_{t+1+i} \in A_1, \widehat{X}_{t+2+i} \in A_2, \ldots, \widehat{X}_{t+1+j+i} \in A_j)$$

$$= \mathbb{P}(\widetilde{X}_{t+1} \in A_0, \widetilde{X}_{t+2} \in A_1, \ldots, \widetilde{X}_{t+1+j} \in A_j),$$

where the periodicity in distribution of $\widehat{\mathcal{X}}$ is used to obtain (9). Specifically, by periodicity in distribution, the last summand of (9) is equal to the first summand of the preceding sum. This tells us that $\mathbb{P}(\widetilde{X}_t \in A_0, \widetilde{X}_{t+1} \in A_1, \ldots, \widetilde{X}_{t+j} \in A_j)$ does not depend on t, and, since j is arbitrary, we see that $\widetilde{\mathcal{X}}$ is stationary. $\quad\square$

In the analysis of shortest paths in $[0,1]^2$, there are three useful distances that one can consider. One has (a) the traditional Euclidean distance, (b) the torus distance where one identifies opposite boundary edges and (c) the "Free-on-Boundary" distance where the cost of travel along any boundary edge is taken to be zero. For the moment, we let the length of the shortest path through the points $x_1, x_2, \ldots, x_n$ under the Euclidean, torus, and "Free-on-Boundary" distances be denoted by $L_E(\cdot)$, $L_T(\cdot)$ and $L_B(\cdot)$ respectively, and we note that one can show that

(10) $L_B(x_1, x_2, \ldots, x_n) \leq L_T(x_1, x_2, \ldots, x_n)$

$$\leq L_E(x_1, x_2, \ldots, x_n) \leq L_B(x_1, x_2, \ldots, x_n) + 4.$$

The first two inequalities are obvious, and the third is easy if one draws the right picture. Specifically, given a path of minimal length under the "Free-on-Boundary" distance, it may or may not connect with the boundary, but if it does connect with the boundary then one can always choose a path of minimal "Free-on-Boundary" length that never traverses any part of the boundary more than once. Consequently, if instead of riding for free one were to pay the full Euclidean cost of this travel, then the added cost of this boundary travel would be generously bounded by the total length of the boundary, and consequently one has the last inequality of (10). This argument of Redmond and Yukich (1994; 1996) has natural analogs in higher dimensions and for other functionals, see Yukich (1998, pp. 12–17) for further discussion and details. Incidently, one should also note that the torus model was considered earlier in the analyses of the minimum spanning tree problem by Avram and Bertsimas (1992) and Jaillet (1993).

An immediate implication of the bounds (10) is that if the asymptotic relation (2) holds for any one of the three choices for the distance on $[0,1]^2$, then it holds for all three. Here we will find it convenient to work with the torus distance, since this choice gives us a convenient additive group structure. In particular, for $X = (\xi, \xi') \in \mathcal{T}$ and $0 < \epsilon < 1$, we can define the ϵ-translation $X(\epsilon)$ of X, by setting

(11) $X(\epsilon) = ((\xi + \epsilon) \mod 1, \xi').$

That is, we get $X(\epsilon)$ by shifting X by ϵ in *just the first coordinate* and the shift is taken modulo 1.

We now consider a doubly infinite stationary process $\mathcal{X} = (\ldots, X_{-1}, X_0, X_1, \ldots)$ where each coordinate X_t takes its value in the flat torus $\mathcal{T}$. Given $N \in \mathbb{N}$, we then define blocks B_k, $k \in \mathbb{Z}$, of length $2N$ by setting

$$(12) \quad B_k = (X_{kN}, X_{kN+1}, \ldots, X_{(k+1)N-1}, X_{kN}(\epsilon), X_{kN+1}(\epsilon), \ldots, X_{(k+1)N-1}(\epsilon)),$$

where the translations $X_t(\epsilon)$ for $t \in [kN : (k+1)N - 1]$, are defined as in (11). We write the doubly infinite concatenation of these blocks as

$$(\ldots, B_{-2}, B_{-1}, B_0, B_1, B_2, \ldots),$$

and we note that this gives us a doubly infinite sequence of $\mathcal{T}$-valued random variables that we may also write as

$$\widehat{\mathcal{X}} = (\ldots, \widehat{X}_{-2}, \widehat{X}_{-1}, \widehat{X}_0, \widehat{X}_1, \widehat{X}_2, \ldots).$$

The process $\widehat{\mathcal{X}} = (\widehat{X}_t : t \in \mathbb{Z})$ is called the *hat process*, the passage from $\mathcal{X}$ to $\widehat{\mathcal{X}}$ is called a $H_{\epsilon,N}$ *transformation*, and we write

$$(13) \qquad\qquad\qquad \widehat{\mathcal{X}} = H_{\epsilon,N}(\mathcal{X}).$$

It is useful to note that the hat process $\widehat{\mathcal{X}} = (\widehat{X}_t : t \in \mathbb{Z})$ is periodic in distribution with period $2N$, so one can use Lemma 3 to construct a closely related stationary sequence $\widetilde{\mathcal{X}} = (\widetilde{X}_t : t \in \mathbb{Z})$. Specifically, we set

$$\widetilde{X}_t = \widehat{X}_{t+I} \quad \text{for all } t \in \mathbb{Z},$$

where the random index I has the uniform distribution on $[0 : 2N - 1]$ and I is independent of the sequence $\widehat{\mathcal{X}}$. The complete passage from $\mathcal{X}$ to $\widetilde{\mathcal{X}}$ is called a $T_{\epsilon,N}$ *transformation*, and it is denoted by

$$\widetilde{\mathcal{X}} = T_{\epsilon,N}(\mathcal{X}).$$

We will make repeated use of the two-step nature of this construction, and we stress that the hat process $\widehat{\mathcal{X}}$ is more than an intermediate product. The properties of the hat process $\widehat{\mathcal{X}}$ are the real guide to our constructions, and the stationary process $\widetilde{\mathcal{X}}$ is best viewed as a polished version of $\widehat{\mathcal{X}}$.

PROPERTIES OF THE $T_{\epsilon,N}$ TRANSFORMATION

The process $\widetilde{\mathcal{X}}$ that one obtains from $\mathcal{X}$ by a $T_{\epsilon,N}$ transformation retains much of the structure of $\mathcal{X}$. We begin with a simple example.

Lemma 4 (Preservation of Uniform Marginals). *If $\mathcal{X} = (X_t : t \in \mathbb{Z})$ is a doubly infinite $\mathcal{T}$-valued process such that X_t has the uniform distribution on $\mathcal{T}$ for each $t \in \mathbb{Z}$ and if*

$$\widetilde{\mathcal{X}} = T_{\epsilon,N}(\mathcal{X}) = (\widetilde{X}_t : t \in \mathbb{Z}) \quad \text{where } 0 < \epsilon < 1 \text{ and } N \in \mathbb{N},$$

then $\widetilde{X}_t$ has the uniform distribution on $\mathcal{T}$ for each $t \in \mathbb{Z}$.

Proof. Here it is immediate from the blocking and shifting steps of the $H_{\epsilon,N}$ transformation that $\widehat{X}_t$ has the uniform distribution on $\mathcal{T}$ for each $t \in \mathbb{Z}$. Also, by construction, the distribution of $\widetilde{X}_t$ is then a mixture of uniform distributions, and therefore $\widetilde{X}_t$ also has the uniform distribution on $\mathcal{T}$. $\qquad\square$

Given a doubly infinite sequence $\mathcal{X} = (\ldots, X_{-1}, X_0, X_1, \ldots)$ the *translation* of $\mathcal{X}$ is the process $\mathcal{X}(\delta)$ defined by setting

$$\mathcal{X}(\delta) = (X_t(\delta) : t \in \mathbb{Z}),$$

where, as before, we have $X_t = (\xi_t, \xi_t') \in \mathcal{T}$ and $X_t(\delta) = (\xi_t + \delta, \xi_t')$ with the addition in the first coordinate that is taken modulo one. We also say that $\mathcal{X}$ is *translation invariant* if for each $\delta > 0$ we have $\mathcal{X} \overset{d}{=} \mathcal{X}(\delta)$. Next, we check that translation invariance of a process $\mathcal{X}$ is preserved under any $T_{\epsilon,N}$ transformation.

Lemma 5 (Preservation of Translation Invariance). *If $\mathcal{X}$ is a doubly infinite, translation invariant, $\mathcal{T}$-valued process, then for each $0 < \epsilon < 1$ and $N \in \mathbb{N}$, the process $\widetilde{\mathcal{X}}$ defined by*

$$\widetilde{\mathcal{X}} = T_{\epsilon,N}(\mathcal{X})$$

is also translation invariant.

Proof. The crux of the matter is that a translation by δ and an application of the $H_{\epsilon,N}$ transformation are pointwise commutative. In symbols one has

$$(14) \qquad H_{\epsilon,N}(\mathcal{X}(\delta)) = (H_{\epsilon,N}(\mathcal{X}))(\delta),$$

where on the left we translate $\mathcal{X}$ by δ and then apply $H_{\epsilon,N}$, and on the right we apply $H_{\epsilon,N}$ and then translate by δ. A formal proof of this identity only requires one to unwrap the definition of $H_{\epsilon,N}$ and to use commutativity of addition.

Now, by our hypothesis, $\mathcal{X}$ is equal in distribution to $\mathcal{X}(\delta)$, and, since equality in distribution is preserved by a $H_{\epsilon,N}$ transformation, we also have

$$(15) \qquad H_{\epsilon,N}(\mathcal{X}) \overset{d}{=} H_{\epsilon,N}(\mathcal{X}(\delta)) = (H_{\epsilon,N}(\mathcal{X}))(\delta),$$

where in the second equality we used (14). When we shift the indices of two processes that are equal in distribution by an independent shift I, the resulting processes are again equal in distribution. Thus, when make the shift on each side of (15) that is required by the definition of the $T_{\epsilon,N}$ transformation, we have

$$\widetilde{\mathcal{X}} = T_{\epsilon,N}(\mathcal{X}) \overset{d}{=} (T_{\epsilon,N}(\mathcal{X}))(\delta) = \widetilde{\mathcal{X}}(\delta),$$

just as we needed. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

The process $\widetilde{\mathcal{X}}$ that one obtains from a doubly infinite stationary sequence $\mathcal{X}$ by a $T_{\epsilon,N}$ transformation is typically singular with respect to $\mathcal{X}$. Nevertheless, on a short segment the two processes are close in distribution. The next lemma makes this precise.

Lemma 6 (Closeness in Distribution). *Let $\mathcal{X}$ be a translation invariant doubly infinite stationary sequence with values in the flat torus $\mathcal{T}$. For each $0 < \epsilon < 1$ and $N \in \mathbb{N}$, the process $\widetilde{\mathcal{X}}$ defined by $\widetilde{\mathcal{X}} = T_{\epsilon,N}(\mathcal{X})$ satisfies*

$$(16) \qquad |\, \mathbb{P}(\widetilde{\mathcal{X}}[0 : m] \in \mathcal{A}) - \mathbb{P}(\mathcal{X}[0 : m] \in \mathcal{A}) \,| \leq \frac{m}{N},$$

for all Borel sets $\mathcal{A} \subseteq \mathcal{T}^{m+1}$ and for all $m = 0, 1, 2, \ldots$.

Proof. Recalling the two-step construction that takes one from $\mathcal{X}$ to $\widetilde{\mathcal{X}}$, we first note that we can write $\widetilde{\mathcal{X}}[0 : m]$ in terms of the hat process $\widehat{\mathcal{X}}$ given by the construction (13); specifically, we have

$$\widetilde{\mathcal{X}}[0 : m] = \widehat{\mathcal{X}}[I : I + m],$$

where the random variable I is independent of $\widehat{\mathcal{X}}$ and uniformly distributed on $\{0, 1, \ldots, 2N - 1\}$. Now we condition on the value i of I. For any i such that $[i : i + m] \subseteq [0 : N - 1]$, the definition of the hat process gives us the distributional identity

$$\text{(17)} \qquad \widehat{\mathcal{X}}[i : i + m] = \mathcal{X}[i : i + m] \stackrel{d}{=} \mathcal{X}[0 : m],$$

where in the last step we used the stationarity of $\mathcal{X}$. Similarly, for i such that $[i : i + m] \subseteq [N : 2N - 1]$, we have

$$\text{(18)} \quad \widehat{\mathcal{X}}[i : i + m] = \mathcal{X}[i - N : i - N + m](\epsilon) \stackrel{d}{=} \mathcal{X}[i - N : i - N + m] \stackrel{d}{=} \mathcal{X}[0 : m],$$

where, in the next-to-last step, we use the translation invariance of $\mathcal{X}$ and in the last step we again used the stationarity of $\mathcal{X}$.

We now consider the "good set" of indices

$$G = \{i : 0 \leq i \leq i + m < N \text{ or } N \leq i \leq i + m < 2N\},$$

where the equalities (17) and (18) hold, and we also consider the complementary "bad set" of indices $G' = [0 : 2N - 1] \backslash G$. If we condition on I and use (17) and (18), then we see that for our Borel set $\mathcal{A} \subseteq \mathcal{T}^{m+1}$ one has

$$\mathbb{P}(\widetilde{\mathcal{X}}[0 : m] \in \mathcal{A}) = \frac{1}{2N} \sum_{i \in G} \mathbb{P}(\widehat{\mathcal{X}}[i : i + m] \in \mathcal{A}) + \frac{1}{2N} \sum_{i \in G'} \mathbb{P}(\widehat{\mathcal{X}}[i : i + m] \in \mathcal{A})$$

$$= \frac{2N - 2m}{2N} \mathbb{P}(\mathcal{X}[0 : m] \in \mathcal{A}) + \frac{1}{2N} \sum_{i \in G'} \mathbb{P}(\widehat{\mathcal{X}}[i : i + m] \in \mathcal{A}),$$

which one can then write more nicely as

$$\mathbb{P}(\widetilde{\mathcal{X}}[0 : m] \in \mathcal{A}) - \mathbb{P}(\mathcal{X}[0 : m] \in \mathcal{A}) = -\frac{m}{N} \mathbb{P}(\mathcal{X}[0 : m] \in \mathcal{A})$$

$$+ \frac{1}{2N} \sum_{i \in G'} \mathbb{P}(\widehat{\mathcal{X}}[i : i + m] \in \mathcal{A}).$$

The last sum has only $|G'| = 2m$ terms, so we have the bounds

$$-\frac{m}{N} \leq \mathbb{P}(\widetilde{\mathcal{X}}[0 : m] \in \mathcal{A}) - \mathbb{P}(\mathcal{X}[0 : m] \in \mathcal{A}) \leq \frac{m}{N}$$

that complete the proof of the lemma. $\qquad\qquad\square$

3. Locally Uniform Processes and BHH in Mean

Our inductive construction requires an extension of the Beardwood, Halton and Hammersley theorem to a certain class of dependent processes that we call *locally uniform processes*. The definition of these processes requires some notation.

First, given any Borel set $A \subseteq \mathcal{T}$ and any set of indices $J \subseteq \mathbb{Z}$, we let

$$N(A, \mathcal{Y}[J]) = \sum_{t \in J} \mathbb{1}(Y_t \in A),$$

so $N(A, \mathcal{Y}[J])$ is the number of elements of $\mathcal{Y}[J] = (Y_t : t \in J)$ that fall in the Borel set A. We also say that a subset $Q \subseteq \mathcal{T}$ is a *subsquare of side length* α if it can be written as $[x, x + \alpha] \times [y, y + \alpha]$ where one makes the usual identifications of the points in the flat torus. Finally, given any $0 < \alpha \leq 1$, we let $\mathcal{Q}(\alpha)$ denote the set of all subsquares of $\mathcal{T}$ that have side length less than or equal to α.

We further let

$$(U_s(A) : 1 \leq s < \infty)$$

denote an infinite sequence of independent random variables with the uniform distribution on the Borel set $A \subseteq \mathcal{T}$; in particular, $\mathcal{S} = \{U_s(A) : 1 \leq s \leq n\}$ is a uniform random sample from A with cardinality $|\mathcal{S}| = n$. Finally, if $\mathcal{S}$ and $\mathcal{S}'$ are two random finite subsets of the Borel set $A \subseteq \mathcal{T}$, we write

$$\mathcal{S} \stackrel{\text{dpp}}{=} \mathcal{S}'$$

to indicate that $\mathcal{S}$ and $\mathcal{S}'$ are equal in distribution as point processes.

Definition 7 (Locally Uniform Processes). If $0 < \alpha \leq 1$ and $M \in \mathbb{N}$ we say that a $\mathcal{T}$-valued process $\mathcal{Y} = (Y_t : t \in \mathbb{Z})$ with uniform marginal distributions is an (α, M) *locally uniform process* provided that it satisfies the two following conditions:

(19)   VARIANCE CONDITION. There is a constant $C < \infty$ that depends only on the distribution of $\mathcal{Y}$ such that for each pair of integers $a \leq b$ and each Borel set $A \subseteq Q \in \mathcal{Q}(\alpha)$ we have

$$\text{Var}[N(A, \mathcal{Y}[a : b])] \leq C|b - a + 1|.$$

(20)   LOCAL UNIFORMITY CONDITION. For each pair of integers $a \leq b$ and each Borel set $A \subseteq Q \in \mathcal{Q}(\alpha)$ there is a random set

$$\mathcal{S} \subseteq \{Y_t : Y_t \in A \text{ and } t \in [a : b]\}$$

for which one has the cardinality bounds

$$N(A, \mathcal{Y}[a : b]) - M \leq |\mathcal{S}| \leq N(A, \mathcal{Y}[a : b]),$$

and the distributional identity

$$\mathcal{S} \stackrel{\text{dpp}}{=} \{U_s(A) : 1 \leq s \leq |\mathcal{S}|\}.$$

Here one should note that if $\alpha = 1$ and $M = 0$, then a locally uniform process is just a sequence of independent random variables with the uniform distribution on the flat torus $\mathcal{T}$. More generally, the parameter α quantifies the scale at which an (α, M) locally uniform process looks almost like a sample of independent uniformly distributed random points, and the parameter M bounds the size of an exception set that can be discarded to achieve exact uniformity.

One should also note that in a locally uniform process $\mathcal{Y} = (Y_t : t \in \mathbb{Z})$ each Y_t has the uniform distribution on $\mathcal{T}$, but $\mathcal{Y}$ is not required to be a stationary process. This will be important to us later. The first observation is that despite the possible lack of stationarity, one can still show that locally uniform processes satisfy a relaxed version of the BHH theorem.

Proposition 8 (BHH in Mean for Locally Uniform Processes). *If the $\mathcal{T}$-valued process $\mathcal{Y} = (Y_t : t \in \mathbb{Z})$ is (α, M) locally uniform for some $0 < \alpha \leq 1$ and some $M \in \mathbb{N}$, then one has*

$$\mathbb{E}[L(Y_1, Y_2, \ldots, Y_n)] \sim \beta n^{1/2} \quad \text{as } n \to \infty,$$

where β is the BHH constant in (2).

Proof. For any integer k such that $k^{-1} \leq \alpha$ we consider the natural decomposition of the flat torus into k^2 subsquares Q_i, $i = 1, 2, \ldots, k^2$, of side length k^{-1}. We then introduce the sets

$$S(Q_i, n) = \{Y_t : Y_t \in Q_i \text{ and } t \in [1 : n]\}$$

and we let $L(S(Q_i, n))$ be the length of the shortest path through the points in $S(Q_i, n)$. If we then stitch these k^2 optimal paths together by considering the subsquares Q_i, $1 \leq i \leq k^2$, in plowman's order — down one row then back the next, then our stitching cost is less than $3k$, but all we need from these considerations is that there is a universal constant $C_1 > 0$ such that one has the pointwise bound,

$$(21) \qquad L(Y_1, Y_2, \ldots, Y_n) \leq C_1 k + \sum_{i=1}^{k^2} L(S(Q_i, n)).$$

More notably, one can also show that there is a universal constant $C_0 > 0$ for which one has

$$(22) \qquad -C_0 k + \sum_{i=1}^{k^2} L(S(Q_i, n)) \leq L(Y_1, Y_2, \ldots, Y_n).$$

This bound is due to Redmond and Yukich (1994), and it may be proved by noticing that the sum of the values $L(S(Q_i, n))$ can be bounded by the length of the optimum path through $Y_1, Y_2, \ldots, Y_n$ and the sum of lengths of the boundaries of the individual squares Q_i, $1 \leq i \leq k^2$. For the details concerning (22), including analogous bounds for $[0, 1]^d$, $d \geq 2$, one can consult Yukich (1998, Chapter 3).

We now recall that

$$S(Q_i, n) = \{Y_t : Y_t \in Q_i, t \in [1 : n]\} \quad \text{and} \quad N(Q_i, \mathcal{Y}[1 : n]) = \sum_{t=1}^{n} \mathbb{1}(Y_t \in Q_i),$$

and we estimate the value of $\mathbb{E}[L(S(Q_i, n))]$ as $n \to \infty$. By the (α, M) local uniformity of $\mathcal{Y}$ and (20), there is a set $\mathcal{S}_n \subseteq S(Q_i, n)$ such that

$$(23) \qquad N(Q_i, \mathcal{Y}[1 : n]) - M \leq |\mathcal{S}_n| \leq N(Q_i, \mathcal{Y}[1 : n])$$

and

$$\mathcal{S}_n \overset{\text{dpp}}{=} \{U_s(Q_i) : 1 \leq s \leq |\mathcal{S}_n|\}.$$

From the cardinality bounds (23) and a crude path length comparison we then have

$$(24) \qquad |L(\mathcal{S}_n) - L(S(Q_i, n))| \leq 2^{3/2} M k^{-1}.$$

Now, given the pointwise bounds (21), (22) and (24), the lemma will follow once we show that

$$(25) \qquad \mathbb{E}[L(\mathcal{S}_n)] \sim \beta k^{-2} n^{1/2} \qquad \text{as } n \to \infty.$$

If we now let $\ell(j)$ denote the expected length of the minimal path through an independent uniform sample of size j in the unit square, then by scaling and by conditioning on the cardinality of $\mathcal{S}_n$ we have

$$\mathbb{E}[L(\mathcal{S}_n)] = k^{-1} \sum_{j=0}^{n} \ell(j) \mathbb{P}(|\mathcal{S}_n| = j).$$

The BHH relation (2) also tells us that we have $\ell(j) = \beta j^{1/2} + o(j^{1/2})$, so

$$(26) \qquad \mathbb{E}[L(\mathcal{S}_n)] = \beta k^{-1} \sum_{j=0}^{n} j^{1/2} \, \mathbb{P}(|\mathcal{S}_n| = j) + o(n^{1/2})$$

$$= \beta k^{-1} \mathbb{E}[|\mathcal{S}_n|^{1/2}] + o(n^{1/2}).$$

Now one just needs to estimate $\mathbb{E}[|\mathcal{S}_n|^{1/2}]$. Linearity of expectation and the fact that each Y_t is uniformly distributed on $\mathcal{T}$ then combine to give us

(27) $\qquad \mathbb{E}[N(Q_i, \mathcal{Y}[1:n])] = k^{-2}n \quad \text{and} \quad \mathbb{E}[N(Q_i, \mathcal{Y}[1:n])^{1/2}] \leq k^{-1}n^{1/2},$

where the second bound comes from Jensen's inequality. For any $0 < \theta < 1$, the Variance Condition (19) and Chebyshev's inequality also give us

$$\mathbb{P}(N(Q_i, \mathcal{Y}[1:n]) \geq \theta \, \mathbb{E}[N(Q_i, \mathcal{Y}[1:n])]) = 1 - O(1/n),$$

so we have the lower bound

(28) $\qquad \theta^{1/2}k^{-1}n^{1/2}(1 - O(1/n)) \leq \mathbb{E}[N(Q_i, \mathcal{Y}[1:n])^{1/2}].$

By the inequalities (27) and (28), we see from the arbitrariness of θ and the bound (23) that we have

$$\mathbb{E}[|\mathcal{S}_n|^{1/2}] \sim k^{-1}n^{1/2} \qquad \text{as } n \to \infty.$$

Finally, from (26) we see this gives us that $\mathbb{E}[L(\mathcal{S}_n)] \sim \beta k^{-2}n^{1/2}$ as $n \to \infty$, so by the observation preceding (25) the proof of the lemma is complete. $\qquad \square$

4. Preservation of Local Uniformity

A key feature of local uniformity is that it is preserved by a $H_{\epsilon,N}$ transformation. Before proving this fact, it is useful to introduce some notation and to make a geometric observation. First, for any Borel set $A \subseteq \mathcal{T}$ we consider two kinds of translations of A by an $\epsilon > 0$; these are given by

$$^{\epsilon}A = \{(x - \epsilon, y) : (x, y) \in A\} \quad \text{and} \quad A^{\epsilon} = \{(x + \epsilon, y) : (x, y) \in A\},$$

so the set $^{\epsilon}A$ is equal to the set A translated to the *left* in $\mathcal{T}$, and A^{ϵ} is equal to the set A shifted to the *right* in $\mathcal{T}$.

Next, we fix a Borel set $A \subseteq \mathcal{T}$ for which we have $^{\epsilon}A \cap A = \emptyset$. We then consider a sequence of independent random variables $Y_1, Y_2, \ldots, Y_n$ with the uniform distribution on $^{\epsilon}A \cup A$ and a sequence of independent random variables $Y'_1, Y'_2, \ldots, Y'_n$ with the uniform distribution on A. In this situation, we then have an elementary distributional identity of point processes,

(29) $\quad \{Y_t : Y_t \in {}^{\epsilon}A, \ t \in [1:n]\}^{\epsilon} \cup \{Y_t : Y_t \in A, \ t \in [1:n]\} \stackrel{\text{dpp}}{=} \{Y'_t : t \in [1:n]\},$

where one should note that in the first term of (29) we have used both a left shift on A and a right shift on the set of points that fall into the shifted set $^{\epsilon}A$. The identity (29) provides an essential step in the proof of the next proposition.

Proposition 9 (Local Uniformity and $H_{\epsilon,N}$ Transformations). *If $\mathcal{X}$ is an (α, M) locally uniform process with $0 < \alpha \leq 1$ and $M \in \mathbb{N}$ and if one has $0 < \epsilon < \alpha$ and $N < \infty$, then the process*

$$\widehat{\mathcal{X}} = H_{\epsilon,N}(\mathcal{X})$$

is $(\widehat{\alpha}, \widehat{M})$ locally uniform with

$$0 < \widehat{\alpha} < \min\{\epsilon, \alpha - \epsilon\} \quad \text{and} \quad \widehat{M} = M + 4N.$$

Proof. The definition of $H_{\epsilon,N}$ tells us that for each $a \leq b$ we can write $\widehat{\mathcal{X}}[a:b]$ in the form

(30) $\qquad \widehat{\mathcal{X}}[a:b] = (L, B_{\ell}, B_{\ell+1}, \ldots, B_r, R).$

where the lengths of the segments L and R are between 0 and $2N - 1$ and where each segment B_k, $k \in [\ell : r]$, is a "complete block" of $\widehat{\mathcal{X}}$ that has the form

$$(31) \quad B_k = (X_{kN}, X_{kN+1}, \ldots, X_{(k+1)N-1}, X_{kN}(\epsilon), X_{kN+1}(\epsilon), \ldots, X_{(k+1)N-1}(\epsilon)).$$

Here, in the case that $\widehat{\mathcal{X}}[a : b]$ does not contain a complete block, we simply write $\widehat{\mathcal{X}}[a : b] = (L, R)$ for any choices of L and R that satisfy the length constraints. In this case we (crudely) note that $|b - a + 1| < 4N$.

If the set of complete blocks in the decomposition (30) is not empty, then we can take

$$(32) \qquad\qquad \widehat{\mathcal{X}}[a' : b'] = (B_\ell, B_{\ell+1}, \ldots, B_r),$$

to be the maximal subsegment of $\widehat{\mathcal{X}}[a : b]$ that contains only complete blocks. In this case, we have

$$a' = \min\{2kN : 2kN \geq a\} \quad \text{and} \quad b' = \max\{2kN - 1 : 2kN - 1 \leq b\},$$

but we will not make use of these explicit formulas beyond noting that we have the bounds $0 \leq |b - a + 1| - |b' - a' + 1| \leq 4N$.

For each $k \in [\ell : r]$ we let B'_k be the first half of the complete block B_k described in (31); that is, we set

$$B'_k = (X_{kN}, X_{kN+1}, \ldots, X_{(k+1)N-1}).$$

We then concatenate these half blocks to obtain a segment of the original $\mathcal{X}$ process; specifically we obtain a segment of $\mathcal{X}$ that can be written as

$$\mathcal{X}[a'' : b''] = (B'_\ell, B'_{\ell+1}, \ldots, B'_r).$$

Here one should note that the length of the segment $\mathcal{X}[a'' : b'']$ is exactly half of the length of the segment $\widehat{\mathcal{X}}[a' : b']$. We will use the correspondence between $[a : b]$, $[a' : b']$, and $[a'' : b'']$ throughout the remainder of this proof.

To confirm the local uniformity of $\widehat{\mathcal{X}}$, we need to check the two conditions (19) and (20). To check (19), we first fix a Borel set $A \subseteq Q \in \mathcal{Q}(\widehat{\alpha})$, and we note that our assumptions $0 < \widehat{\alpha} < \min\{\epsilon, \alpha - \epsilon\}$ and $0 < \epsilon < \alpha$ give us the relations

$$(33) \qquad\qquad {}^\epsilon A \cap A = \emptyset \quad \text{and} \quad {}^\epsilon A \cup A \subseteq Q' \quad \text{for a } Q' \in \mathcal{Q}(\alpha),$$

which are illustrated in Figure 1 on page 14.

Next, we consider an integer pair $a < b$, and we assume for the moment that $\widehat{\mathcal{X}}[a : b]$ contains at least one complete block. By our decomposition (30) of $\widehat{\mathcal{X}}[a : b]$ we have $0 \leq |b - a + 1| - |b' - a' + 1| \leq 4N$, so there is a random variable $W(A, \widehat{\mathcal{X}}[a : b])$ such that $0 \leq W(A, \widehat{\mathcal{X}}[a : b]) \leq 4N$ and

$$(34) \qquad N(A, \widehat{\mathcal{X}}[a, b]) = N(A, \widehat{\mathcal{X}}[a' : b']) + W(A, \widehat{\mathcal{X}}[a : b])$$
$$= N({}^\epsilon A \cup A, \mathcal{X}[a'' : b'']) + W(A, \widehat{\mathcal{X}}[a : b]),$$

where in the second equality we just used the definition of $\widehat{\mathcal{X}}$ and the definition of the set of indices $[a'' : b'']$.

When we compute variances in (34), we get

$$(35) \qquad \mathrm{Var}[N(A, \widehat{\mathcal{X}}[a : b])] \leq 2\,\mathrm{Var}[N({}^\epsilon A \cup A, \mathcal{X}[a'' : b''])] + 2(4N)^2$$
$$\leq 2C|b'' - a'' + 1| + 32N^2 \leq C|b - a + 1| + 32N^2,$$

FIGURE 1. Local Uniformity and $H_{\epsilon,N}$ Transformations
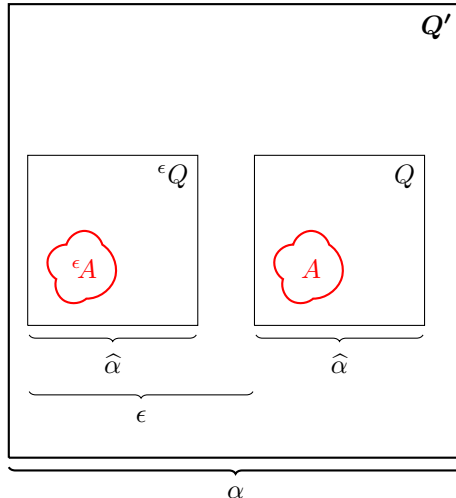
The figure illustrates the relations (33). The size constraint on $\widehat{\alpha}$ guarantees the existence of a subsquare Q' of side length α that contains both a given $Q \in \mathcal{Q}(\widehat{\alpha})$ and its left translate ${}^\epsilon Q$. Similarly, the size constraint on ϵ guarantees that Q and ${}^\epsilon Q$ are disjoint, so the Borel sets $A \subseteq Q$ and ${}^\epsilon A \subseteq {}^\epsilon Q$ are also disjoint.

where in the second inequality we used the variance condition (19) for $\mathcal{X}$ and the bounds $2|b'' - a'' + 1| = |b' - a' + 1| \leq |b - a + 1|$. By (35) we see that $\widehat{\mathcal{X}}$ satisfies the variance condition (19), and one can take $\widehat{C} = C + 32N^2$ as a generous choice for the required constant. Finally, we note that in case $\widehat{\mathcal{X}}[a : b]$ does not contain a complete block, then we have $|b - a + 1| < 4N$. In this case, it is trivial that one has the variance condition (19) for $\widehat{\mathcal{X}}[a : b]$ when one takes $\widehat{C}$ for the required constant.

To check the local uniformity condition (20) for $\widehat{\mathcal{X}}$, we fix $a \leq b$ and A as before so we continue to have the relations (33). In the case when $\widehat{\mathcal{X}}[a : b]$ does not contain a complete block, then $|b - a + 1| < 4N$ and we can simply take $\mathcal{S}$ required by (20) to be the empty set. In this case the conditions required by (20) all hold trivially. Thus, without loss of generality, we assume that $\widehat{\mathcal{X}}[a : b]$ contains at least one complete block.

By our hypothesis, we have condition (20) for the process $\mathcal{X}$, and in this case we can apply (20) to the Borel set ${}^\epsilon A \cup A \subseteq Q \in \mathcal{Q}(\alpha)$ and the interval $[a'' : b'']$. We are then guaranteed the existence of a set

(36) $$\mathcal{S} \subseteq \{X_t : X_t \in {}^\epsilon A \cup A \text{ and } t \in [a'' : b'']\}$$

for which we have both the cardinality bounds

(37) $$N({}^\epsilon A \cup A, \mathcal{X}[a'' : b'']) - M \leq |\mathcal{S}| \leq N({}^\epsilon A \cup A, \mathcal{X}[a'' : b''])$$

and the distributional identity $\mathcal{S} \overset{\mathrm{dpp}}{=} \{U_s({}^\epsilon A \cup A) : 1 \leq s \leq |\mathcal{S}|\}$. Now it only remains to check that if we take

$$\widehat{\mathcal{S}} = \{\mathcal{S} \cap {}^\epsilon A\}^\epsilon \cup \{\mathcal{S} \cap A\} \subseteq A$$

then the set $\widehat{\mathcal{S}}$ will suffice to confirm that we have the local uniformity condition (20) for $\widehat{\mathcal{X}}$.

First, we note from the inclusion (36) that we have

(38) $\quad \widehat{\mathcal{S}} \subseteq \{X_t : X_t \in {}^\epsilon A \text{ and } t \in [a'' : b'']\}^\epsilon \cup \{X_t : X_t \in A \text{ and } t \in [a'' : b'']\}$

$\qquad = \{\widehat{X}_t : \widehat{X}_t \in A \text{ and } t \in [a' : b']\} \subseteq \{\widehat{X}_t : \widehat{X}_t \in A \text{ and } t \in [a : b]\},$

where, for the middle equality we used the fact that $\widehat{\mathcal{X}}[a' : b']$ contains only complete blocks and the fact that $\mathcal{X}[a'' : b'']$ contains only the unshifted halves of the same blocks. The last inclusion follows simply because $a \leq a' < b' \leq b$.

Now, we already know from (29) that the point process $\widehat{\mathcal{S}}$ has the same distribution as an independent sample of $|\widehat{\mathcal{S}}|$ points with the uniform distribution on A, so to complete the proof of the lemma, we just need to check the double bound

(39) $\qquad N(A, \widehat{\mathcal{X}}[a, b]) - M - 4N \leq |\widehat{\mathcal{S}}| \leq N(A, \widehat{\mathcal{X}}[a, b]).$

The upper bound follows from the last inclusion of (38). For the lower bound, we start with the decomposition (34). The uniform bound $W(A, \widehat{\mathcal{X}}[a : b]) \leq 4N$ and the left inequality of (37) then give us

$$N(A, \widehat{\mathcal{X}}[a, b]) - M - 4N \leq |\mathcal{S}| = |\widehat{\mathcal{S}}|,$$

and this lower bound completes the check of (39). $\qquad \square$

Proposition 9 takes us much closer to the proof of Theorems 1 and 2. In particular, it is now easy to show that $T_{\epsilon, N}$ also preserves local uniformity. We establish this as a corollary to a slightly more general lemma.

Lemma 10 (Independent Random Shift of Index). *Let $\mathcal{Y}$ be an (α, M) locally uniform process with $0 < \alpha \leq 1$ and $M \in \mathbb{N}$. If the random variable I has the uniform distribution on $\{0, 1, \ldots, K - 1\}$ and if I is independent of $\mathcal{Y}$, then the process $\mathcal{Y}' = (Y'_t : t \in \mathbb{Z})$ defined by setting*

$$Y'_t = Y_{t+I} \quad \text{for all } t \in \mathbb{Z}$$

is an $(\alpha, M + K)$ locally uniform process.

Proof. If $b < a + K$, then the result is trivial since one can simply take $\mathcal{S} = \emptyset$. On the other hand, if $a + K \leq b$ then $0 \leq I < K$, and we have

$$a \leq a + I < a + K \leq b \leq b + I < b + K.$$

Thus, for any $A \subseteq Q \in \mathcal{Q}(\alpha)$ we have

$$N(A, \mathcal{Y}[a + K : b]) \leq N(A, \mathcal{Y}'[a : b]) \leq N(A, \mathcal{Y}[a : b + K]),$$

and we can write

(40) $\qquad N(A, \mathcal{Y}'[a : b]) = N(A, \mathcal{Y}[a + K : b]) + W(A, \mathcal{Y}'[a : b]),$

where $W(A, \mathcal{Y}'[a : b])$ is a random variable such that $0 \leq W(A, \mathcal{Y}'[a : b]) \leq K$. Taking the variance in (40) then gives

$$\mathrm{Var}[N(A, \mathcal{Y}'[a : b])] \leq 2\,\mathrm{Var}[N(A, \mathcal{Y}[a + K : b])] + 2\,\mathrm{Var}[W(A, \mathcal{Y}'[a : b])]$$

$$\leq 2C|b - a - K + 1| + 2K^2 \leq (2C + 2K^2)|b - a + 1|,$$

where in the second inequality we used the variance condition (19) for the process $\mathcal{Y}$ together with the almost sure bound on $W(A, \mathcal{Y}'[a : b])$. The last inequality confirms the variance condition (19) for the process $\mathcal{Y}'$.

To check the local uniformity condition (20) for $\mathcal{Y}'$, we take integers $a \leq b$, and we fix a Borel set $A \subseteq Q \in \mathcal{Q}(\alpha)$. By the (α, M) local uniformity of $\mathcal{Y}$ we know that there is a set $\mathcal{S} \subseteq \{Y_t : Y_t \in A \text{ and } t \in [a+K:b]\}$ such that

$$(41) \qquad N(A, \mathcal{Y}[a+K:b]) - M \leq |\mathcal{S}| \leq N(A, \mathcal{Y}[a+K:b])$$

and $\mathcal{S} \overset{\text{dpp}}{=} \{U_s(A) : 1 \leq s \leq |\mathcal{S}|\}$.

We now just need to check we have the local uniformity condition (20) for $\mathcal{Y}'$ if we take $\mathcal{S}' = \mathcal{S}$. With this choice, the bounds $0 \leq I < K$ give us the inclusion

$$\mathcal{S}' \subseteq \{Y_t : Y_t \in A \text{ and } t \in [a+K:b]\}$$
$$\subseteq \{Y_t : Y_t \in A \text{ and } t \in [a+I:b+I]\} = \{Y_t' : Y_t' \in A \text{ and } t \in [a:b]\}.$$

The set $\mathcal{S}'$ inherits the distributional requirement of (20) from $\mathcal{S}$, so we only need to check its cardinality requirements. From (40) and (41) we get the upper bound

$$|\mathcal{S}'| = |\mathcal{S}| \leq N(A, \mathcal{Y}[a+K:b]) \leq N(A, \mathcal{Y}'[a:b]),$$

as well as the lower bound

$$N(A, \mathcal{Y}'[a:b]) - K - M \leq N(A, \mathcal{Y}[a+K:b]) - M \leq |\mathcal{S}| = |\mathcal{S}'|,$$

so the cardinality requirements of (20) also hold. $\qquad\square$

A $T_{\epsilon,N}$ transformation is just a $H_{\epsilon,N}$ transformation followed by an index shift by an independent random variable I that has the uniform distribution on the set of indices $[0 : 2N-1]$. Proposition 9 and Lemma 10 tell us that each of these actions preserves local uniformity. These observations give us the final lemma of the section, where, for simplicity, we take a generous value for the size $\widetilde{M}$ of the exception set.

Lemma 11 (Local Uniformity and $T_{\epsilon,N}$ Transformations). *If $\mathcal{X}$ is an (α, M) locally uniform process for some $0 < \alpha \leq 1$ and $M \in \mathbb{N}$, then for each $0 < \epsilon < \alpha$ and $N < \infty$ the process*

$$\widetilde{\mathcal{X}} = T_{\epsilon,N}(\mathcal{X}),$$

is an $(\widetilde{\alpha}, \widetilde{M})$ locally uniform process where

$$0 < \widetilde{\alpha} < \min\{\epsilon, \alpha - \epsilon\} \quad and \quad \widetilde{M} = M + 6N.$$

5. Iterated $T_{\epsilon,N}$ Transformations and a Limit Process

We now consider the construction of a process $\mathcal{X}^* = (X_t^* : t \in \mathbb{Z})$ that we obtain as a limit of iterated $T_{\epsilon,N}$ transformations. First we fix a sequence of integers $1 \leq N_1 < N_2 < \cdots$ and a sequence of real numbers that we write as $1 > \epsilon_1 > \epsilon_2 > \cdots > 0$. Next, we let $\mathcal{X}^{(0)} = (X_t^{(0)} : t \in \mathbb{Z})$ be the doubly infinite sequence of independent random variables with the uniform distribution on $\mathcal{T}$, and we consider the infinite sequence of stationary processes $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \mathcal{X}^{(3)}, \ldots$ that one obtains by successive applications of appropriate $T_{\epsilon,N}$ transformations:

$$(42) \quad \mathcal{X}^{(1)} = T_{\epsilon_1, N_1}(\mathcal{X}^{(0)}), \quad \mathcal{X}^{(2)} = T_{\epsilon_2, N_2}(\mathcal{X}^{(1)}), \quad \mathcal{X}^{(3)} = T_{\epsilon_3, N_3}(\mathcal{X}^{(2)}), \quad \cdots.$$

We now let $\mathcal{T}^\infty$ be the set of doubly infinite sequences $\boldsymbol{x} = (\ldots, x_{-1}, x_0, x_1, \ldots)$ where $x_t \in \mathcal{T}$ for each $t \in \mathbb{Z}$, and we view $\mathcal{T}^\infty$ as a topological space with respect to the product topology. By $\mathcal{B}(\mathcal{T}^\infty)$ we denote the σ-field of Borel sets of $\mathcal{T}^\infty$, and we let $\mathcal{B}(\mathcal{T}^{[-m:m]})$ denote the smallest sub-σ-field of $\mathcal{B}(\mathcal{T}^\infty)$ such that each continuous function $f : \mathcal{T}^\infty \mapsto \mathbb{R}$ of the form $\boldsymbol{x} \mapsto f(x_{-m}, \ldots, x_{-1}, x_0, x_1, \ldots, x_m)$

is $\mathcal{B}(\mathcal{T}^{[-m:m]})$ measurable. In less formal language, $\mathcal{B}(\mathcal{T}^{[-m:m]})$ is the subset of elements of $\mathcal{B}(\mathcal{T}^{\infty})$ that do not depend on x_t for $|t| > m$.

Next, we take $\mathcal{M}$ to be the set of Borel probability measures on $\mathcal{T}^{\infty}$, and we note that $\mathcal{M}$ is a complete metric space if one defines the distance $\rho(\mu, \mu')$ between the Borel measures μ and μ' by setting

$$(43) \qquad \rho(\mu, \mu') = \sum_{m=1}^{\infty} 2^{-m} \sup\{\, |\mu(\mathcal{A}) - \mu'(\mathcal{A})| \, : \, \mathcal{A} \in \mathcal{B}(\mathcal{T}^{[-m:m]}) \,\}.$$

To show that the sequence $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \mathcal{X}^{(3)}, \ldots$ converges in distribution to a process $\mathcal{X}^*$, it suffices to show that if we define the measures $\mu_1, \mu_2, \ldots$ on $\mathcal{B}(\mathcal{T}^{\infty})$ by setting

$$(44) \qquad \mu_j(\mathcal{A}) = P(\mathcal{X}^{(j)} \in \mathcal{A}), \qquad \text{for each } \mathcal{A} \in \mathcal{B}(\mathcal{T}^{\infty}),$$

then the sequence $\mu_1, \mu_2, \ldots$ is a Cauchy sequence under the metric ρ. Fortunately, the Cauchy criterion can be verified under a mild condition on the defining sequence of integers $N_1, N_2, \ldots$.

Lemma 12 (A Condition for Convergence). *If the processes* $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \mathcal{X}^{(3)}, \ldots$ *are defined by the iterated* $T_{\epsilon,N}$ *transformations* (42) *and if*

$$\sum_{j=1}^{\infty} \frac{1}{N_j} < \infty,$$

then the sequence of processes $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \mathcal{X}^{(3)}, \ldots$ *converges in distribution to a stationary, translation invariant process* $\mathcal{X}^* = (\ldots, X_{-1}^*, X_0^*, X_1^*, \ldots)$ *such that* X_1^* *has the uniform distributed on* $\mathcal{T}$.

Proof. By the closeness inequality (16) and the definition (44) of μ_j, we have for all $m = 1, 2, \ldots$ that

$$\sup\{|\mu_j(\mathcal{A}) - \mu_{j+1}(\mathcal{A})| \, : \, \mathcal{A} \in \mathcal{B}(\mathcal{T}^{[-m:m]})\} \leq \frac{2m}{N_{j+1}}.$$

The definition (43) of the metric ρ and a simple summation then give us

$$\rho(\mu_j, \mu_{j+1}) \leq \frac{4}{N_{j+1}},$$

so, by the completeness of the metric space $(\mathcal{M}, \rho)$, the sequence of processes $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \mathcal{X}^{(3)}, \ldots$ converges in distribution to a process $\mathcal{X}^*$. By Lemmas 3, 4, and 5, we know that each of the processes $\mathcal{X}^{(j)}$ is stationary and translation invariant. Moreover, each of these processes has uniform marginal distributions. The process $\mathcal{X}^*$ inherits all of these properties through convergence in distribution. $\square$

6. Path Lengths for the Limit Process

The next lemma expresses a kind of Lipschitz property for the shortest path functional. Specifically, it bounds the absolute difference in the expected value of $L(Z)$ and $L(\widetilde{Z})$, where Z and $\widetilde{Z}$ are arbitrary n-dimensional random vectors with values in $\mathcal{T}^n$ and where for $(z_1, z_2, \ldots, z_n) \in \mathcal{T}^n$, we write $L(z_1, z_2, \ldots, z_n)$ for the length of the shortest path through the points $z_1, z_2, \ldots, z_n$.

The lemma is stated and proved for general Z and $\widetilde{Z}$, but our typical choice will be $Z = \mathcal{X}[0 : n-1]$ and $\widetilde{Z} = \widetilde{\mathcal{X}}[0 : n-1]$. We also recall that if $\mathcal{B}(\mathcal{T}^n)$ denotes the

set of all Borel subsets of $\mathcal{T}^n$, then the *total variation distance* between Z and $\widetilde{Z}$ is given by

$$d_{\mathrm{TV}}(Z, \widetilde{Z}) = \sup\{|\mathbb{P}(Z \in \mathcal{A}) - \mathbb{P}(\widetilde{Z} \in \mathcal{A})| : \mathcal{A} \in \mathcal{B}(\mathcal{T}^n)\}.$$

We also recall that the function $(z_1, z_2, \ldots, z_n) \mapsto n^{-1/2}L(z_1, z_2, \ldots, z_n)$ is uniformly bounded; in fact, by early work of Few (1955), it is bounded by 3.

Lemma 13. *For all random vectors Z and $\widetilde{Z}$ with values in $\mathcal{T}^n$ we have*

$$|\mathbb{E}[L(Z)] - \mathbb{E}[L(\widetilde{Z})]| \leq 3n^{1/2}d_{\mathrm{TV}}(Z, \widetilde{Z}).$$

Proof. By the maximal coupling theorem (Lindvall, 2002, Theorem 5.2) there exists a probability space and a random pair $(Z', \widetilde{Z}')$ such that $Z' \overset{d}{=} Z$, $\widetilde{Z}' \overset{d}{=} \widetilde{Z}$ and

$$\mathbb{P}(Z' \neq \widetilde{Z}') = d_{\mathrm{TV}}(Z, \widetilde{Z}).$$

Now, if we set $L_n^* = \max\{L(z_1, z_2, \ldots, z_n) : z_t \in \mathcal{T} \text{ and } t \in [1:n]\}$, then we have

$$\begin{aligned}
|\mathbb{E}[L(Z)] - \mathbb{E}[L(\widetilde{Z})]| &= |\mathbb{E}[L(Z')] - \mathbb{E}[L(\widetilde{Z}')]| \\
&\leq \mathbb{P}(Z' \neq \widetilde{Z}')L_n^* \\
&\leq 3n^{1/2}d_{\mathrm{TV}}(Z, \widetilde{Z}),
\end{aligned}$$

where in the last line we used the classic bound $L_n^* \leq 3n^{1/2}$ from Few (1955). $\square$

The immediate benefit of Lemma 13 is that it gives us a way to estimate the cost of a minimal path through the points $X_0^*, X_1^*, \ldots, X_{n-1}^*$.

Lemma 14 (Shortest Path Differences in the Limit). *For all $0 \leq j < \infty$ and all $n \geq 1$ we have*

$$(45) \qquad |\mathbb{E}[L(\mathcal{X}^{(j)}[0:n-1])] - \mathbb{E}[L(\mathcal{X}^*[0:n-1])]| \leq 3\,n^{3/2}\sum_{k=1}^{\infty} \frac{1}{N_{j+k}}$$

Proof. Using the shorthand $\mathcal{X}_n^{(j+k)} = \mathcal{X}^{(j+k)}[0:n-1]$ for all $k \in \{0, 1, 2, \ldots\}$, one has by the triangle inequality that

$$(46) \qquad |\mathbb{E}[L(\mathcal{X}_n^{(j)})] - \mathbb{E}[L(\mathcal{X}_n^*)]| \leq \sum_{k=1}^{\infty} |\mathbb{E}[L(\mathcal{X}_n^{(j+k)})] - \mathbb{E}[L(\mathcal{X}_n^{(j+k-1)})]|.$$

Lemma 13 then tells us that

$$(47) \qquad |\mathbb{E}[L(\mathcal{X}_n^{(j+k)})] - \mathbb{E}[L(\mathcal{X}_n^{(j+k-1)})]| \leq 3\,n^{1/2}d_{\mathrm{TV}}(\mathcal{X}_n^{(j+k)}, \mathcal{X}_n^{(j+k-1)}),$$

and Lemma 6 implies that

$$(48) \qquad d_{\mathrm{TV}}(\mathcal{X}_n^{(j+k)}, \mathcal{X}_n^{(j+k-1)}) \leq \frac{n}{N_{j+k}},$$

so using (47) and (48) in the sum (46) completes the proof of the lemma. $\square$

7. Parameter Choices

To pass from the general iterative construction (42) to the process required by Theorem 2, we need to make parameter choices that go beyond those required by Lemma 12 on a sufficient condition for convergence.

First we fix a sequence $\eta_1, \eta_2, \ldots$ of values in $(0, 1)$ that decrease monotonically to zero as $j \to \infty$; these values just serve to provide us with a measure of smallness of scale. We then inductively define the values N_j and ϵ_j through which we finally define $\mathcal{X}^*$ by the sequence (42) of transformations T_{ϵ_j, N_j}, $j = 1, 2, \ldots$. To begin the construction, we can take any $\epsilon_1 \in (0, 1)$ and any integer $N_1 \geq 2$. Subsequent values with $j \geq 2$ are determined by two rules:

(49)   RULE 1. We choose an integer N_j such that $N_j > j^2 N_{j-1}$ and such that

$$| \mathbb{E}[L(\mathcal{X}^{(j-1)}[0 : n-1])] - \beta n^{1/2} | \leq \eta_j n^{1/2} \qquad \text{for all } n \geq \lfloor j^{-1} N_j \rfloor.$$

(50)   RULE 2. We choose an $\epsilon_j \in (0, \epsilon_{j-1})$ such that

$$\epsilon_j j^{1/2} N_j^{1/2} \leq \eta_j.$$

The first rule leans on the fact that each of the processes $\mathcal{X}^{(j)}$, $1 \leq j < \infty$, is locally uniform so, from Proposition 8, we have that

$$n^{-1/2} \mathbb{E}[L(\mathcal{X}^{(j-1)}[0 : n-1])] \to \beta \qquad \text{as } n \to \infty.$$

This guarantees the existence of the integer N_j required by (49), and, once N_j is determined, it is trivial to choose ϵ_j to satisfy the second rule.

8. Estimation of Expected Path Lengths: Proof of Theorem 2

The j'th stage of the construction (42) takes a doubly infinite sequence $\mathcal{X}^{(j-1)}$ to another doubly infinite sequence $\mathcal{X}^{(j)}$ by a T_{ϵ_j, N_j} transformation that we define in two steps. Specifically, if I is a random variable with the uniform distribution on $\{0, 1, \ldots, 2N_j - 1\}$ that is independent of $\mathcal{X}^{(j-1)}$, then we have

$$\widehat{\mathcal{X}}^{(j-1)} = H_{\epsilon_j, N_j}(\mathcal{X}^{(j-1)}) \quad \text{and} \quad X_t^{(j)} = \widehat{X}_{t+I}^{(j-1)} \quad \text{for all } t \in \mathbb{Z}.$$

We now claim that we have the bounds

(51) $L(\mathcal{X}^{(j)}[0 : 2jN_j - 1]) \leq L(\widehat{\mathcal{X}}^{(j-1)}[0 : 2(j+1)N_j - 1])$

$$\leq L(\mathcal{X}^{(j-1)}[0 : (j+1)N_j - 1]) + 2\epsilon_j(j+1)N_j.$$

To check the first inequality, we recall that $0 \leq I < 2N_j - 1$ so the set

$$\mathcal{W} = \{ \widehat{X}_t^{(j-1)} : 0 \leq t \leq 2(j+1)N_j - 1 \}$$

is a superset of $\{ \widehat{X}_t^{(j-1)} : I \leq t \leq I + 2jN_j - 1 \} = \{ X_t^{(j)} : 0 \leq t \leq 2jN_j - 1 \}$.

To check the second inequality of (51) takes more work. We first note that the segment $\widehat{\mathcal{X}}^{(j-1)}[0 : 2(j+1)N_j - 1]$ contains only complete blocks, so one also has

$$\mathcal{W} = \{ X_t^{(j-1)} : 0 \leq t \leq (j+1)N_j - 1 \} \cup \{ X_t^{(j-1)}(\epsilon) : 0 \leq t \leq (j+1)N_j - 1 \},$$

and, to prove the second inequality of (51), we construct a suboptimal path through the points of $\mathcal{W}$.

To build this path, we just take an optimal path $\mathcal{P}$ through the set of points $\{ X_t^{(j-1)} : 0 \leq t \leq (j+1)N_j - 1 \}$ and then for each $t \in [0 : (j+1)N_j - 1]$ we adjoin to $\mathcal{P}$ the loop that goes from $X_t^{(j-1)}$ to its shifted twin $X_t^{(j-1)}(\epsilon_j)$ and back

to $X_t^{(j-1)}$. The suboptimal path has $(j+1)N_j$ loops and each of these has length $2\epsilon_j$, so the suboptimal path through $\mathcal{W}$ has a cost that is bounded by the last sum in (51).

Now, when we take expectations in (51), the rules (49) and (50) give us

$$(2jN_j)^{-1/2}\mathbb{E}[L(\mathcal{X}^{(j)}[0\!:\!2jN_j-1])] \leq (\beta+\eta_j)(1+j^{-1/2})2^{-1/2}+2\epsilon_j(j^{1/2}+j^{-1/2})N_j^{1/2}$$

$$\leq (\beta+\eta_j)(1+j^{-1/2})2^{-1/2}+4\eta_j.$$

This bound together with (45) for $n=2jN_j$ then gives us

$$(2jN_j)^{-1/2}\mathbb{E}[L(\mathcal{X}^*[0:2jN_j-1])] \leq (\beta+\eta_j)(1+j^{-1/2})2^{-1/2}+4\eta_j+6jN_j\sum_{k=1}^{\infty}\frac{1}{N_{j+k}}$$

$$\leq (\beta+\eta_j)(1+j^{-1/2})2^{-1/2}+4\eta_j+6/j,$$

where, in the second inequality, we estimate the sum using the strict inequality $j^{2k}N_j < N_{j+k}$ which holds for all $k \geq 1$ and $j \geq 2$ by the first part of our first parameter formation rule (49). This last displayed bound is more than one needs to complete the proof of the first inequality (4) of Theorem 2.

The second inequality (4) of Theorem 2 is easier. If we take $n = \lfloor j^{-1}N_j \rfloor$ in the estimate (49) given by our first parameter rule, then we have

$$\beta - \eta_j \leq \lfloor j^{-1}N_j \rfloor^{-1/2}\mathbb{E}[L(\mathcal{X}^{(j-1)}[0:\lfloor j^{-1}N_j \rfloor-1])].$$

Now, if we use the bound (45) for $n = \lfloor j^{-1}N_j \rfloor$ and if we estimate the infinite sum exactly as we did before, then we have

$$\beta - \eta_j - 6/j \leq \lfloor j^{-1}N_j \rfloor^{-1/2}\mathbb{E}[L(\mathcal{X}^*[0:\lfloor j^{-1}N_j \rfloor-1])],$$

and this bound more than one needs to complete the proof of second inequality (4) of Theorem 2.

9. THEOREM 2 IMPLIES THEOREM 1

We now show that Theorem 1 follows from Theorem 2 and Choquet's representation theorem. To set this up, we first consider the set $\mathcal{U} \subset \mathcal{M}$ of Borel probability measures on $\mathcal{T}^\infty$ that are shift invariant and that have uniform marginals. Here the shift transformation $\Theta : \mathcal{T}^\infty \to \mathcal{T}^\infty$ is defined by setting $\Theta(\boldsymbol{x})_t = x_{t+1}$ where $\boldsymbol{x} = (\ldots, x_{-1}, x_0, x_1, \ldots) \in \mathcal{T}^\infty$, and a Borel measure ν on $\mathcal{T}^\infty$ is said to be shift invariant if $\nu(\Theta^{-1}(\mathcal{A})) = \nu(\mathcal{A})$ for each Borel set $\mathcal{A} \subseteq \mathcal{T}^\infty$. Finally, to say that a measure ν on $\mathcal{T}^\infty$ has uniform marginals just means that for each Borel set $A \subseteq \mathcal{T}$ and for each fixed $t \in \mathbb{Z}$ we have $\nu(\mathcal{A}_t) = \lambda(A)$ when we set $\mathcal{A}_t = \{\boldsymbol{x} : x_t \in A\}$ and where $\lambda(A)$ denotes the Lebesgue measure of the set $A \subseteq \mathcal{T}$.

Under the product topology the space $\mathcal{T}^\infty$ is compact, and, under the topology of weak convergence, the space $\mathcal{M}$ of probability measures on $\mathcal{T}^\infty$ is metrizable, locally convex, and compact. The set $\mathcal{U} \subset \mathcal{M}$ of shift invariant measures with uniform marginals is trivially convex, and $\mathcal{U}$ is a closed subset of $\mathcal{M}$. Hence, under the topology of weak convergence, the set $\mathcal{U}$ is also compact.

We now consider the process $\mathcal{X}^*$ given by Theorem 2, and we define a measure μ on $\mathcal{T}^\infty$ by setting

(52) $$\mu(\mathcal{A}) = P(\mathcal{X}^* \in \mathcal{A}) \quad \text{for each Borel set } \mathcal{A} \subseteq \mathcal{T}^\infty.$$

Stationarity of $\mathcal{X}^*$ tells us that μ is invariant under the shift transformation Θ, and, since μ inherits the uniform marginal property from $\mathcal{X}^*$, we have $\mu \in \mathcal{U}$.

Next, we denote the symmetric difference between the sets $\mathcal{A}$ and $\mathcal{A}'$ by $\mathcal{A}\Delta\mathcal{A}'$, and we recall that a probability measure ν on $\mathcal{T}^\infty$ is *ergodic* if for any Borel $\mathcal{A} \subseteq \mathcal{T}^\infty$, one has

$$\nu(\mathcal{A}\Delta\Theta^{-1}(\mathcal{A})) = 0 \quad \text{if and only if} \quad \nu(\mathcal{A}) = 0 \text{ or } \nu(\mathcal{A}) = 1.$$

Moreover, if we let $\mathcal{U}_e$ denote the set of of extreme points of the convex set $\mathcal{U}$, then it is a very pleasing fact that $\mathcal{U}$ is exactly equal to the set of ergodic measures in $\mathcal{U}$, (cf. Simon, 2011, Example 8.17, pp. 129–131.)

Choquet's representation theorem (cf. Simon, 2011, Theorem 10.7) now tells us that there is a probability measure D_μ with support on $\mathcal{U}_e \subset \mathcal{U}$ such that

$$(53) \qquad \mu(\mathcal{A}) = \int_{\mathcal{U}_e} \nu(\mathcal{A})\, D_\mu(d\nu) \qquad \text{for each Borel } \mathcal{A} \subseteq \mathcal{T}^\infty.$$

In other words, any shift invariant probability measure on $\mathcal{T}^\infty$ with uniform marginal distributions is an average of ergodic, shift invariant probability measures on $\mathcal{T}^\infty$ that have uniform marginal distributions. This representation will lead us quickly to Theorem 1.

For an $\boldsymbol{x} = (\ldots, x_{-1}, x_0, x_1, \ldots) \in \mathcal{T}^\infty$ we write the $[1:n]$-segment of $\boldsymbol{x}$ with the shorthand $\boldsymbol{x}_n = (x_1, x_2, \ldots, x_n)$, and, for constants $c_1 < c_2$, we consider the set

$$\Gamma[c_1, c_2] = \big\{ \boldsymbol{x} \in \mathcal{T}^\infty : \liminf_{n\to\infty} n^{-1/2} L(\boldsymbol{x}_n) \le c_1 \text{ and } c_2 \le \limsup_{n\to\infty} n^{-1/2} L(\boldsymbol{x}_n) \big\}.$$

If $\mu(\Gamma[c_1, c_2]) = 0$ for all $c_1 < c_2$, then $n^{-1/2} L(\boldsymbol{x}_n)$ converges with μ-probability one. Since this ratio is bounded, the dominated convergence theorem implies the convergence of the expectations $n^{-1/2}\mathbb{E}_\mu[L(\boldsymbol{x}_n)] = n^{-1/2}\mathbb{E}[L(\mathcal{X}^*[1:n])]$, where in the last equality we just used the definition (52) of μ. By Theorem 2 we know that we do not have the convergence of these expectations, so there must be some pair $c_1 < c_2$ for which we have $0 < \mu(\Gamma[c_1, c_2])$. Finally, by the representation (53) we have

$$0 < \mu(\Gamma[c_1, c_2]) = \int_{\mathcal{U}_e} \nu(\Gamma[c_1, c_2])\, D_\mu(d\nu),$$

so there is some ergodic measure $\nu \in \mathcal{U}_e$ for which we have $0 < \nu(\Gamma[c_1, c_2])$.

The sets $\Gamma[c_1, c_2]$ and $\Theta^{-1}(\Gamma[c_1, c_2])$ are identical, so $\Gamma[c_1, c_2]$ is an invariant set for ν, or any other measure. Since $0 < \nu(\Gamma[c_1, c_2])$, the ergodicity of $\nu \in \mathcal{U}_e$ gives us $\nu(\Gamma[c_1, c_2]) = 1$. Finally, if we take $\mathcal{X}$ to be the stationary process determined by the shift transformation Θ and the ergodic measure ν, then we have

$$(54) \qquad P(\mathcal{X} \in \mathcal{A}) = \nu(\mathcal{A}) \quad \text{for each Borel set } \mathcal{A} \subseteq \mathcal{T}^\infty.$$

By construction, the process $\mathcal{X}$ is stationary and ergodic, and $\mathcal{X}$ also inherits from ν the property of uniform marginal distributions. Finally, by (54) and the definition of $\Gamma[c_1, c_2]$, we see that $\mathcal{X}$ has all of the features required by Theorem 1.

Remark 15. Here, instead of using Choquet's theorem, one could consider using the ergodic decomposition theorems of Krylov and Bogolioubov (cf. Dynkin 1978 or Kallenberg 2002, Theorem 10.26, p. 196), but, since the usual statements of these theorems to not immediately accommodate the restriction to measures with uniform marginal distributions, it is simpler to work directly with Choquet's theorem where no modifications are required.

For a full treatment of Choquet's theorem one can consult Phelps (2001) or Simon (2011), but for the existence theorem in the metrizable case one can appeal

more easily to the short proof of Bonsall (1963) which uses little more than the Hahn-Banach theorem and the Riesz representation theorem.

10. Extensions, Refinements, and Problems

There are easily proved analogs of Theorems 1 and 2 for many of the functionals of combinatorial optimization for which one has the analog of the Beardwood-Halton-Hammersley theorem. In particular, one can show that the analogs of Theorems 1 and 2 hold for the minimal spanning tree (MST) problem studied in Steele (1988) and for the minimal matching problem studied in Rhee (1993). In these cases, the construction of the processes in Theorems 1 and 2 needs almost no alteration. The main issue is that one needs to establish a proper analog of Proposition 8, but this is often easy. Once an analog of Proposition 8 is in hand, one only needs to make few cosmetic changes to the arguments of Section 8.

Still, there are interesting functionals for which it is not as clear how one can adapt the proofs of Theorems 1 and 2. One engaging example is the sum of the edge lengths of the Voronoi tessellation. In this case, the analog of the BHH theorem was developed by Miles (1970) for Poisson sample sizes, and later by McGivney and Yukich (1999) for fixed sample sizes (and with complete convergence). A second, much different example, is the length of the path that one obtains by running the Karp-Held algorithm for the TSP. The expressly algorithmic nature of this functional introduces several new twists, but, nevertheless, Goemans and Bertsimas (1991) obtained the analog of the BHH theorem.

These two functionals are "less local" than the TSP, MST, or minimal matching functionals; in particular, they are not as amenable to suboptimal patching bounds such as those we used in Section 8. Nevertheless, these functionals are sufficiently local to allow for analogs of the BHH theorem, so it seems probable that the natural analogs of Theorems 1 and 2 would hold as well.

There are two further points worth noting. First, at the cost of using more complicated versions of the transformations $H_{\epsilon,N}$ and $T_{\epsilon,N}$, one can replace the infimum bound $2^{-1/2}\beta$ of Theorem 2 with a smaller constant. Since the method of Section 9 shows that any infimum bound less that β suffices to prove Theorem 1, we did not pursue the issue of a minimal infimum bound.

Finally, it is feasible that the process $(X_t^* : t \in \mathbb{Z})$ that was constructed for the proof of Theorem 2 is itself ergodic — or even mixing. If this could be established, then one would not need the Choquet's integral representation argument of Section 9. Unfortunately, it does not seem easy to prove that the process $(X_t^* : t \in \mathbb{Z})$ is ergodic, even though this may be somewhat intuitive.

Acknowledgments

References

Applegate, D. L., Bixby, R. E., Chvátal, V. and Cook, W. J. (2006), *The traveling salesman problem*, Princeton Series in Applied Mathematics, Princeton University Press, Princeton, NJ. A computational study.

Arora, S. (1998), 'Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems', *J. ACM* **45**(5), 753–782.

Avram, F. and Bertsimas, D. (1992), 'The minimum spanning tree constant in geometrical probability and under the independent model: a unified approach', *Ann. Appl. Probab.* **2**(1), 113–130.

Beardwood, J., Halton, J. H. and Hammersley, J. M. (1959), 'The shortest path through many points', *Proc. Cambridge Philos. Soc.* **55**, 299–327.

Bertsimas, D. J., Jaillet, P. and Odoni, A. R. (1990), 'A priori optimization', *Oper. Res.* **38**(6), 1019–1033.

Bertsimas, D. J. and van Ryzin, G. (1993), 'Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles', *Oper. Res.* **41**(1), 60–76.

Bonsall, F. F. (1963), 'On the representation of points of a convex set', *J. London Math. Soc.* **38**, 332–334.

Dynkin, E. B. (1978), 'Sufficient statistics and extreme points', *Ann. Probab.* **6**(5), 705–730.

Few, L. (1955), 'The shortest path and the shortest road through n points', *Mathematika* **2**, 141–144.

Finch, S. R. (2003), *Mathematical constants*, Vol. 94 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, Cambridge.

Goemans, M. X. and Bertsimas, D. J. (1991), 'Probabilistic analysis of the Held and Karp lower bound for the Euclidean traveling salesman problem', *Math. Oper. Res.* **16**(1), 72–89.

Gutin, G. and Punnen, A. P., eds (2002), *The traveling salesman problem and its variations*, Vol. 12 of *Combinatorial Optimization*, Kluwer Academic Publishers, Dordrecht.

Hwang, F. K., Richards, D. S. and Winter, P. (1992), *The Steiner tree problem*, Vol. 53 of *Annals of Discrete Mathematics*, North-Holland Publishing Co., Amsterdam.

Jaillet, P. (1993), 'Cube versus torus models and the Euclidean minimum spanning tree constant', *Ann. Appl. Probab.* **3**(2), 582–592.

Kallenberg, O. (2002), *Foundations of modern probability*, Probability and its Applications (New York), second edn, Springer-Verlag, New York.

Karp, R. M. (1977), 'Probabilistic analysis of partitioning algorithms for the traveling-salesman problem in the plane', *Math. Oper. Res.* **2**(3), 209–224.

Karp, R. M. and Steele, J. M. (1985), Probabilistic analysis of heuristics, *in* 'The traveling salesman problem', Wiley-Intersci. Ser. Discrete Math., Wiley, Chichester, pp. 181–205.

Lawler, E. L., Lenstra, J. K., Rinnooy Kan, A. H. G. and Shmoys, D. B., eds (1985), *The traveling salesman problem*, John Wiley & Sons Ltd., Chichester.

Lindvall, T. (2002), *Lectures on the coupling method*, Dover Publications Inc., Mineola, NY. Corrected reprint of the 1992 original.

McGivney, K. and Yukich, J. E. (1999), 'Asymptotics for Voronoi tessellations on random samples', *Stochastic Process. Appl.* **83**(2), 273–288.

Miles, R. E. (1970), 'On the homogeneous planar Poisson point process', *Math. Biosci.* **6**, 85–127.

Pavone, M., Bisnik, N., Frazzoli, E. and Isler, V. (2007), Decentralized vehicle routing in a stochastic and dynamic environment with customer impatience, *in* 'Proceedings of the 1st International Conference on Robot Communication and Coordination', RoboComm '07, IEEE Press, Piscataway, NJ, USA, pp. 24:1–24:8.

Phelps, R. R. (2001), *Lectures on Choquet's theorem*, Vol. 1757 of *Lecture Notes in Mathematics*, second edn, Springer-Verlag, Berlin.

Redmond, C. and Yukich, J. E. (1994), 'Limit theorems and rates of convergence for Euclidean functionals', *Ann. Appl. Probab.* **4**(4), 1057–1073.

Redmond, C. and Yukich, J. E. (1996), 'Asymptotics for Euclidean functionals with power-weighted edges', *Stochastic Process. Appl.* **61**(2), 289–304.

Rhee, W. T. (1993), 'A matching problem and subadditive Euclidean functionals', *Ann. Appl. Probab.* **3**(3), 794–801.

Schmidt, W. M. (1964), 'Metrical theorems on fractional parts of sequences', *Trans. Amer. Math. Soc.* **110**, 493–518.

Simon, B. (2011), *Convexity*, Vol. 187 of *Cambridge Tracts in Mathematics*, Cambridge University Press, Cambridge. An analytic viewpoint.

Steele, J. M. (1980), 'Shortest paths through pseudorandom points in the d-cube', *Proc. Amer. Math. Soc.* **80**(1), 130–134.

Steele, J. M. (1981), 'Complete convergence of short paths and Karp's algorithm for the TSP', *Math. Oper. Res.* **6**(3), 374–378.

Steele, J. M. (1988), 'Growth rates of Euclidean minimal spanning trees with power weighted edges', *Ann. Probab.* **16**(4), 1767–1787.

Steinerberger, S. (2010), 'A new lower bound for the geometric traveling salesman problem in terms of discrepancy', *Oper. Res. Lett.* **38**(4), 318–319.

Yukich, J. E. (1998), *Probability theory of classical Euclidean optimization problems*, Vol. 1675 of *Lecture Notes in Mathematics*, Springer-Verlag, Berlin.