in *Performance Evaluation of Computer and Communications Systems*, (L. Donatiello and R. Nelson, Eds.), 359–393, Springer-Verlag Lecture Notes in Computer Science, 1993.

AN OVERVIEW OF TES PROCESSES AND MODELING METHODOLOGY

BENJAMIN MELAMED

NEC USA, Inc.
C&C Research Laboratories
4 Independence Way
Princeton, New Jersey 08540

Abstract

TES (Transform-Expand-Sample) is a versatile methodology for modeling stationary time series with general marginal distributions and a broad range of dependence structures. From the viewpoint of Monte Carlo simulation, TES constitutes a new and flexible input analysis approach whose principal merit is its potential ability to simultaneously capture first-order and second-order statistics of empirical time series. That is, TES is designed to fit an arbitrary empirical marginal distribution (histogram), and to simultaneously approximate the leading empirical autocorrelations. This paper is a tutorial introduction to the theory of TES processes and to the modeling methodology based on it. It employs a didactic approach which relies heavily on visual intuition as a means of conveying key ideas and an aid in building deep understanding of TES. This approach is in line with practical TES modeling which itself is based on visual interaction under software support. The interaction takes on the form of a heuristic search in a large parameter space, and it currently relies on visual feedback supplied by computer graphics. The tutorial is structured around an illustrative example both to clarify the modeling methodology and to exemplify its efficacy.

Keywords and Phrases: Autocorrelated Processes, Input Analysis, Modeling and Simulation, Software Systems, Source Modeling, Stochastic Processes, TES Processes, TES Sequences.

1 INTRODUCTION

TES (Transform-Expand-Sample) is a recent approach to modeling stationary time series [15, 7, 8, 9]. The TES approach is non-parametric in the sense that it makes no assumptions on marginal distributions, though the underlying temporal dependence structure is postulated to be Markovian, with a continuous state space. Nevertheless, its modeling scope is quite broad: Additional transformations may be applied, leading to non-Markovian processes; and stationary TES processes may be combined into new ones, e.g., via modulation of TES processes by another process, often a discrete-state Markov process. The main application of TES to date has been to create source models (of incoming traffic or workload), in order to drive Monte Carlo simulations [2, 10].

What is new about TES is its potential ability to capture (fit) both the marginal distribution (a first-order statistic) and the autocorrelation function (a second-order statistic) of empirical data. Most importantly, TES aims to fit both marginals and autocorrelations *simultaneously*. This goal is not new; in fact, engineers have attempted such simultaneous fitting, mainly in the context of signal processing (see, e.g., [13] and references therein). The TES variation on this theme is to precisely fit the empirical marginal distribution (typically an empirical histogram), and at the same time capture *temporal dependence* proxied by the autocorrelation function (a measure of linear dependence). Being able to do this is no small feat. In fact, other modeling approaches to time series are able to do either one or the other but not both. For example, autoregression can fit a variety of autocorrelation functions, but not arbitrary marginal distributions. Conversely, the minification/maxification approach [12] can fit general marginal distributions, but has a relatively limited repertoire of autocorrelation functions. And approaches that use Gaussian processes as in [20] are difficult to implement in practice. All modeling approaches which try to capture marginals and autocorrelations simultaneously (TES included) appear to lack an effective algorithmic modeling procedure. TES itself employs a heuristic search over a large parameter space; nevertheless, experience shows that this is actually a viable approach which utilizes human visual perception to guide the search process.

Why should one insist on modeling both first-order and second-order properties simultaneously? From a purely philosophic viewpoint, it is clear that fitting more statistical aspects of empirical data is a priori desirable, as it can only increase one's confidence in a model. From a more practical vantage point, ignoring temporal dependence in data can often carry serious modeling risks. The perils inherent in failing to capture temporal dependence will be illustrated in the sequel. In recognition of these perils, the TES approach stipulates that fitting both first-order and second-order statistics is a minimal modeling requirement.

1.1 Temporal Dependence and Autocorrelations

Dependence, temporal and spatial, pervades many real-world random phenomena. Temporal dependence is a major cause of burstiness in telecommunications traffic, especially in emerging high-speed communications networks; typical examples are file transfers and compressed VBR (variable bit rate) video. Spatial dependence underpins reference of locality in caches and data bases. The combined effect of temporal and spatial dependencies underlies fault cascades observed in network management.

The autocorrelation function is a convenient measure of linear temporal dependence in real-valued stochastic processes, frequently used by engineers [1]. For a discrete-time, real-valued stationary stochastic process $\{X_n\}_{n=0}^{\infty}$, the autocorrelation function $\rho_X(\tau)$ consists of the lagged correlation coefficients

$$\rho_X(\tau) \;=\; \frac{E[X_n\, X_{n+\tau}] - \mu_X^2}{\sigma_X^2}, \quad \tau = 1, 2, \ldots \tag{1.1}$$

where $\mu_X < \infty$ and $\sigma_X^2 < \infty$ are the common mean and variance, respectively, of the X_n. While autocorrelation is a measure of linear dependence, it is often used to proxy for general dependence, and this is satisfactory in most cases.

1.2 The Importance of Modeling Temporal Dependence

Many systems, encountered in practice, have both deterministic and stochastic components. The deterministic components are usually easy to capture faithfully in a simulation model. However, the stochastic components tend to be the weak link in the chain of modeling assumptions in both simulation models and analytic ones. This tendency is most pronounced in analytical models, where oversimplified assumptions are routinely made to facilitate tractability. A typical case in point is renewal queueing, namely, GI/GI/m queues. Here dependence is eliminated or overlooked for the sake of analytical or numerical tractabilty, or because it is not clear how to model temporal dependence in arrivals or services. Occasionally, this attitude is justified, for example, when no additional knowledge is available, and one seeks qualitative insights and does not insist on quantitative accuracy. Analytical models often trade quantitative accuracy for analytical or numerical tractability. But where more accurate predictions are needed (often in the context of a Monte Carlo simulation model), this attitude cannot be routinely justified.

Since dependence is frequently ignored by modelers, even when it is clearly present, it is important to understand the modeling risk of inadequately capturing dependence or ignoring it altogether; we refer to this aspect as *undermodeling*. We note, in passing, that *overmodeling* is also undesirable; one is merely interested in the simplest possible model that would yield adequate predictions (the Principle of Parsimony). Queueing systems, so prevalent in computer and communications modeling, illustrate the risk of this kind of undermodeling. Consider a queueing system with bursty arrivals. Note that burstiness occurs because short interarrival intervals tend to follow in succession, until interrupted by a lull. Burstiness can be due either to the shape of the marginal interarrival distribution (say, due to a high coefficient of variation), or more commonly, because significant positive autocorrelations are present in the interarrival time process. A little reflection reveals that burstiness can make waiting times arbitrarily high, without increasing the arrival rate. For illustration, consider the effect of merely increasing the average number of customers that arrive in a single burst (think of batch arrivals), while spacing the bursts farther and farther apart on the average. Clearly, we can keep the arrival rate constant, this way, but the effect on mean waiting times will be disastrous: customers arriving in larger and larger bursts will experience increasing waiting times, while the longer and longer lulls separating bursts just waste the server's work potential. While this simple thought experiment should convince the reader that temporal dependence is important, the magnitude of its effect in queueing systems can be quantitatively startling. Indeed, various studies [5, 14, 18] have shown that when autocorrelated traffic is offered to a queueing system, the resulting performance measures are considerably worse than those corresponding to renewal traffic. In fact, mean waiting times can differ by orders of magnitude!

A growing realization of the impact of temporal dependence in traffic streams on queueing system performance has provided a prime motivation for devising input analysis methods that can capture such dependence. In general, prudent modelers ought to conclude that temporal dependence cannot be ignored! They should heed the cautionary admonition in the closing sentence of [21]: *"Would you rather be elegant and precisely wrong, or messy and vaguely right?"*

1.3 Goodness Criteria for Modeling Empirical Time Series

The TES modeling approach stipulates a number of requirements — precise requirements as well as heuristic ones — for the goodness of a candidate time series model based on empirical sample paths (see also [13, 12, 20] for related views):

Requirement 1: The marginal distribution of the model should match its empirical counterpart.

Requirement 2: The autocorrelation function of the model should approximate its empirical counterpart. Because the empirical data is finite, the model need only approximate the significant leading autocorrelations. For an empirical sample size of N, a rule of thumb for the maximal lag is the rounded value of $\sqrt{N}$; see [3], Chapter 5.

Requirement 3: Sample paths generated by a Monte Carlo simulation of the model should "resemble" the empirical data.

Note that these requirements are arranged in decreasing stringency. Requirements 1 and 2 constitute quantitative goodness-of-fit criteria, requiring that first-order and second-order properties of the empirical data be adequately captured. These are precise and well-defined requirements, although the particular metric of goodness is left up to the analyst. Requirement 3, however, is a heuristic qualitative requirement which cannot be defined with mathematical rigor. It is simply too subjective, and is better left to human cognitive judgment. Nevertheless, human cognition excells at pattern recognition; we are routinely called to judge for visual similarity, without having to be precise or even able to verbalize how we reached that judgment. Furthermore, the same judgment is often shared by different observers. Recall that it does not make sense to reproduce the empirical sample paths; one merely wishes to approximate their statistical signature. If a model can imitate the qualitative character of the empirical data, so much the better, as this would increase a practitioner's confidence in the model. It is important to realize that qualitative similarity *should not* substitute for the two preceding quantitative requirements. Rather, it is *in addition*, not *instead*. In short, Requirement 3 is merely the qualitative icing on the quantitative cake.

To get a more concrete idea regarding the three modeling requirements above, examine Figure 1. It displays a graphics screen produced by a TES modeling tool, called TEStool

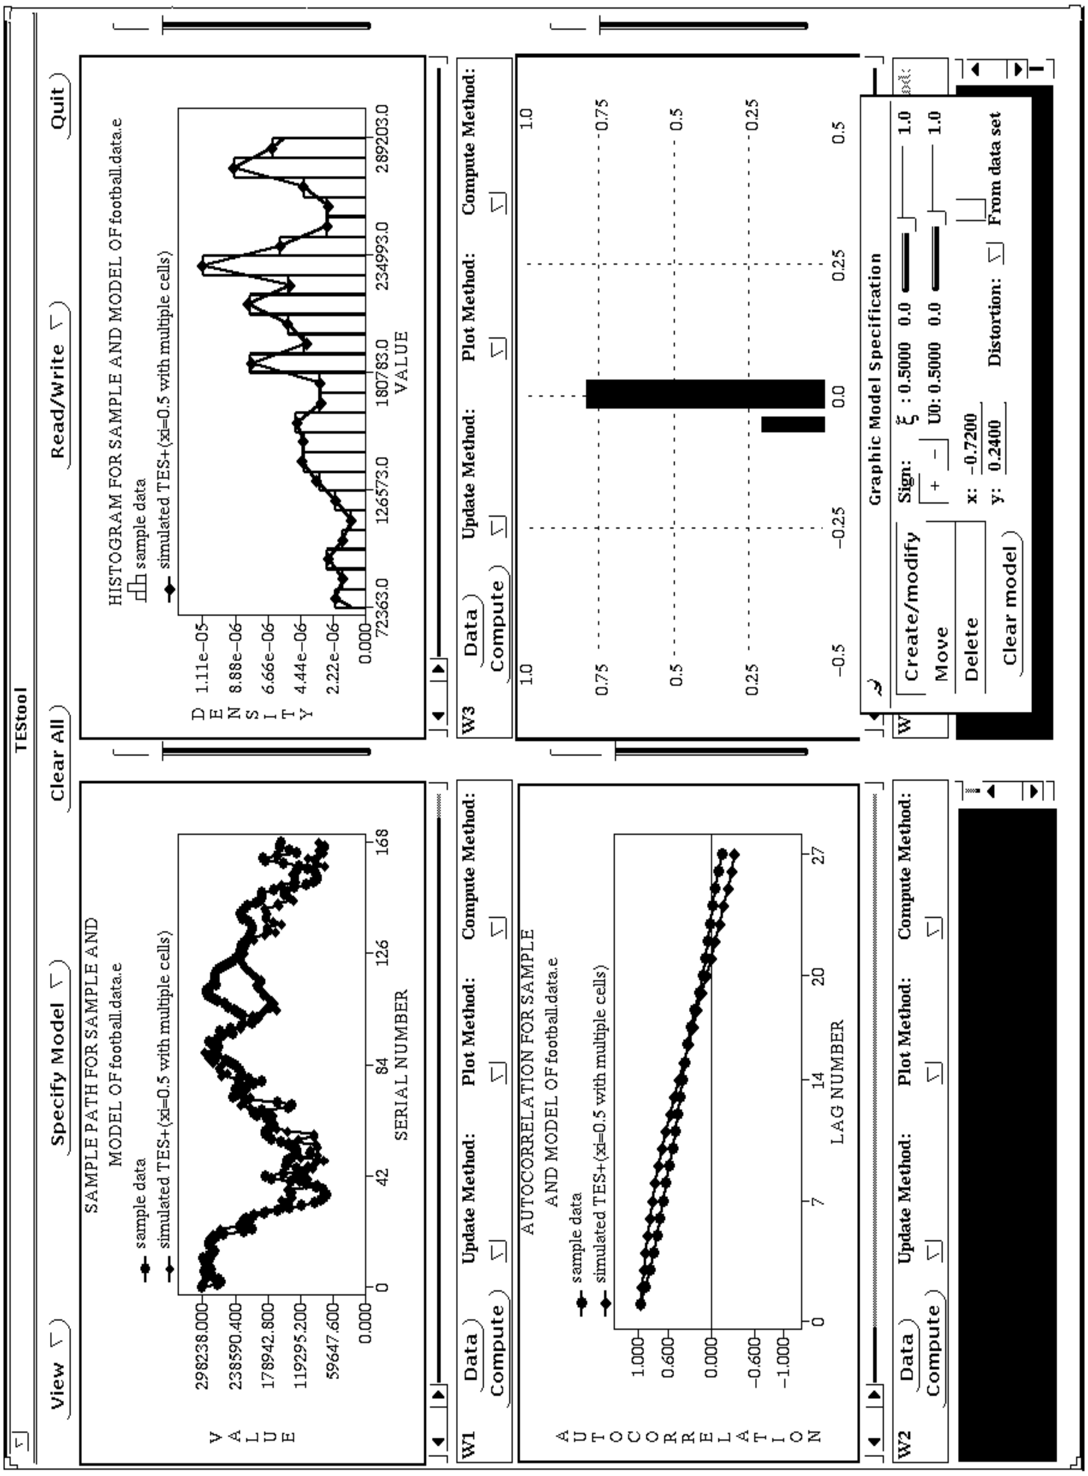

Figure 1: A TEStool Screen

[6, 16], to be described in Section 4.2. The screen consists of four tiled canvases (sub-windows). The lower-right canvas contains a visual specification of a TES model, to be explained in Section 4.2. The remaining three canvases display various statistics, illustrating the modeling requirements above. In each of them, the corresponding statistics are superimposed on each other for comparison; the empirical statistics are marked with bullets, and their TES counterparts by diamonds (see the legend at the top of each canvas). The upper-left canvas displays an empirical sample path against a TES model sample path generated by Monte Carlo simulation; the TES model was constructed from the emprical sample path data, so as to approximate the empirical statistics. The upper-right canvas displays two histograms, corresponding to the two sample paths in the upper-left canvas. Finally, the lower-left canvas displays the empirical autocorrelation function against its TES model counterpart; the latter was numerically computed from an analytical formula, to be presented in Section 3.3. Note, that Requirements 1 and 2 are apparently satisfied, as evidenced by the excellent agreement of the curves in the upper-right and lower-left canvases. It is also interesting to note that the upper-left canvas exhibits considerable "similarity" between corresponding sample paths, in apparent compliance with Requirement 3. Altogether, Figure 1 displays the results of a successful TES modeling effort, when judged against the three modeling requirements above.

1.4 The Merits of the TES Modeling Methodology

It appears that TES is the only modeling method that is designed to simultaneously meet the three goodness-of-fit requirements above in a systematic way. First, TES guarantees an exact fit to arbitrary marginals. More accurately, we are assured of an arbitrarily close fit, provided the simulation run of any of its sample paths is long enough. In particular, TES can match any empirical density (histogram). Second, TES possesses a large degree of freedom in approximating empirical autocorrelations, even as it maintains an exact match to the empirical marginal. TES autocorrelation functions have diverse functional forms including monotone, oscillatory, alternating and others. And third, TES processes span a wide qualitative range of sample paths, including cyclical as well as non-cyclical paths. Altogether, TES defines a very large class of models, encompassing both Markovian and non-Markovian processes.

TES processes enjoy two important computational advantages. To begin with, TES sequences are easily generated on a computer, and their periods are much longer than the underlying pseudo-random number stream. Their generation time complexity is small compared to that of the underlying pseudo-random number generator, and its space complexity is negligible. Furthermore, TES autocorrelations (and spectral densities) can be computed from accurate and fast (near real-time) analytical formulas without requiring simulation.

Because TES autocorrelations can be calculated in near real time on a modern workstation, TES modeling of empirical data can be carried out interactively, with guidance from visual feedback provided by a Graphical User Interface. This observation motivated the implementation of the TEStool modeling software package [6, 16] (recall Figure 1). TEStool makes heavy use of visualization in order to provide a pleasant interactive environment for TES modeling of general autocorrelated stationary time series. It supports visual modeling in the sense that the user can immediately see the statistics for each TES model (obtained by incremental modifications), superimposed on their empirical counterparts, and thereby judge the goodness of the current model. The environment speeds up the modeling search process, cuts down on modeling errors, and relieves the tedium of repetitive search. Modeling interactions are easy to grasp by experts and non-experts alike, since the search problem and search activities are cast in intuitive visual terms.

While TES is very versatile and its statistics exhibit rich behavior, the definition of TES processes is surprisingly simple. In essence, a TES process is a modulo-1 reduction of a simple linear autoregressive scheme, followed by additional transformations. The basic TES formulation is Markovian; however, the aforementioned transformations usually result in non-Markovian processes. Thus, TES is a non-linear autogressive scheme, encompassing Markovian and non-Markovian processes, which may explain in part its diversity and versatility.

1.5 Organization

This paper is a tutorial introduction to TES processes and the TES modeling methodology. As such it contains a certain amount of redundancy to facilitate the presentation.

The rest of the paper is organized as follows. Section 2 puts together the technical preliminaries required in the sequel. Section 3 contains a tutorial overview of TES processes. Section 4 explains the TES modeling methodology and overviews the TEStool modeling environment. Finally, Section 5 contains the conclusion of this paper.

2 TECHNICAL PRELIMINARIES

Several key technical concepts are essential to the understanding of TES processes and the TES modeling methodology. These will be intuitively explained in this section.

2.1 The Inversion Method

The *inversion method* is a standard technique of long standing for transforming a uniform random variable to one with an arbitrary prescribed distribution F. It has been used by Monte Carlo simulation analysts from its earliest days [2, 10].

For a random variable X and distribution F, let $X \sim F$ denote the fact that F is the distribution of X. In particular, $X \sim \text{Uniform}(c, d)$ means that X is uniformly distributed over the interval $[c, d)$.

Lemma 1 (Inversion Method) *Let F be any distribution function and $U \sim \text{Uniform}(0,1)$. Then the random variable $X = F^{-1}(U)$ satisfies $X \sim F$.*

For a proof, see *ibid.* Lemma 1 provides a very simple method of converting a marginally uniform sequence $\{U_n\}$ into a sequence $\{X_n\}$ with an arbitrary marginal distribution F. Simply set for each n,

$$X_n = F^{-1}(U_n). \tag{2.1}$$

A transformation of the form (2.1) is called an *inversion*. It always exists because F, being a cumulative distribution function (cdf), is non-decreasing, and therefore can always be inverted; although the inversion is not unique (unless F is strictly increasing), all choices of an inverse produce the same effect when applying the inversion method.

2.2 Histogram Inversions

In practical modeling, one usually estimates the empirical marginal density of an empirical sample by an empirical histogram statistic of the form $\hat{\mathcal{H}} = \{(l_j, r_j, \hat{p}_j) : 1 \leq j \leq J\}$, where J is the number of histogram cells, $[l_j, r_j)$ is the interval of cell j with width $w_j = r_j - l_j > 0$, and $\hat{p}_j$ is the probability estimator of cell j (the relative frequency of observations that fell into that cell). Note that hats signify that the corresponding quantity is a sample-based estimate, whereas the others are user-specified parameters. The empirical probability density function (pdf) is estimated as

$$\hat{h}(y) \;=\; \sum_{j=1}^{J} 1_{[l_j,r_j)}(y)\,\frac{\hat{p}_j}{w_j}, \quad -\infty < y < \infty, \tag{2.2}$$

where $1_A(x)$ denotes the indicator function of set A. Observe that, mathematically, Eq. (2.2) is just the density of a probabilistic mixture of uniform variates, where component j in the mixture is Uniform(l_j, r_j) and occurs with probability $\hat{p}_j$. The corresponding cumulative distribution function (cdf) is the piecewise linear function

$$\hat{H}(y) \;=\; \sum_{j=1}^{J} 1_{[l_j,r_j)}(y)\,[\hat{C}_{j-1} + (y - l_j)\frac{\hat{p}_j}{w_j}], \quad -\infty < y < \infty, \tag{2.3}$$

where $\{\hat{C}_i\}_{i=0}^{J}$ is the cdf of $\{p_j\}_{j=1}^{J}$, i.e., $\hat{C}_j = \sum_{i=1}^{j} \hat{p}_i,\; 1 \leq j \leq J$ ($\hat{C}_0 = 0$ and $\hat{C}_J = 1$).

The *histogram inversion* corrsponding to $\hat{H}$ is the piecewise linear function

$$D_H(x) \;=\; \hat{H}^{-1}(x) \;=\; \sum_{j=1}^{J} 1_{[\hat{C}_{j-1},\hat{C}_j)}(x)\,[l_j + (x - \hat{C}_{j-1})\,\frac{w_j}{\hat{p}_j}], \quad 0 \leq x \leq 1. \tag{2.4}$$

The TES modeling methodology uses histogram inversions of the form (2.4) to construct random variables with prescribed histogram densities (estimated from empirical data) via the inversion method. However, the TEStool environment admits other distributions as well.

Finally, a philosophical remark on the inversion approach is in order at this juncture. Analysts often try to fit an analytically-known density to empirical histograms. This approach has merit if it leads to tractable analytical models, or if it speeds up the inversion method on a computer. However, "smoothing" empirical histograms by analytic fitting may introduce inaccuracies instead of purportedly removing "noise" from the empirical data. Unless prior knowledge is availabele to justify this removal, it is best to stick with the empirical histogram. Not only do we save the labor of the fitting stage, but we work with a representation which contains more information than its fitted counterpart. Furthermore, the price paid for the fidelity of histogram inversions is low, since implementing the inversion method on a computer for a moderate-size histogram (20-25 cells) involves a very modest computational cost.

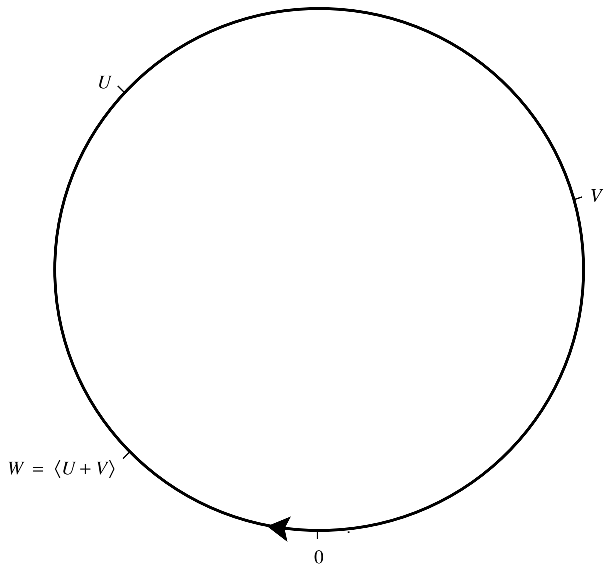

Figure 2: A Geometric Representation of Modulo-1 Addition

2.3 Modulo-1 Arithmetic

Modulo-1 arithmetic is simply the operation of taking the fractional part, after ordinary arithmetic is performed on real numbers. In order to extend the fractional part notion from its familiar form for positive numbers to negative ones, we need a bit of notation.

For any real x, let the *floor operator* $\lfloor\cdot\rfloor$ be defined as $\lfloor x\rfloor = \max\{\text{integer } n : \; n \le x\}$, i.e., the integral part of x. Note that the floor operator always truncates downwards. Thus, $\lfloor 1.2\rfloor = 1$, but $\lfloor -1.2\rfloor = -2$. The *modulo-1* (fractional part) operator $\langle\cdot\rangle$ can now be defined for any real x by $\langle x\rangle = x - \lfloor x\rfloor$. Thus, $\langle 1.2\rangle = 0.2$, but $\langle -1.2\rangle = 0.8$. Note, that a fractional part always lies in the interval $[0, 1)$, even for negative numbers.

An intuitive understanding of modulo-1 arithmetic can be gained from the geometric representation of Figure 2. Here, the interval $[0, 1)$ (the range of fractional parts) has been topologically transformed into a circle in a distance-preserving manner; the origin was arbitrarily selected at the bottom, and represents the fractional part 0. Henceforth, this circle will be referred to as the *unit circle* (note that it is the circumference which is unity, not the radius). A number $0 < y < 1$ is represented by a point on the unit circle at distance y from the origin (the distance is measured by the length of the arc connecting the origin clockwise to y). Two such numbers, U and V, are marked on the unit circle in Figure 2. Suppose $0 < n \le x < n+1$, for some positive integer n. To find $y = \langle x\rangle$ on the unit circle, perform the following thought experiment. Think of x as a virtual inelastic string of length x and zero width; the string can be bent, but cannot be stretched or compressed in length. Now, anchor one end of the string at the origin and proceed to overlay it precisely over the circle in a clockwise direction. You may have to wrap it around the circle several times (actually, $\lfloor x\rfloor = n$ times) before running out of slack, but wherever the string end comes to rest on the unit circle marks the required y. Pictorially, the wrapping action around the origin corresponds to discarding the integral part of x, so this procedure automatically yields the fractional part. For negative x, use a virtual string of length $|x|$, anchor at the origin as before, but overlay the unit circle in the counter-clockwise direction. It is now easy to see that for negative numbers $-1 < x < 0$, we have $\langle x \rangle = 1 - |x|$. Figure 2 illustrates the modulo-1 addition $\langle U + V \rangle$. First, sum $U + V$ in the ordinary way, and then perform the modulo-1 reduction as described before.

Just as modulo-1 arithmetic is an adaptation of ordinary arithmetic to the unit circle, the notion of *circular intervals* is an adaptation of ordinary intervals to the unit circle. The circular (modulo-1) interval $\langle a, b \rangle$ is defined for all $a, b \in [0, 1)$ by

$$\langle a, b \rangle = \begin{cases} [a, b), & a < b \\ [0, 1) \setminus [b, a), & b < a \\ \{a\}, & a = b \end{cases} \tag{2.5}$$

Intuitively, for $a \neq b$, the circular interval $\langle a, b \rangle$ consists of all the points on the unit circle from a and clockwise on the circle to b (including a but excluding b). If $a < b$, the result is the ordinary interval $[a, b)$; but if $b < a$, the resulting circular interval straddles point 0, and becomes in effect a union $[0, 1) \setminus [b, a) = [0, b) \bigcup [a, 1)$ of two intervals. This situation is depicted in Figure 2, with $a = V$ and $b = U$. The case $a = b$ is defined (somewhat inconsistently) as the singleton set rather than the empty set, for reasons of notational convenience.

2.4 Iterated Uniformity

Iterated uniformity refers to the ability to create sequences of variates, say $\{U_n\}_{n=0}^{\infty}$, where each U_n is distributed uniformly on $[0, 1)$. Modulo-1 arithmetic is intimately connected to iterated uniformity. Specifically, it can be shown that uniformity is closed under modulo-1 addition of independent random variables. While this result is not new (see, e.g., [4], p. 64), it plays such a central role in TES modeling that it merits rigorous quoting.

Lemma 2 (General Iterated Uniformity) *Let $U \sim$ Uniform$(0, 1)$, and let V be any random variable, independent of U. Define $W = \langle U + V \rangle$. Then $W \sim$ Uniform$(0, 1)$.*

Figure 2 illustrates how W is generated via modulo-1 addition on the unit circle, in accordance with Lemma 2. A plausibility argument goes as follows. Start by sampling U (independently) many times on the unit circle. Then any region of the unit circle is equally likely to have the same (uniform) density of points. Next, add a constant c modulo-1 to each point. This simply translates every point by the same amount on the unit circle either clockwise or counter-clockwise, depending on the sign of c. But the symmetry of the unit circle leaves the translated points still uniformly distributed on it. Since the argument holds for every c, it will also hold for any random variable V, as long as V is independent of U. A rigorous proof may be found in [7].

The significance of the Lemma stems from the fact that it gives us a simple prescription for generating a wide variety of marginally uniform sequences: Let $U_0 \sim$ Uniform$(0, 1)$, and let further $\{V_n\}_{n=1}^{\infty}$ be a sequence of iid (independent identically distributed) random variables with arbitrary marginal density f_V, and independent of U_0. The random variables V_n are referred to as *innovations*. Then the recursive scheme

$$U_n \;=\; \langle U_{n-1} + V_n \rangle, \quad n > 0, \tag{2.6}$$

is marginally uniform on $[0, 1)$ by Lemma 2. Note carefully the surprising fact that the distribution of the V_n is entirely irrelevant! Any sequence of mutually independent innovations (even non-stationary ones) will do, as long as its members are independent of U_0. In fact, the V_n may be random, deterministic and even assume negative values. Thus, the space of all innovations constitutes a huge degree of freedom for creating marginally uniform sequences! Notice carefully, that every choice of an innovation density f_V gives rise to a different *dependence structure* in $\{U_n\}$, as defined by Eq. (2.6). This flexibility is a major reason for the modeling versatility of TES, and will be revisited in Section 3.

2.5 Step-Function Innovation Densities

While the class of all innovation densities affords broad freedom in defining TES sequences, that class is simply too large to handle operationally. What is needed is a suitable restriction which attains sufficient manageability without sacrificing too much generality.

The class of *step-function* innovation densities on the unit circle (represented by the interval $[-0.5, 0.5)$ for reasons to be explained later) is a natural choice. First and foremost, step-function densities can approximate any density arbitrarily closely, so considerable generality is retained. The second reason is mathematical and conceptual simplicity. Step-function densities, being piecewise constant, correspond to probabilistic mixtures of uniform densities. Mathematically, these densities can be characterized by simply specifying their steps as triplet sets of the form $\{(L_k, R_k, P_k)\}_{k=1}^{K}$, where K is the number of steps, $[L_k, R_k)$ is the support of step k ($-0.5 \le L_k < R_k < 0.5$), and $0 < P_k \le 1$ is the mixing probability of step k ($\sum_{k=1}^{K} P_k = 1$). It is convenient to require that steps do not overlap (i.e., $R_k \le L_{k+1},\ k = 1, \ldots, K-1$). Since the density value over step k is fixed at $P_k/(R_k - L_k)$, we formally have

$$f_V(x) \;=\; \sum_{k=1}^{K} 1_{[L_k,R_k)}(x)\,\frac{P_k}{R_k - L_k}, \quad -0.5 \le x \le 0.5. \tag{2.7}$$

The conceptual simplicity of step functions is immensely valuable in facilitating their visual (graphic) specification and manipulation on a computer screen. This point will be explained in detail in Section 4.2.

Notice the similarity between Eqs. (2.7) and (2.2). In fact, both are step-function densities, but the former is theoretical and the latter is empirical. We use similar but distinct notation to emphasize this point.

2.6 Foreground / Background Schemes

The reader can now see how the previous subsections on modulo-1 arithmetic and iterated uniformity combine with Lemma 1 to provide a generic scheme for generating stationary sequences with arbitrary marginal distribution F.

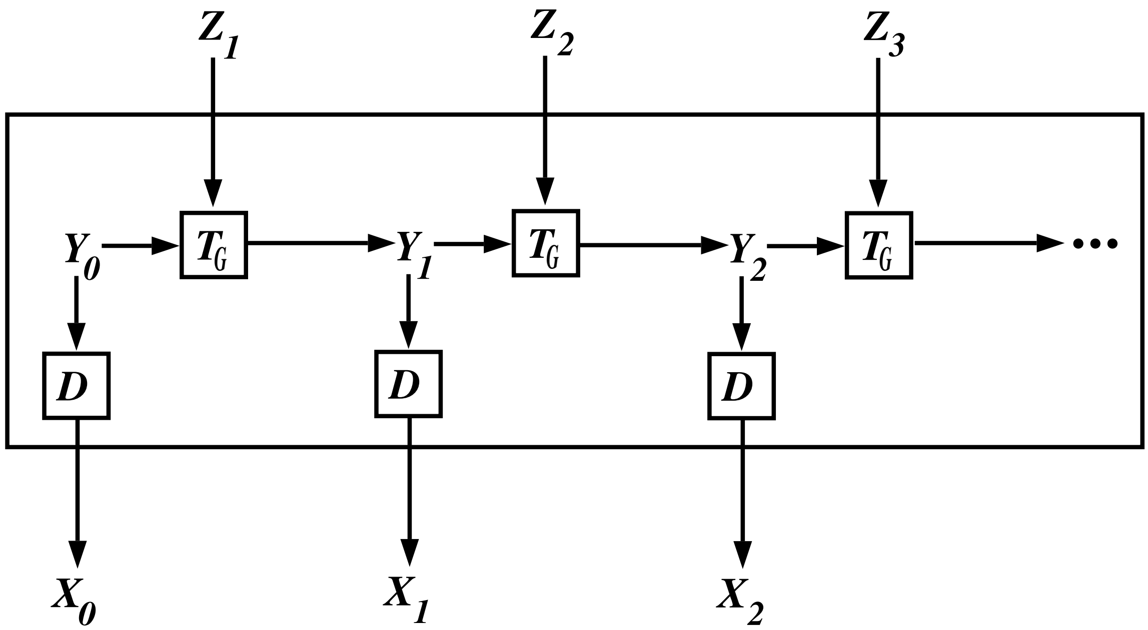

Figure 3: The Generic Foreground/Background Scheme

Figure 3 sketches the overall structure of such schemes, referred to as *foreground / background schemes*. The $\{Z_n\}$ sequence is a pseudo-random number stream (iid Uniform$(0, 1)$ variates), available on most computers. The sequence $\{Y_n\}$ is the *background sequence*, obtained as a general autoregression T_G (that is, Y_n is a function of Y_{n-1} and Z_n, where the latter plays the role of innovation), in such a way that the resulting sequence is stationary with some marginal distribution G. A deterministic transformation D, called *distortion*, transforms each Y_n to a corresponding X_n, and the resulting sequence $\{X_n\}$ is the *foreground sequence*. Intuitively, the background sequence is an auxiliary one which runs "unobserved in the background". Its foreground counterpart is the real target, since the goal is to endow it with the desired statistical properties (a prescribed marginal distribution F is one of them). The tranformation D has the curious name distortion, because in a sense, uniform variates are "ideal" (being so very simple...), and any deviation from uniformity constitutes a "distortion". This connotation not withstanding, distortions are extremely useful, operationally.

As we shall later see, TES generation methods fall within the class of foreground / uniform background schemes.

2.7 Stitching Transformations

Imagine a random walk on the unit circle as depicted in Figure 2. Such random walks will have inherent sample path "discontinuities" due to modulo-1 "crossings" of point 0 in the clockwise or counter-clockwise direction. The term "discontinuities" is used here figuratively; it actually refers to transitions from large fractions to small ones and vice versa, when the random walk wraps around point 0 modulo-1 in each direction.

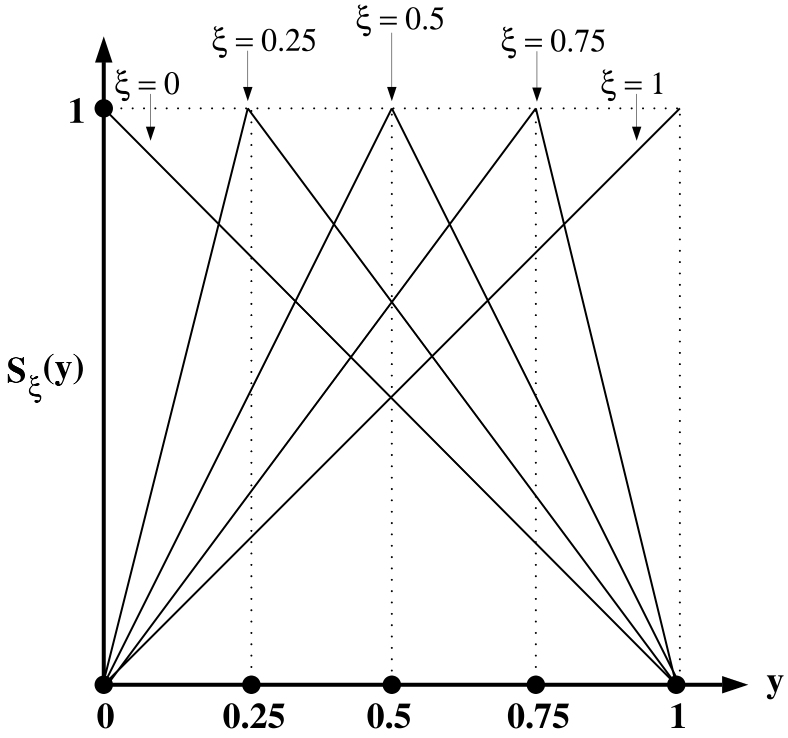

Figure 4: Select Stitching Transformations

In some select situations, such sample path "discontinuities" could be endowed with a valid modeling interpretation. For example, crossing point 0, properly defined, could model a catastrophe — perhaps in a cyclical economic model. But in most cases, such "discontinuities" in the sample path are undesirable, as they have no qualitative counterparts in empirical data. Their presence may violate Requirement 3 in Section 1.3, even though Requirements 1 and 2 may well be complied with. What is needed are "smoothing" transformations to achieve compliance with the former, but without sacrificing compliance with the latter. Most desirable are "smoothing" transformations which are uniformity-preserving, in order to fit arbitrary marginal distributions via the inversion method. The so-called *stitching transformations* [7] precisely fit the bill.

Formally, a stitching transformation S_ξ maps the interval $[0, 1)$ to itself, and is determined by a *stitching parameter* ξ in the range $0 \leq \xi \leq 1$. For a given ξ, S_ξ is defined by

$$S_\xi(y) \;=\; \begin{cases} y/\xi, & 0 \leq y \leq \xi \\ (1-y)/(1-\xi), & \xi \leq y < 1 \end{cases} \tag{2.8}$$

Figure 4 depicts select stitching transformations plotted for various values of ξ. The sobriquet "stitching" is motivated by the fact that for $0 < \xi < 1$, S_ξ is continuous on the unit circle. To see that, notice that all S_ξ are continuous in the ordinary sense in the interior of the unit interval. However, as the unit interval is closed into a unit circle by "stitching" together the edge points 0 and 1, the images of S_ξ at the edge points will be similarly stitched, in the sense that $S_\xi(0) = S_\xi(1) = S_\xi(1-)$, $0 < \xi < 1$. Now imagine that a highly autocorrelated, marginally uniform sequence $\{U_n\}$, is represented by points on the abcissa of the graph in Figure 4. Intuitively, this means that successive points of $\{U_n\}$ are not too far apart on the abcissa, except when they wrap around 0. But even then, S_ξ ensures that for $0 < \xi < 1$, successive images of $\{S_\xi(U_n)\}$ are similarly spaced on the ordinate. Thus, $\{S_\xi(U_n)\}$ is "smoother" than the original sequence $\{U_n\}$. Note that no "smoothing" takes place for $S_1(x) = x$ (the identity) and $S_0(x) = 1 - x$ (the antithetic transformation); indeed, Figure 4 shows that these are not "continuous" at the edges of the unit circle.

The primary utility of stitching transformations is that they all preserve uniformity, in addition to their "smoothing" effect.

Lemma 3 *If* $U \sim$ *Uniform*$(0,1)$*, then* $S_\xi(U) \sim$ *Uniform*$(0,1)$*, for all* $0 \le \xi \le 1$.

A simple proof may be found in [15]. It follows that a compound distortion of the form

$$D(x) \;=\; F^{-1}(S_\xi(x)), \quad x \in [0,1], \tag{2.9}$$

employing stitching and inversion in succession, will allow us to to fit arbitrary marginal distributions, and simultaneously attain sample path "smoothing". The empirical TES modeling methodology, to be described in Section 4.1, employs compound distortions of the form (2.9).

An example of the "smoothing" effect of stitching on TES sequences is deferred until Section 3.3.

2.8 The TES Modeling Approach

We can now gather the strands of the foregoing discussion and sketch the main elements of the TES modeling approach.

Consider again Figure 3, and make the following specializations. Let $Y_0 = U_0$ be uniform on $[0,1)$, and let T_G implement the iterated uniformity scheme (2.6), where each V_n is a function of the corresponding Z_n (the V_n could be possibly generated via the inversion method). The background sequence $\{Y_n\} = \{U_n\}$ will then be marginally uniform on $[0,1)$, and so each U_n could be transformed by a distortion D to yield a foreground sequence $\{X_n\}$ with prescribed marginal distribution F. To this end, one may use an inversion distortion or a compound one of the form (2.9).

The TEStool modeling environment makes two additional specializations. It uses histogram inversions $\hat{H}^{-1}$ of the form (2.4) in constructing compound distortions of the form (2.9), and it restricts the range of innovation densities to the class of step functions (2.7) whose support is contained in the interval $[-0.5, 0.5)$. A detailed description of TES sequences will be presented in Section 3.

Note carefully, again, that for any distortion $D = F^{-1}$, we have $X_n \sim F$ *regardless* of the innovation sequence selected! Thus, the TES approach guarantees that Requirement 1 in Section 1.3 is always satisfied. How then should one approximate a prescribed autocorrelation function? Here, we take advantage of the fact, that the choice of an innovation density, completely determines the temporal dependence structure of the corresponding background sequence (for the foreground sequence, the distortion also participates in this determination). A choice of innovations (and to a lesser extent, of a stitching parameter) can then allow us to satisfy Requirement 2 in Section 1.3. Experience shows that these choices also have a strong impact on the qualitative nature of TES Monte Carlo sample paths, in line with Requirement 3 in Section 1.3.

All in all, the TES approach allows us to decompose the selection of the marginal distribution and autocorrelation function, rendering them largely orthogonal choices. Therein lie both the novelty of TES and its modeling power.

3 OVERVIEW OF TES SEQUENCES

Throughout the remainder of this paper we assume the following setup, over a common probability space:

1. $U_0 \sim$ Uniform$(0, 1)$ is the initial variate in a background TES sequence.

2. $\{V_n\}_{n=1}^{\infty}$ is a sequence of innovations (iid random variables) with common density f_V, independent of U_0. Without loss of generality, we assume that the support of f_V is the interval $[-0.5, 0.5)$, representing the unit circle. In fact, any interval of length 1 can serve as an equivalent support for f_V, because of the modulo-1 arithmetic involved. The reason for this particular choice will become evident later.

3. D is a distortion (a measurable function from $[0, 1]$ to the real line). As in Eq. (2.9), we assume that D is composed of a stitching transformation S_ξ (for some stitching parameter $\xi \in [0, 1]$), followed by an inversion F_X^{-1} of a distribution function F_X.

These are the determining parameters for a TES process.

3.1 Basic Background TES Sequences

We start out by introducing two classes of the simplest TES background sequences: TES$^+(L, R)$ and TES$^-(L, R)$. Collectively called *basic* TES sequences, these were first studied in [15], and their construction motivated the TES (*Transform-Expand-Sample*) acronym. The superscripts (plus and minus) are suggestive of the fact that basic TES sequences cover all lag-1 autocorrelation; TES$^+$ sequences cover the positive range $[0, 1]$, while TES$^-$ sequences cover the negative range $[-1, 0]$ (see below). Hereafter, we shall consistently append the proper superscript (plus or minus) to distinguish between similar mathematical objects associated with TES$^+$ and TES$^-$, respectively; the superscript is deleted in objects common to TES$^+$ and TES$^-$, or those for which the distinction is immaterial. For example, TES$^+$ and TES$^-$ sequences are denoted $\{U_n^+\}$ and $\{U_n^-\}$, respectively, but U_0 is the common initial variate for all TES classes.

Basic TES sequences are parameterized by pairs (L, R), $-0.5 \le L < R < 0.5$. These are used, in fact, to parameterize uniform innovation variates of the form

$$V_n \;=\; L + (R - L) Z_n, \tag{3.1}$$

where $\{Z_n\}$ is a sequence of iid Uniform$(0, 1)$ variates. The (L, R) parameterization is equivalent to the (α, ϕ) parameterization given by

$$\alpha \;=\; R - L, \quad \phi \;=\; \frac{R + L}{\alpha}, \tag{3.2}$$

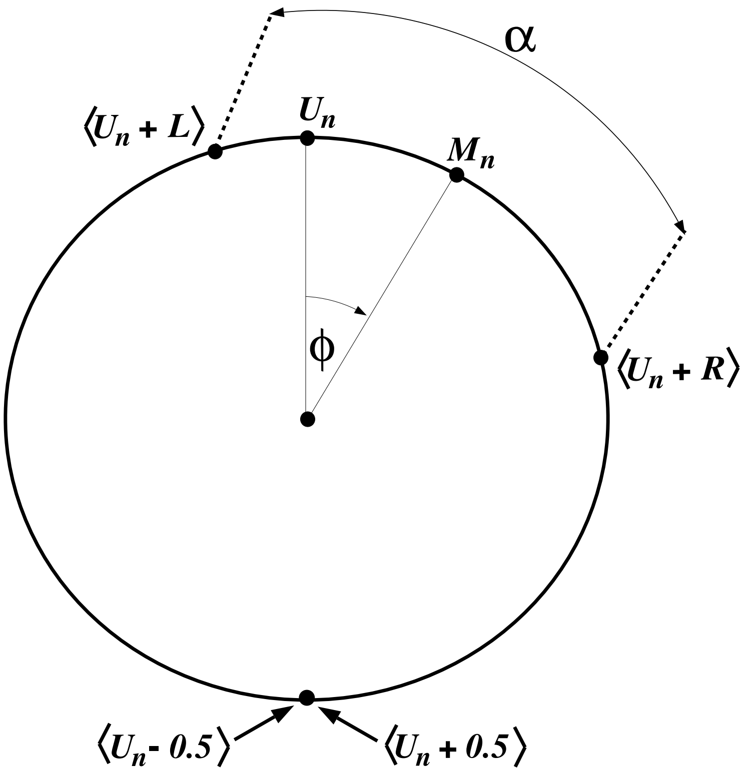

Figure 5: Geometric Interpretation of Basic Background TES Sequences

and the corresponding TES classes are denoted TES$^+(\alpha, \phi)$ and TES$^-(\alpha, \phi)$, respectively. As will be presently seen, the (α, ϕ) parameterization is analoguous to polar coordinates, whereas its (L, R) counterpart is analoguous to Cartesian coordinates. More importantly, the (α, ϕ) parameterization simplifies the representation of the autocorrelation function of TES sequences [7].

Figure 5 ascribes a simple geometric interpretation to TES$^+(\alpha, \phi)$. Here, $\langle U_n + L\rangle$ and $\langle U_n + R\rangle$ are translates of the current TES iterate $U_n = U_n^+$ by offsets $L < 0$ and $R > 0$ on the unit circle; $\alpha = R - L$ is the length of the circular interval $\langle U_n + L, U_n + R\rangle$; and ϕ is an indication of the rotation of that interval away from symmetric straddle of U_n (the angle between U_n and the interval midpoint M_n). The analogy of the (α, ϕ) parameterization to polar coordinates is now evident.

The construction of basic TES sequences is recursive, each resulting in a stationary Markovian sequence whose marginal distribution is uniform on $[0, 1)$; see [15] for formulas of the respective transition densities. Start with $U_0 \sim$ Uniform$(0, 1)$ as the recursion basis. Next, assume that $U_n \sim$ Uniform$(0, 1)$ has already been constructed in the recursion step. The next TES iterate U_{n+1} is constructed as follows (refer to Figure 5).

For TES$^+(\alpha, \phi)$, first construct the circular interval $\langle U_n + L, U_n + R\rangle$ about $U_n = U_n^+$, (the *Expand* stage of TES), and then sample U_{n+1} uniformly in $\langle U_n + L, U_n + R\rangle$, independently of the past (the *Sample* stage of TES). The *Transform* stage in TES$^+$ is the identity. Notice that this is just the geometrical interpretation of the iterated uniformity relation $U_{n+1}^+ = \langle U_n^+ + V_n\rangle$ from Eq. (2.6), for V_n of the form (3.1).

For TES$^-(\alpha, \phi)$, the recursive construction of $U_{n+1} = U^-_{n+1}$ is a bit more involved (*ibid.*) The *Transform* stage maps U_n to its antithetic variate $1 - U_n$. The *Expand* stage is step dependent: for odd n construct $\langle U_n + L, U_n + R \rangle$, but for even n construct $\langle 1 - U_n + R, 1 - U_n + L \rangle$. The *Sample* stage to obtain $U_{n+1} = U^-_{n+1}$ is the same as in TES$^+(\alpha, \phi)$. A little reflection should reveal that this construction for TES$^-$ is equivalent to $U^-_{n+1} = \langle U^-_n + V_n \rangle$ for odd n , and $U^-_{n+1} = \langle 1 - U^-_n - V_n \rangle$ for even n.

The qualitative behavior of basic TES sequences as a function of α and ϕ can be readily grasped with the aid of Figure 5. As $\alpha \downarrow 0$, the limiting circular intervals $\langle U_n + L, U_n + R \rangle$ shrink to a singleton set by (2.5). Consequently, $\{U^+_n\}$ approaches $\{U_0, U_0, \ldots\}$ with $\rho^+_U(\tau) \equiv 1$, while $\{U^-_n\}$ approaches $\{U_0, 1-U_0, U_0, 1-U_0, \ldots\}$, with $\rho^-_U(\tau) = -1,\ \tau \geq 1$. As $\alpha \uparrow 1$, the limiting circular intervals expand to $[0, 1)$ by (2.5), and both $\{U^+_n\}$ and $\{U^-_n\}$ approach iid Uniform(0,1) variates, so $\rho^+_U(\tau) \equiv \rho^-_U(\tau) \equiv 0$. In between these extremal cases lie intermediate values $\alpha \in (0, 1)$ which give rise to autocorrelated TES sequences, such that $\rho^+_U(1)$ covers $(0, 1)$ and $\rho^-_U(1)$ covers $(-1, 0)$, as L and R range over $[-0.5, 0.5)$; see [15] for details.

Consider now a TES$^+(\alpha, \phi)$ sequence, and refer to Figure 5. If $\phi = 0$, then $R = -L$, so that the next iterate U^+_{n+1} is equally likely to fall to the left or to the right of the current iterate $U_n = U^+_n$. In this case, $E[V_n] = 0$, and $\{U^+_n\}$ is a random walk with zero net drift around the unit circle. If $\phi > 0$, then $R > |L|$, so U^+_{n+1} is more likely to fall to the right than to the left of $U_n = U^+_n$. In this case, $E[V_n] > 0$, and $\{U^+_n\}$ is a random walk with positive (clockwise) net drift around the unit circle. As we shall see later on, this results in sawtooth-shaped cyclical sample paths. The case $\phi < 0$ is analogous but with counter-clockwise net drift, resulting from $E[V_n] < 0$. Clearly, the innovation variates V_n serve as differences between successive iterates U_n and U_{n+1}, and the average difference determines whether the TES sequence is driftless or directional as well as the drift direction. However, note carefully that these are no ordinary differences; rather, these are *differences on the unit circle*. In fact, one cannot deduce the innovation variates V_n from $\{U^+_n\}$ in a unique way, because the rules of modulo-1 arithmetic allow us to write $U^+_{n+1} = \langle U^+_n + V_n \rangle$ as well as $U^+_{n+1} = \langle U^+_n + V_n - 1 \rangle$. It follows that both V_n and $V_n - 1$ qualify as innovation variates. This is just a reflection of the fact that to get from one point to another on a circle, one can either go to the right or to the left with equal effect.

Although basic TES sequences have just two parameters, they already possess a considerable range of sample path behavior, and a variety of functional forms of associated autocorrelation functions. This variety is demonstrated in Figures 6–13 which depict four fundamental behavioral modes of basic TES(α, ϕ) background sequences.

An examination of the sample paths (Figures 6, 8, 10 and 12) reveals a qualitative difference between TES$^+$ and TES$^-$ sample paths. The latter are characterized by a visual "shadow" appearance (Figures 10, 12), due to the alternation of large and small values, which in turn gives rise to an alternating autocorrelation function with negative lag-1 autocorrelation. Another fundamental qualitative difference can be discerned between sequences without drift ($\phi = 0$), as in Figures 6 and 10, and those with drift ($\phi = 0.5$), as in Figures 8 and 12. The latter exhibit a distinct cyclical structure resulting from periodic drifting across the "origin" on the unit circle. This cyclical structure is absent for $\phi = 0$, because the resulting random walk is not directional. Subsequently, the behavior of the corresponding autocorrelation functions (Figures 7, 9, 11 and 13) displays four functional

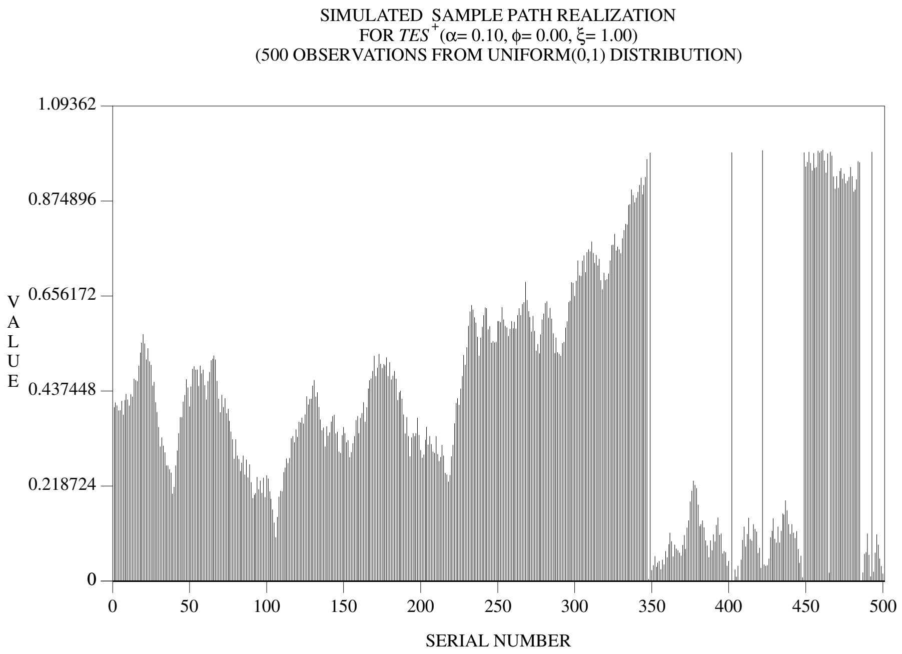

Figure 6: Sample Path of a Uniform TES$^+(\alpha, \phi)$ Process Without Drift

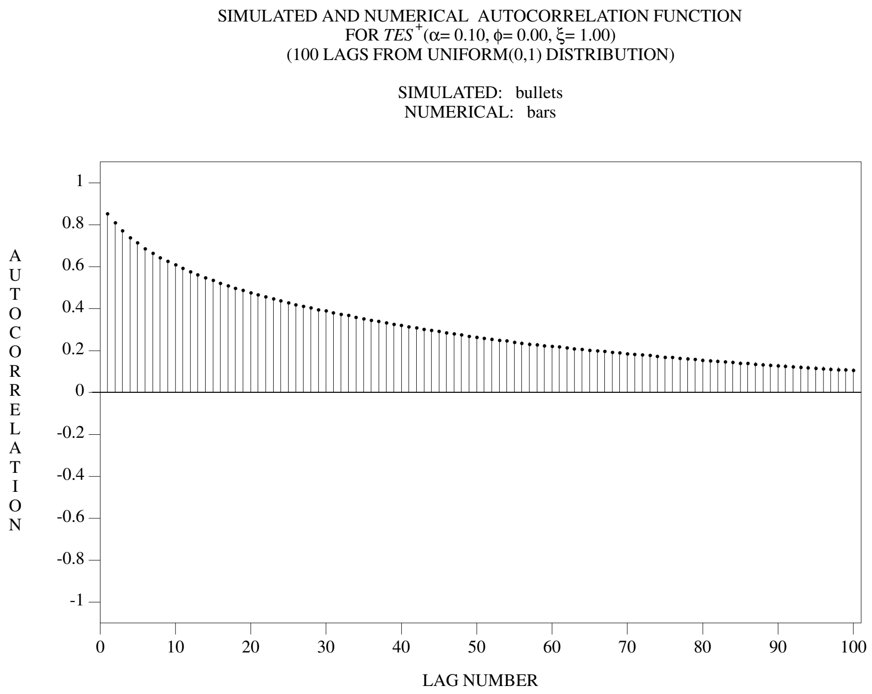

Figure 7: Autocorrelation Functions of a Uniform TES$^+$ Process Without Drift

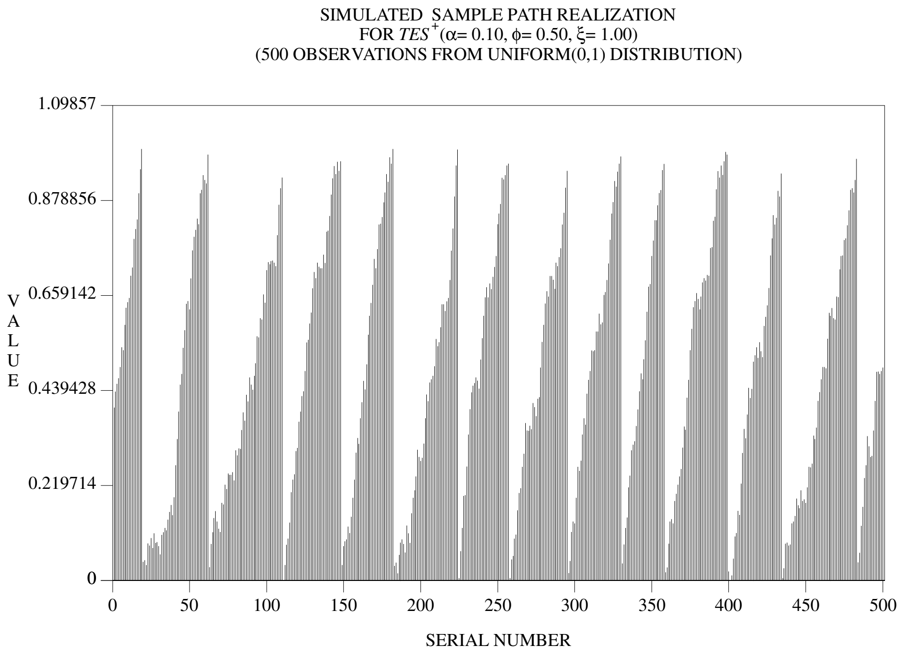

Figure 8: Sample Path of a Uniform TES$^+(\alpha, \phi)$ Process With Drift

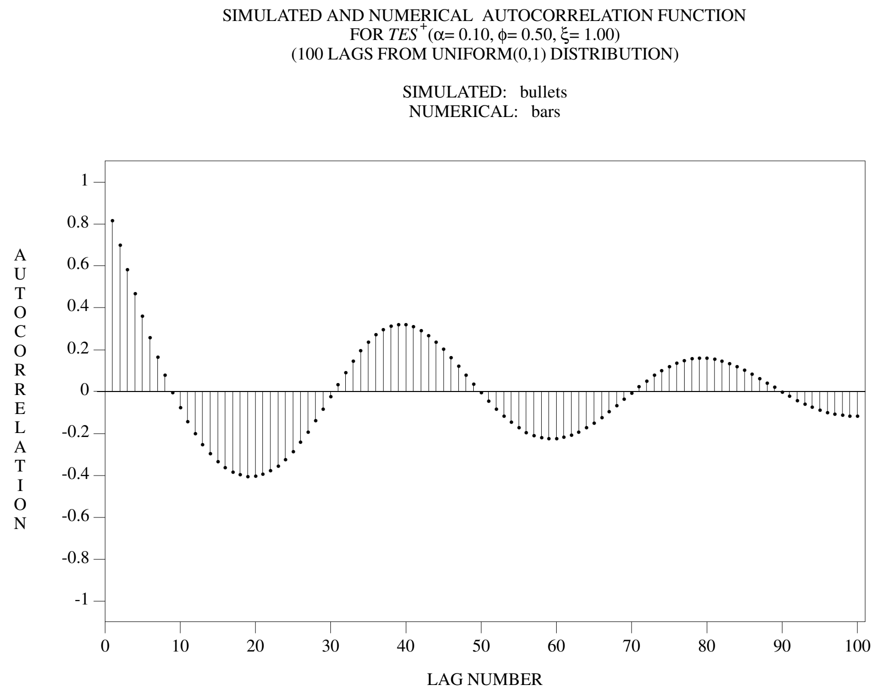

Figure 9: Autocorrelation Functions of a Uniform TES$^+(\alpha, \phi)$ Process With Drift

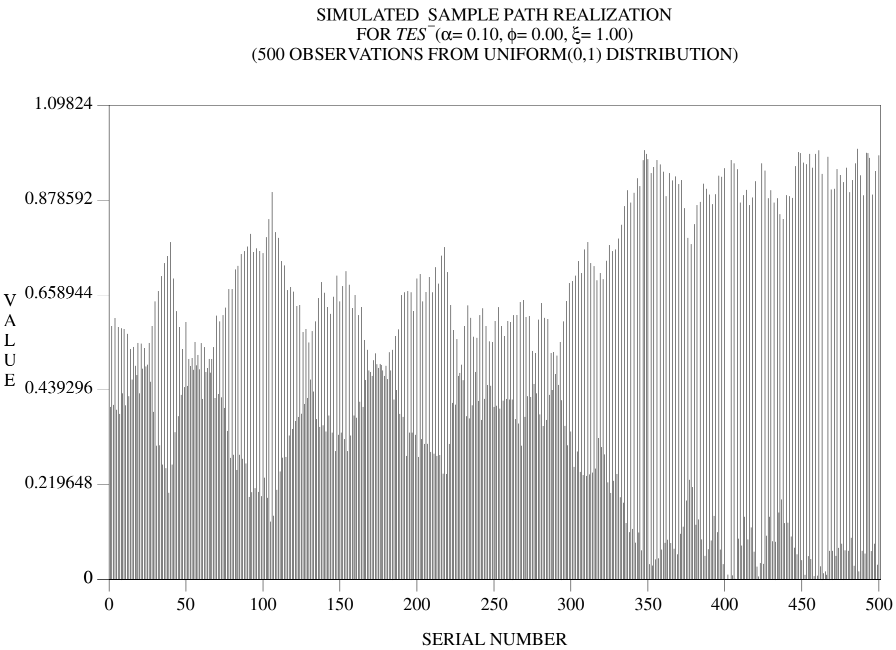

Figure 10: Sample Path of a Uniform TES$^-(\alpha, \phi)$ Process Without Drift

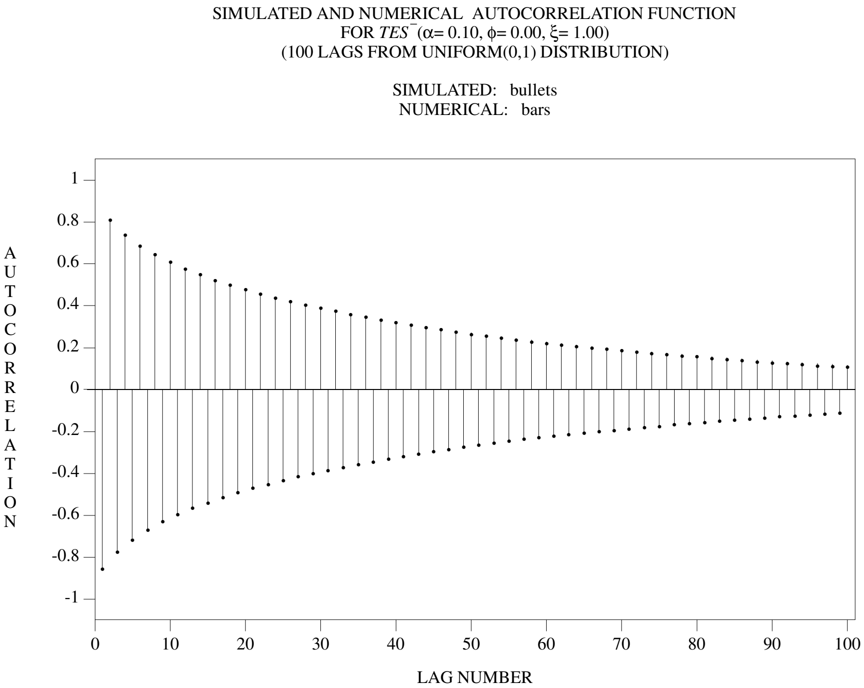

Figure 11: Autocorrelation Functions of a Uniform TES$^-(\alpha, \phi)$ Process Without Drift

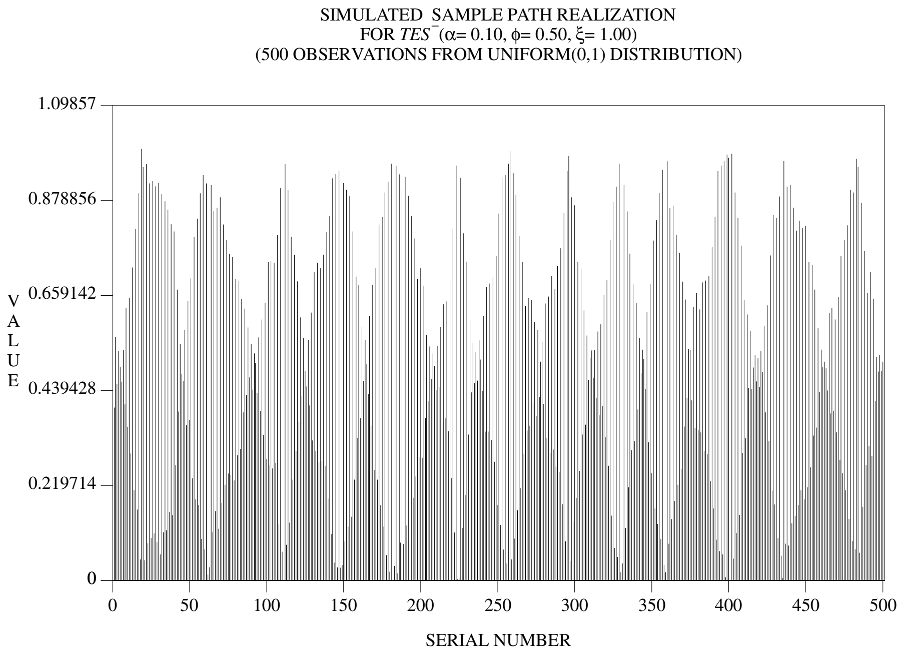

Figure 12: Sample Path of a Uniform TES$^-(\alpha, \phi)$ Process With Drift

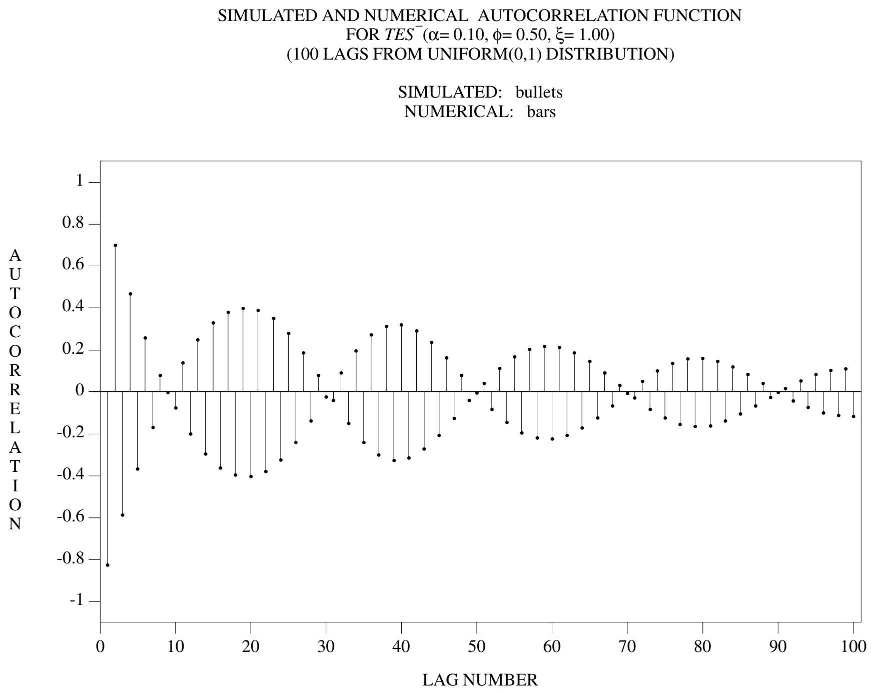

Figure 13: Autocorrelation Functions of a Uniform TES$^-(\alpha, \phi)$ Process With Drift

forms. The driftless TES$^+$ case in Figure 7 is monotone decreasing, whereas the driftless TES$^-$ case in Figure 11 is alternating with monotone decreasing envelope. The TES$^+$ case with drift in Figure 9 is oscillatory, whereas its TES$^-$ counterpart in Figure 13 is simultaneously alternating and oscillatory. All autocorrelation functions are bounded by envelopes with monotone decreasing magnitude. The effect of increasing the magnitude of ϕ on sample path behavior is to decrease the mean "cycle time" of the directional random walk around the unit circle, with the corresponding autocorrelation function oscillating more rapidly. As expected, an increase in α causes a decrease in the magnitude of the resultant autocorrelation functions for all lags.

When the full generality of TES sequences (general innovations and distortions) is invoked, the repertoire of TES sequences and their behavioral range is enhanced far more. This flexibility explains in a nutshell why one can hope a priori to define TES sequences which simultaneously approximate first-order and second-order properties of empirical sequences.

3.2 General TES Background Sequences

This section generalizes the scope of innovations of background TES sequences from uniform innovation densities to general ones.

General TES background sequences fall into two classes, denoted TES$^+$ and TES$^-$. Class TES$^+$ consists of random sequences $\{U_n^+\}$ of the form

$$U_n^+ \;=\; \begin{cases} U_0, & n = 0 \\ \langle U_{n-1}^+ + V_n \rangle, & n > 0 \end{cases} \tag{3.3}$$

and class TES$^-$ consists of random sequences $\{U_n^-\}$ of the form

$$U_n^- \;=\; \begin{cases} U_n^+, & n \text{ even} \\ 1 - U_n^+, & n \text{ odd} \end{cases} \tag{3.4}$$

It immediately follows that both $\{U_n^+\}$ and $\{U_n^-\}$ are Markovian sequences. Their transition densities were derived in [7], in terms of Fourier expansions, but these are beyond the scope of this paper and will not be reproduced here. By iterated uniformity (see Section 2.4), it further follows that both TES$^+$ and TES$^-$ processes are stationary sequences with Uniform$(0, 1)$ marginals. Eq. (3.4) readily implies that for any given innovation sequence $\{V_n\}$, the autocorrelation functions of the corresponding $\{U_n^+\}$ and $\{U_n^-\}$ sequences satisfy the relation

$$\rho_U^-(\tau) \;=\; (-1)^\tau \rho_U^+(\tau), \quad \tau = 0, 1, \ldots \tag{3.5}$$

It is clear from Eq. (3.5) that if TES$^+$ sequences cover the positive range $[0, 1]$ of lag-1 autocorrelations, then TES$^-$ sequences cover the negative range $[-1, 0]$.

Figure 14 provides a 3-dimensional geometric interpretation of the construction of background TES sequences with step-function innovation densities. Here the unit circle lies at the bottom, in the 2-dimensional plane, with the current TES variate U_n in the north/north-easterly sector. A step-function innovation density f_V was erected over the unit circle, with values in the third dimension (the "up" direction). The origin for the

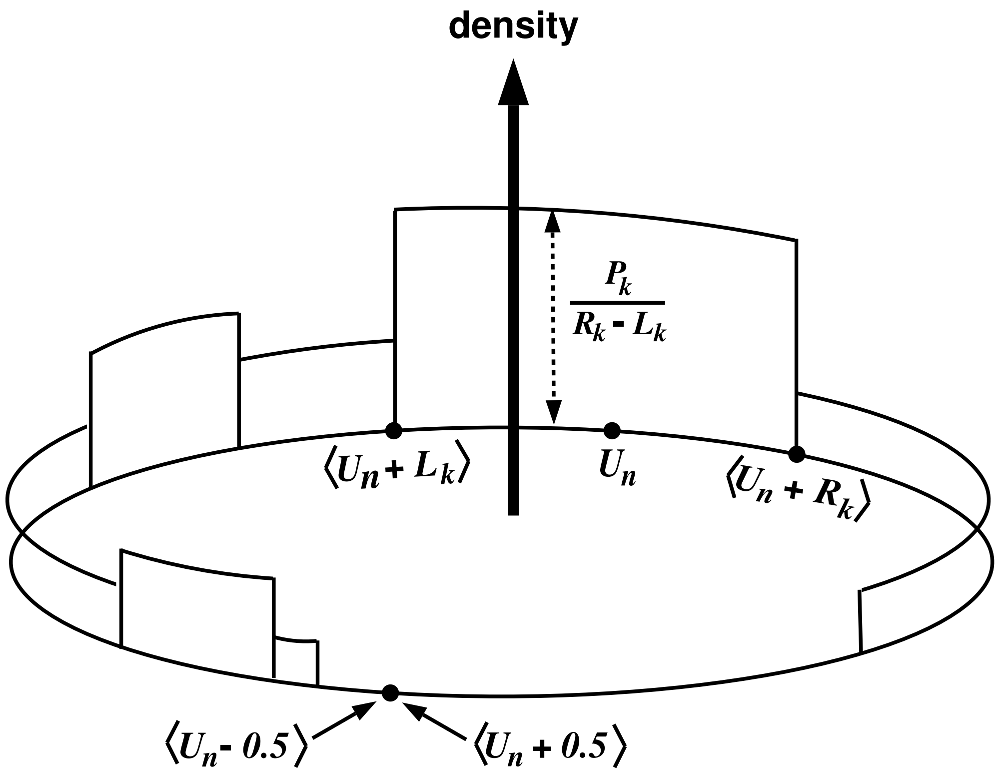

Figure 14: Geometric Interpretation of Background TES Sequences With Step-Function Innovations

support of f_V was set at the current TES variate U_n, implying that the support of f_V is the interval $[-0.5, 0.5)$. To obtain the next TES$^+$ variate U_{n+1}, just sample a value on the unit circle from the innovation density f_V. Indeed, because the origin of f_V was set at U_n, this procedure implies $U_{n+1} = \langle U_n + V_n \rangle$, in agreement with Eq. (3.3). To map a TES$^-$ method into a geometrical interpretation in Figure 14, simply use alternately f_V for even indices n, and its antithetic counterpart (corresponding to sampling $U^-_{n+1} = 1 - U^+_n$) for odd indices n, as prescribed by Eq. (3.4).

Figure 14 clarifies the reason for selecting the interval $[-0.5, 0.5)$ as the support set of f_V. Note carefully, that any interval of lenght 1 will do, due to the modulo-1 reduction perfomed in Eqs. (3.3)–(3.4). But Eq. (3.3) implies (and Figure 14 illustrates graphically) that we can think about general innovation variates V_n as *modulo-1 differences* between successive TES$^+$ variates, that is, *differences on the unit circle*. This interpretation is only valid for the choice of the natural support $[-0.5, 0.5)$ for f_V.

Although Figure 14 does not exhaust the full generality of TES processes, it embodies the level of generality of background TES sequences, adopted by the TES modeling methodology, and implemented in TEStool. For all practical purposes, we may, henceforth, adopt Figure 14 as the mental image characterizing the construction of general background TES sequences.

3.3 General TES Foreground Sequences

Recall that the background TES sequences, $\{U_n^+\}$ and $\{U_n^-\}$ from Eqs. (3.3)–(3.4), play an auxiliary, albeit important, role in TES modeling. The real modeling interest lies in general foreground TES sequences $\{X_n^+\}$ and $\{X_n^-\}$, obtained from Eq. (3.3) and Eq. (3.4), respectively, via a general distortion, as

$$
\begin{aligned}
X_n^+ &= D(U_n^+), \qquad (3.6)\\
X_n^- &= D(U_n^-). \qquad (3.7)
\end{aligned}
$$

Figures 15 and 16 summarize the generation scheme of TES$^+$ and TES$^-$ sequences, respectively, in a self-explanatory manner.

Recall that a compound distortion D of the form (2.9) serves a dual goal: It "smoothes" sample paths of TES sequences, and it guarantees their marginal distribution to match, in principle, any prescribed one. We still need, though, to approximate a prescribed autocorrelation function; consequently, we need a way to calculate or estimate TES autocorrelations.

Naturally, a Monte Carlo simulation of the foreground TES sequences (3.6) and (3.7) can always provide a good estimate of the corresponding autocorrelation function, provided a sufficient sample size is generated. Unfortunately, this approach can be costly in terms of time complexity, especially when high autocorrelations necessitate large sample sizes for adequate statistical reliability (it takes some 10-15 minutes of elapsed time to estimate 100 lags of the autocorrelation function for a million-observation sample on a standard workstation). This pretty much precludes the use of simulation-based estimators in interactive heuristic searches for suitable TES models. Fortunately, fast and accurate numerical algorithms are available for general foreground TES processes $\{X_n^+\}$ and $\{X_n^-\}$, with general distortions and innovations.

Let $\tilde{f}(s) = \int_{-\infty}^{\infty} \exp(-sx)\, f(x)\, dx$ denote the Laplace Transform of a function $f(x)$. For a given lag τ, the corresponding autocorrelation functions are given, respectively, by

$$\rho_X^+(\tau) = \frac{2}{\sigma_X^2} \sum_{\nu=1}^{\infty} Re[\tilde{f}_V^\tau(i2\pi\nu)]\, |\tilde{D}(i2\pi\nu)|^2 \qquad (3.8)$$

and

$$\rho_X^-(\tau) = \begin{cases} \rho_X^+(\tau), & \tau \text{ even} \\ \dfrac{2}{\sigma_X^2} \displaystyle\sum_{\nu=1}^{\infty} Re[\tilde{f}_V^\tau(i2\pi\nu)]\, Re[\tilde{D}(i2\pi\nu)^2], & \tau \text{ odd} \end{cases} \qquad (3.9)$$

See [7] for the mathematical derivations. These formulas are given in terms of the Fourier coefficients $\tilde{f}_V(i2\pi\nu)$ and $\tilde{D}(i2\pi\nu)$, and so the effect of innovations and distortions on the autocorrelation function is conveniently separated. On the other hand, these infinite expansions do not constitute closed-form representations. In order to gain insight into the effect of innovations and distortions on the corresponding autocorrelation functions, one must perform a numerical summation for a sufficiently large number of terms. Fortunately, Eqs. (3.8) and (3.9) are eminently computable, and the sums converge rather rapidly for all popular and useful innovations and distortions, including compound distortions of the form (2.9) and step-function innovation densities of the form (2.7), both of which are utilized

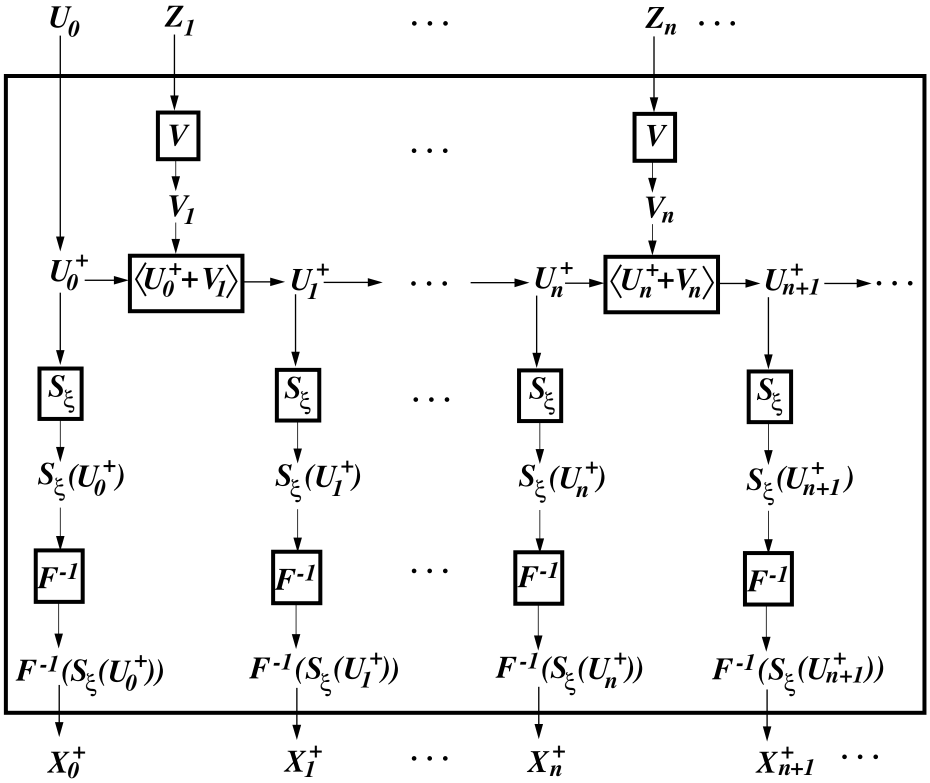

Figure 15: The General TES$^+$ Generation Scheme

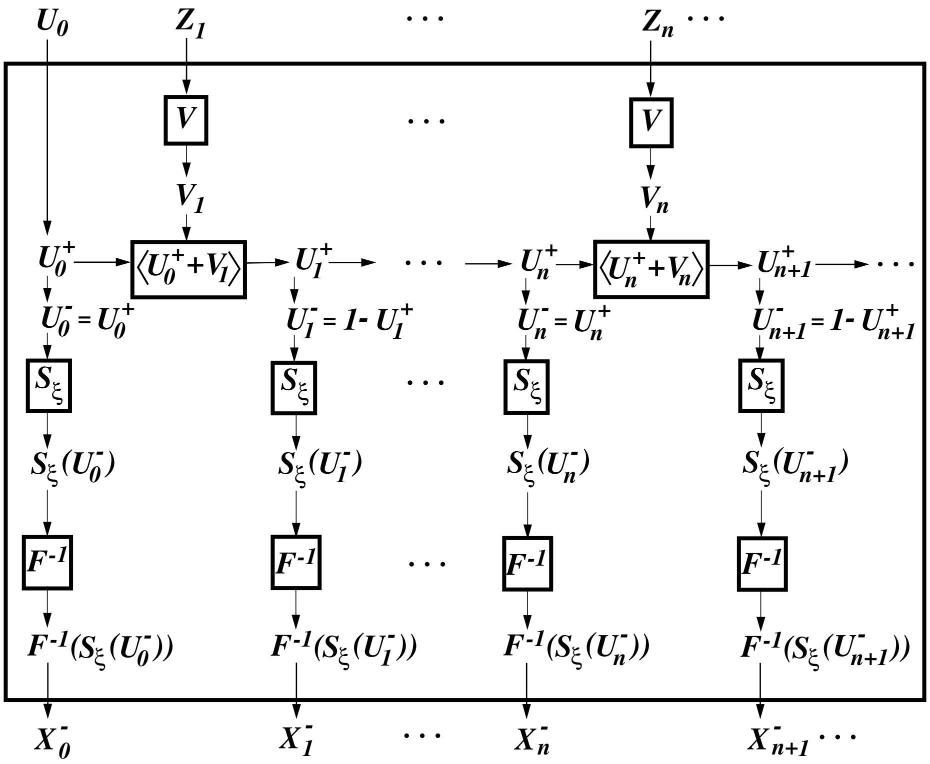

Figure 16: The General TES$^-$ Generation Scheme

in TEStool for empirical TES modeling. Formulas for these and other innovations and distortions were explicitly calculated in [8], and will not be reproduced here.

Figures 6–13 and 17–22 serve two purposes. They exhibit the efficacy of the numerical algorithms for calculating TES autocorrelations, and simultaneously demonstrate their range of functional forms, as well as the variety of the associated sample paths. The figures are arranged in pairs, where the first member displays a sample path obtained from a simulation run, and the next member displays the corresponding pair of TES autocorrelation functions of which one is estimated statistically from a simulation run of 10^6 observations (bullets), and the other is computed numerically from specializations of Eqs. (3.8)–(3.9). As evidenced by the graphs, the agreement is excellent (tabulated printouts confirm the accuracy of the algorithms for a large variety of additional cases). Recall that Figures 6–13 have already been discussed in Section 3.1.

The "smoothing" effect of stitching on basic TES sequences is demonstrated in Figures 6 and 17. Figure 17 is a stitched version ($\xi = 0.5$) of the sample path shown in Figure 6. Notice how the "discontinuities" in the sample path of Figure 6 (the two spikes around serial numbers 400 and 420, and the dip around serial number 470) have been eliminated in Figure 17; the latter has an overall "continuous" appearance. The effect of stitching on the corresponding autocorrelation functions is demonstrated in Figures 7–18. They clearly show that stitching has the effect of increasing the magnitude of the autocorrelation function. This phenomenon stems from the fact, that crossing point 0 in either direction (from large values to small ones or vice versa) introduces a negative component into the autocorrelations of unstitched sequences, and this negative component is eliminated in stitched ones. Refer to Figure 5 as a pictorial aid.

The effect of inversions (see Section 2.1) on sample paths and autocorrelation functions of background TES sequences is exemplified in Figures 19–22. Figures 19–20 correspond to exponential marginals (rate 1), while Figures 21–22 correspond to geometric marginals (parameter 0.5). These particular examples were selected to emphasize that TES can model discrete-valued sequences in addition to continuous-valued ones. A comparison of each distorted foreground sample path with its background counterpart reveals that certain visual features are preserved under unstitched inversions; this is due to the fact that inversion transformations (of a cdf) are monotonic; furthermore, the inversions in the aforementioned figures were applied to the same background sequence. For example, comparing Figures 19 and 21 to Figure 6 reveals that the general appearance of the sample path envelopes is quite similar, though the scales are different. In particular, the sample path "discontinuities" alluded to before are easily identifiable at the same serial numbers. An examination of the corresponding autocorrelation functions (Figures 20, 22, and 7) show that the functional form is largely preserved under inversion (up to scaling), though more complicated effects are described in [8].

The reader is referred to [8] for more information on TES sample paths and autocorrelations, and to [9] for the corresponding information on TES spectral densities.

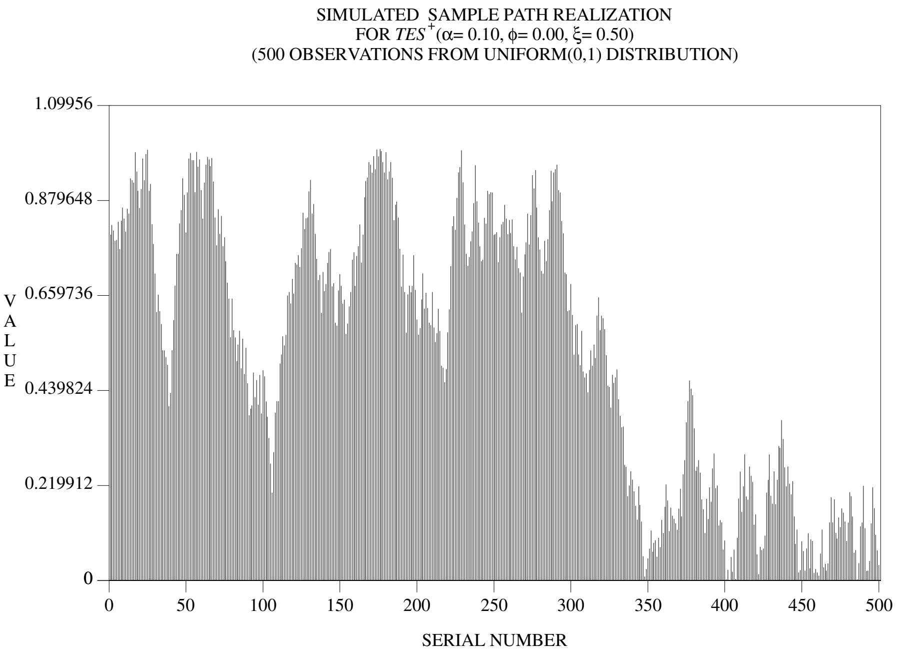

Figure 17: Sample Path of a Uniform Stitched Driftless TES$^+(\alpha, \phi)$ Process

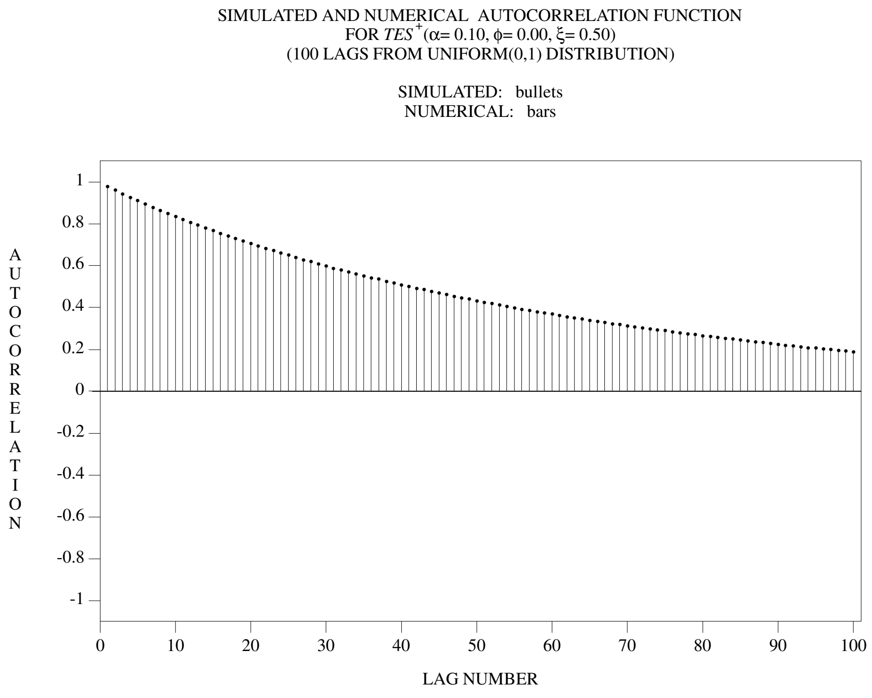

Figure 18: Autocorrelation Functions of a Uniform Stitched Driftless TES$^+(\alpha, \phi)$ Process

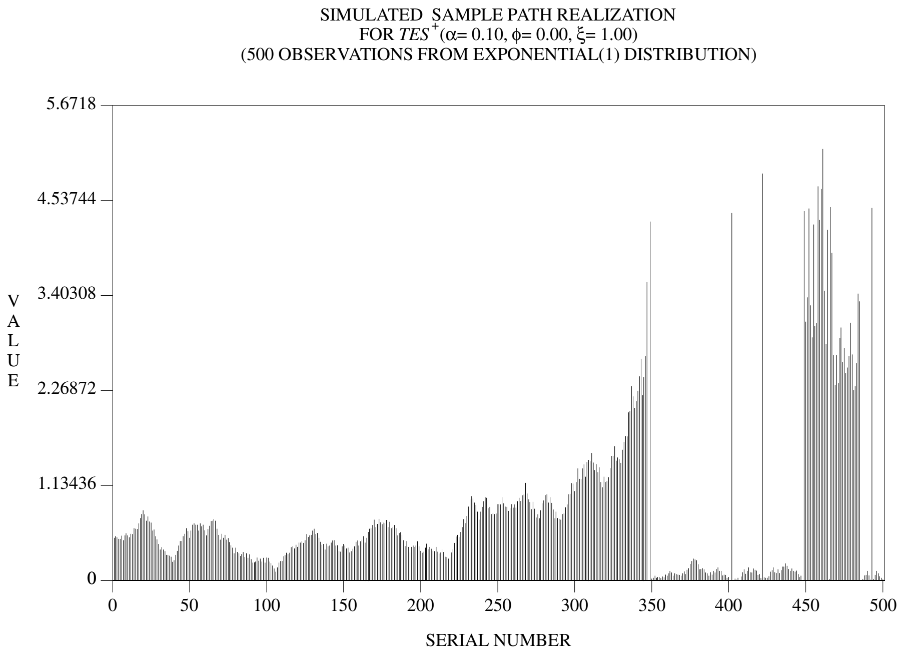

Figure 19: Sample Path of an Exponential TES$^+(\alpha, \phi)$ Process Without Drift

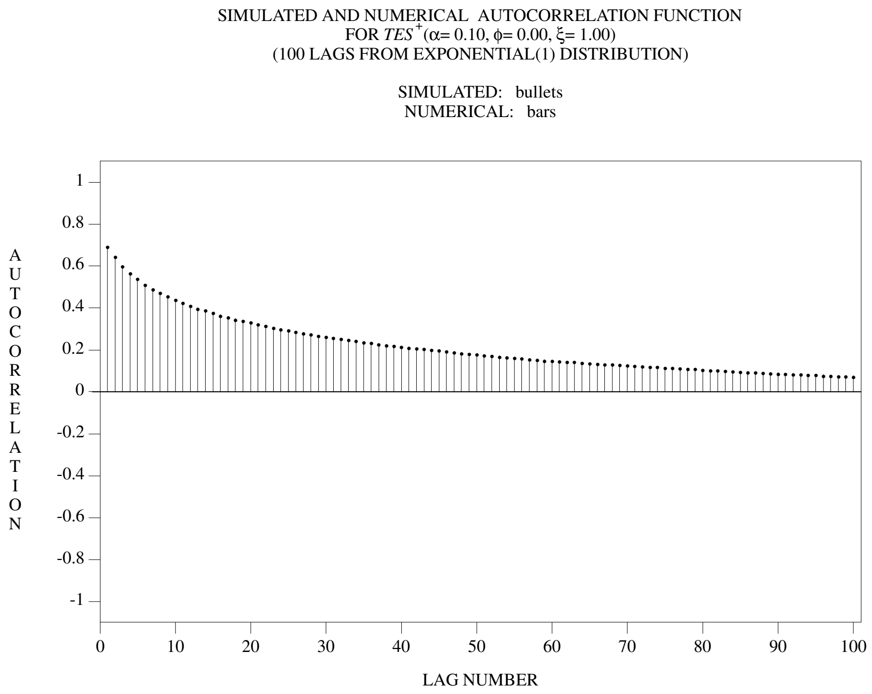

Figure 20: Autocorrelation Functions of an Exponential TES$^+(\alpha, \phi)$ Process Without Drift

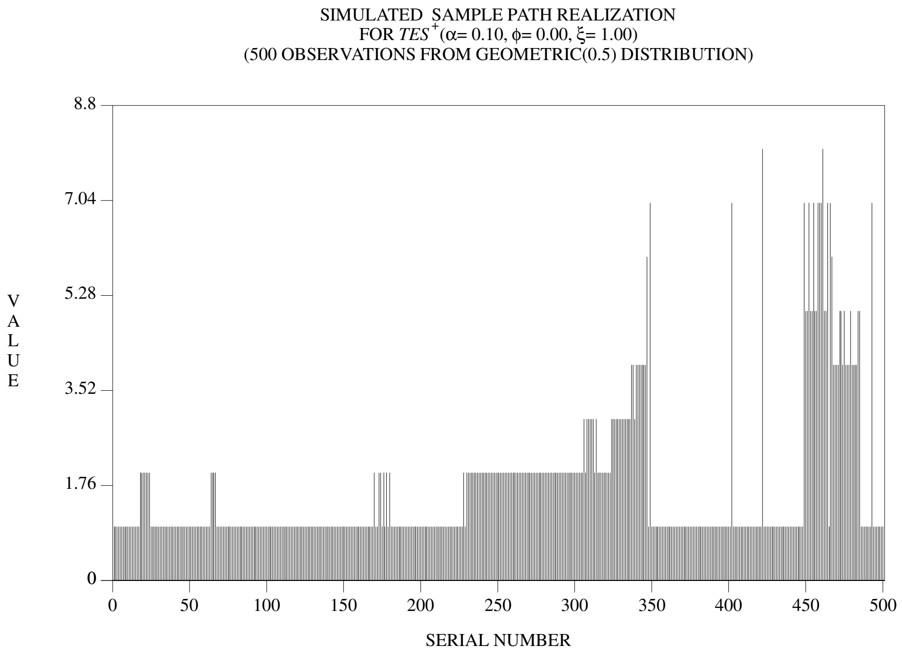

Figure 21: Sample Path of a Geometric TES$^{+}(\alpha, \phi)$ Process Without Drift

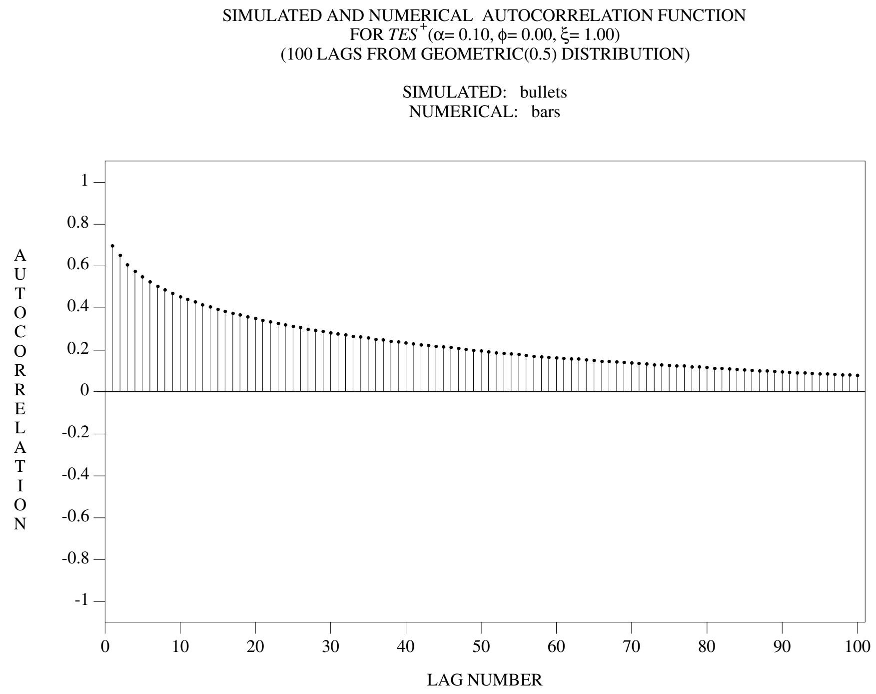

Figure 22: Autocorrelation Functions of a Geometric TES$^{+}(\alpha, \phi)$ Process Without Drift

4 TES MODELING AND SOFTWARE SUPPORT

The TES modeling methodology, as currently implemented in TEStool, uses specialized innovation and distortions. Unless otherwise stated, all innovation densities are assumed to be step functions (see Section 2.5), and all distortions are restricted to stitched or unstitched histogram inversions (see Section 2.2).

Recall that Section 2.2 covers continuous empirical histograms. If the empirical histogram is known to contain discrete components (that is, discrete values occur with non-zero probabilities), then a simple modification of Eqs. (2.2)–(2.4) yields analoguous formulas; the corresponding transforms required to calculate autocorrelations in Eqs. (3.8) and (3.9) can be found in [8]. We shall not elaborate on the discrete histogram case, because it occurs relatively rarely and does not introduce significant new wrinkles.

4.1 The Empirical TES Modeling Methodology

Modeling empirical time series is the most common and useful application of the TES modeling methodology. Assume that we have at our disposal some empirical sample path, representing a partial time series history, to which we wish to fit a TES model. Assume further that an empirical histogram and autocorrelation function statistics have been computed from the empirical sample path. A complete specification of a TES sequence requires the modeler to specify a flavor of TES process (TES$^+$ or TES$^-$), a distortion D composed of a stitching parameter ξ and histogram inversion H^{-1}, and a step-function innovation density f_V. An outline of a typical modeling scenario is sketched below.

Selecting the TES Process sign: The selection of the TES process sign is based on the modeler's experience and knowledge of TES sequences.

Selecting the inversion: The natural inversion is the histogram distortion D_H, given in Eq. (2.4). Since this is completely determined by the empirical data, no additional choices are required on the modeler's part. Using a histogram inversion ensures that the TES sequence will match the empirical marginal distribution, regardless of other parameter selections.

Selecting the stitching parameter and innovation density: The core activity of empirical TES modeling is a heuristic search for a suitable stitching parameter and innovation density. The modeler searches through stitching parameters in the range $[0, 1]$ and innovation densities in the space of step-function densities whose support is contained in $[-0.5, 0.5)$. These selections fix the autocorrelation structure of a TES sequence.

Reference [8] provides the ingredients for a numerical calculation of TES$^+$ and TES$^-$ autocorrelation functions from the Fourier representations of Eqs. (3.8) and (3.9). Specifically, it contains the formulas for the Laplace transform of the histogram inversion (2.4), as well as stitched histogram inversions of the form $D(x) = H^{-1}(S_\xi(x))$ (a special case of Eq. (2.9)). It also contains formulas for stitched and unstitched inversions corresponding to a host of other distributions, such as uniform, exponential and geometric. Similarly, [8] also provides formulas for the Laplace transform of step-function densities, as well as other innovation densities.

We stress that the modeling outline above is highly heuristic. The modeler may loop back to any step, based on the quantitative fit of the current model's autocorrelation function and the qualitative fit of Monte Carlo sample paths to the respective empirical counterparts. Nevertheless, several rules of thumb have been gleaned from experimentation with various modeling applications in TEStool. Some important ones are summarized below.

Effect of the sign of a TES sequence: Experience shows that TES$^+$ models are the most common choice, in practice. Consequently, the rest of the comments below address this case. TES$^-$ sequences should be considered, however, when empirical sample paths or autocorrelation functions have an alternating (zigzag) appearance.

Effect of the width of an innovation density support: The width of the innovation density support affects the amplitude of the autocorrelation function. The larger the support, the smaller the amplitude. The intuitive reason is that a larger support permits larger distances between successive TES iterates, resulting in a reduced autocorrelation magnitude.

Effect of the location of an innovation density support: The location of the innovation density support controls the frequency of oscillation of the autocorrelation function. The farther the support from the origin, the higher the frequency of oscillation. This is more easily seen for basic TES$^+(\alpha, \phi)$ sequences (see Section 3.1). When $\phi = 0$, the random walk is driftless. The autocorrelation function is monotone decreasing, and a spectral analysis reveals no periodicities. When $\phi \neq 0$, the random walk is directional, and the autocorrelation function is oscillatory. Consequently, the sample paths have a cyclical appearance, and the presence of periodicities can be confirmed by spectral analysis.

Effect of the stitching parameter: The fundamental effect of $0 < \xi < 1$ is to “smooth” the sample paths of TES sequences, whereas for $\xi = 1$ or $\xi = 0$, no "smoothing" takes place. In cyclical TES processes, ξ can be used to skew sample paths in accordance with the corresponding stitching transformation. Applying S_ξ for ξ ranging in [0,1] will shift the cycle peaks in the corresponding $\{S_\xi(U_n^+)\}$ in proportion to ξ. In particular, S_0 corresponds to the antithetic transformation $(S_0(y) = 1 - y)$; S_1 is the identity ($S_1(y) = y$); and $S_{0.5}$ gives rise to stitched TES$^+$ processes with symmetrical cycles. Recall that a stitching transformation increases the magnitude of the autocorrelation function of a background sequence as ξ increases or decreases towards 0.5 (see Section 3.3). This effect is strictly monotonic. It can also be shown that for any background TES sequence $\{U_n\}$, the autocorrelation function of any compound distortion (2.9) has the symmetry property that the autocorrelation functions of $\{F^{-1}(S_\xi(U_n))\}$ and $\{F^{-1}(S_{1-\xi}(U_n))\}$ are identical [7].

Effect of inversion distortions: An inversion distortion affects the background sample paths through a "scaling" effect. Interestingly, it has a similar quantitative effect on the corresponding autocorrelation function. However, it does not affect it qualitatively, in the sense that it leaves the functional form of the background autocorrelation function unchanged. These effects can be attributed to the monotonic nature of inversion distortions.

4.2 The TEStool Modeling Support Environment

Undirected heuristic searches for a TES model can be both time-consuming and mentally tedious. In the absence of an algorithmic fitting method, a human modeler requires computerized support with visual interactive facilities. Because the TEStool software environment, developed for this purpose, is such an integral part of the TES modeling activities, it will be briefly overviewed in this section.

TEStool is a visual interactive software environment for fitting TES models to empirical data [6, 16]. It provides services to generate and modify TES models, and to examine their statistics, in the context of the heuristic search scenario outlined in Section 4.1.

TEStool distinguishes between three types of statistics. *Empirical* statistics are those associated with the empirical data (sample paths, histogram, autocorrelation function or spectral density). *Simulated* statistics are similarly calculated (estimated) from Monte Carlo simulations of TES sequences. *Analytical* statistics consist of numerical computations of autocorrelation functions and spectral densities of TES models, based on analytical formulas developed in [8].

Figure 1 displays the now familiar TEStool screen comprised of four tiled subwindows (canvases). The upper-left subwindow contains sample paths, the upper-right subwindow contains histograms, the lower-left subwindow contains autocorrelation functions, and the lower-right subwindow contains a graphical specification of a TES model. The latter consists of a joint specification of a TES sign (plus or minus), stitching parameter ξ, an innovation density f_V, and an initial value U_0 for the backgroundTES sequence. An inverse distribution distortion can be selected from a menu, including the histogram distortion (2.4) constructed from the empirical histogram. A TES model can also be specified in TEStool in standard text mode, by filling out text fields in a form. The buttons in the top border of the display and at the bottom of each subwindow control various modeling services. These include reading and writing datasets, subdividing the screen real estate, opening a TES specification window or menu, performing various computations and quitting the system.

The most important service is the visual specification of a TES model and the interactive computations associated with it. The advantage of a graphical specification is that it can be grasped intuitively; more importantly, it enables the user to interact flexibly with the modeling environment, significantly increasing modeling efficiency and substantially decreasing modeling tedium. In the visual specification mode, the user operates the mouse to draw non-overlapping rectangles; each rectangle represents the support and probability of a step, and together they define a step-function density. In the interactive mode, changes in model specification trigger an immediate recomputation and redisplay of model statistics. The modeler may then compare the new statisics to their empirical counterparts for goodness of fit, and decide whether the fit is satisfactory or whether another iteration is required. The heuristic search can then proceed efficiently, and the visual feedback has the additional advantage of aiding the modeler in learning from experience. Finally, since the search is cast in a visual-interactive mold, it can be carried out by experts and non-experts alike in the same way that a player of an arcade game can concentrate on the task at hand without having to understand the underlying software details.

The lower-right subwindow in Figure 1 serves to clarify the advantage of dealing with step-function densities, from a user interface vantage point. A step-function is particularly easy to specify on a compuer screen, as each step is visually represented as an ordinary rectangle — an exceptionally simple geomertical form. Contemporary computer graphics enable intuitive modifications of rectangles in a natural way with the aid of a mouse. It is a simple matter to change the size of a rectangle, by "stretching" a side or a corner. In a similar vein, translating a rectangle on the horizontal axis is accomplished by the familiar operation of "dragging" an icon. The same subwindow also clarifies the choice of the interval $[-0.5, 0.5)$ as the support set for innovation densities. Any such support set represents the unraveling of the unit circle into a linear interval, so any interval of length one is an equivalent choice due to the modulo-1 arithmetic employed in the definition of TES sequences. However, among all such support sets, the interval $[-0.5, 0.5)$ is most compatible with the intuitive interpretation of innovation variates as modulo-1 increments (decrements) of a TES$^+$ sequence, as explained in Sections 3.1 and 3.2. Because the origin represents the relative location of the current TES iterate on the unit circle, the innovation density steps have an obvious interpretation: Step (L_k, R_k, P_k) simply means that the next increment (decrement) of the sequence is uniform on $[L_k, R_k)$ with probability P_k.

4.3 Example: A TES Model of Compressed Video Traffic

Figure 1 is reproduced from an actual workstation display screen to illustrate an example in which TEStool has been used to model real-world data. Here the empirical data is a (random) sequence of encoded (compressed) video frames [11, 19]. Video information is rarely transmitted over a network in its raw form. Rather, engineers take advantage of the considerable visual redundancy inherent in digitized pictures to compress each frame into a fraction of its original size. The compressed frames have random sizes (bit rates) which are then transported over the network and decoded at their destination. The term VBR (variable bit rate) video is used to refer to this kind of video traffic. The VBR video frames modeled in Figure 1 were generated by a video sequence from a football scene, whose raw frames were compressed by a variant of DCT (discrete cosine transform) [19]. See [17] for a review of compressed video modeling, using the TES methodology.

Recall that the upper-left subwindow in Figure 1 displays the empirical data (bullets); superimposed on it is a typical sample path (diamonds), generated by a Monte Carlo simulation of the TES model exhibited in the lower-right subwindow. The histogram and autocorrelation function of the TES model are similarly plotted against their empirical counterparts in the upper-right and lower-left subwindows, respectively. Recall that while the TES model histogram was computed from the TES Monte Carlo sample path, the TES model autocorrelations were calculated numerically with the aid of Eq. (3.8). The lower-right subwindow contains a visual specification of a TES$^+$ model with a stitching parameter $\xi = 0.5$ and a step-function density consisting of two steps.

Observe how the modeling results in Figure 1 satisfy the modeling requirements of Section 1.3. Indeed, the TES model closely matches the empirical histogram and approximates the leading autocorrelations quite well (there is no point here in trying to approximate more than about 10 lags, as there are only some 210 data points in the empirical time series). In addition, note that the qualitative "resemblance" of the TES model sample path to the empirical path is remarkable. Thus, the empirical TES modeling summarized in Figure 1 represents a successful modeling effort, by the criteria of Section 1.3.

5 CONCLUSION

This paper presented a fairly comprehensive tutorial on TES models, an outline of the empirical TES modeling methodology, and a brief overview of the TEStool visual interactive softare environment, designed to support TES modeling.

The TES modeling methodology is a novel input analysis approach which aims to provide more faithful models of temporally-dependent empirical time series; it strives to fit *simultaneously* both marginal distributions and the leading autocorrelations of empirical samples, and to imitate their qualitative character. Its primary purpose is to generate realistic synthetic source models to drive Monte Carlo simulations. Its approach to capturing first-order and second-order empirical statistics also qualifies it for "black-box" forecasting applications.

Practical TES modeling currently relies on software support to carry out heuristic searches for adequate models. The TEStool visual interactive software was created with this goal in mind. Its Graphical User Interface casts the heuristic search into an intuitive interactive activity of modifying a visual parametric representation of a TES model in small incremental stages; it, further, provides visual feedback to guide the search process. The TEStool modeling environment has been shown to yield remarkably accurate TES models in a reasonable amount of time. It also alleviates the tedium of repetitive search, transforming the modeling process into a pleasant activity.

Further information on TES-related topics may be found in the reference list enclosed. The reader is referred to [15, 7, 8, 9] for TES theory; to [11, 17, 14] for TES modeling applications; and to [6, 16] for the TEStool modeling environment.

Acknowledgments

I am indebted to David Jagerman and Jon Hill for carefully reading and commenting on the manuscript.

References

[1] J.S. Bendat and A.G. Piersol. *Random Data*. Wiley, 1986.

[2] P. Bratley, B.L. Fox and L.E. Schrage. *A Guide to Simulation*. Springer-Verlag, New York, NY, 1987.

[3] D.R. Cox and P.A.W. Lewis. *The Statistical Analysis of Series of Events*. Methuen, 1968.

[4] W. Feller. *An Introduction to Probability Theory and Its Applications*. Vol. 2, 2nd edition, Wiley, 1971.

[5] K.W. Fendick, V.R. Saksena and W. Whitt. Dependence in Packet Queues. *IEEE Trans. on Comm.*, Vol. 37, 1173–1183, 1989.

[6] D. Geist and B. Melamed, “TEStool: An Environment for Visual Interactive Modeling of Autocorrelated Traffic”, *Proceedings of the 1992 IEEE International Conference on Communications*, Chicago Illinois, 1285–1289, 1992.

[7] D. Jagerman and B. Melamed. The Transition and Autocorrelation Structure of TES Processes Part I: General Theory. *Stochastic Models*, Vol. 8, No. 2, 193–219, 1992.

[8] D. Jagerman and B. Melamed. The Transition and Autocorrelation Structure of TES Processes Part II: Special Cases. *Stochastic Models*, Vol. 8, No. 3, 499–527, 1992.

[9] D. Jagerman and B. Melamed. Spectral Analysis of Basic TES Processes. *Preprint*, NEC USA, Inc., Princeton, New Jersey, 1992.

[10] A.M. Law and W.D. Kelton. *Simulation Modeling & Analysis*, (second edition). McGraw-Hill, 1991.

[11] D.-S. Lee, B. Melamed, A. Reibman and B. Sengupta. TES Modeling for Analysis of a Video Multiplexor. *Performance Evaluation*, Vol. 16, 21–34, 1992.

[12] P.A.W. Lewis and E. McKenzie. Minification Processes and Their Transformations. *J. Appl. Prob.*, Vol. 28, 45–57, 1991.

[13] B. Liu and D.C. Munson. "Generation of a Random Sequence Having a Jointly Specified Marginal Distribution and Autocovariance". *IEEE Transactions on Acoustics, Speech and Signal Processing* Vol. 30, No. 6, 973–983, 1982.

[14] M. Livny, B. Melamed and A.K. Tsiolis. The Impact of Autocorrelation on Queuing Systems. *Management Science*, Vol. 39, No. 3, 322–339, 1993.

[15] B. Melamed. TES: A Class of Methods for Generating Autocorrelated Uniform Variates. *ORSA J. on Computing*, Vol 3, 317-329, 1991.

[16] B. Melamed, D. Goldsman and J.R. Hill. The TES Methodology: Nonparametric Modeling of Stationary Time Series. *Proceedings of the 1992 Winter Simulation Conference*, Arlington, Virginia, 135–144, 1992.

[17] B. Melamed and B. Sengupta. "TES Modeling of Video Traffic". *IEICE Transactions on Communications*, Vol. E75-B, No. 12, 1291–1300, 1992.

[18] B.E. Patuwo, R.L. Disney and D.C. McNickle. The Effects of Correlated Arrivals on Queues. *IIE Tranactions*, Vol. 25, No. 3, 105–110, 1993.

[19] A.R. Reibman, “DCT-based Embedded Coding for Packet Video”, *Signal Processing: Image Communication*, Vol. 3, 231–237, 1991.

[20] B.W. Schmeiser. Simulation Experiments. In *Handbook of Operations Research and Management Science: Vol. 2* (D. Heyman and M. Sobel, editors). North Holland, New York, 1990.

[21] R.H. Thaler. *The Winner’s Curse*. The Free Press, 1992.