Some Properties of Operators for Unevenly Spaced Time Series*

Andreas Eckner†

First version: August 2012
Current version: April 3, 2017

Abstract

An unevenly spaced time series is a sequence of observation time and value pairs with strictly increasing observation times. This paper examines some properties of such time series and of their operators.

Theorem 6.13 uniquely characterizes the structure of the most widely used class of linear operators, showing that each such linear operator is associated with a finite signed measure, and vice versa, that each finite signed measured defines such a linear operator.

1 Introduction

There is extensive theory for the analysis of equally spaced data. However, much fewer methods exist specifically for unevenly spaced time series. The emphasis of this literature has primarily been on achieving a specific goal, such as modeling univariate autoregressive-moving-average (ARMA) processes; see Jones (1981), Jones (1985), Belcher et al. (1994) and Brockwell (2008). In astronomy, much effort has been devoted to the task of estimating the spectrum of irregular time series data; see Lomb (1975), Scargle (1982), Bos et al. (2002), Thiebaut and Roques (2005), and Broersen (2008). Müller (1991), Gilles Zumbach (2001), and Dacorogna et al. (2001) examined unevenly spaced time series in the context of high-frequency financial data.

This paper takes a broader approach by examining the general properties of unevenly spaced time series and their operators, without any specific application in mind. A special effort is devoted to analyzing the structure of linear operators.

Section 2 introduces unevenly spaced time series and some elementary operations for them. Section 3 defines time series operators and examines their commonly encountered structural features. Section 4 introduces convolution operators, which form a particularly tractable and widely used class of time series operators. Section 5 presents arithmetic, return, and rolling operators as examples. Section 6 extensively analyzes linear time series operators, while Section 7 turns to multivariate time series and associated vector time series operators. Moving averages, which summarize the average value of a time series over a certain horizon, are examined in Section 8. An Appendix summarizes frequently used notation.

*This paper previously circulated under the title "A Framework for the Analysis of Unevenly Spaced Time Series Data." Based on feedback from readers I have shortened the paper and narrowed the scope.

†Comments are welcome at andreas@eckner.com

2 Unevenly Spaced Time Series

An unevenly spaced time series is a sequence of observation time and value pairs (t_n, X_n) with strictly increasing observation times. This notion is made precise by the following

Definition 2.1 *For $n \geq 1$, we call*

(i) $\mathbb{T}_n = \{(t_1 < t_2 < \ldots < t_n) : t_k \in \mathbb{R}, 1 \leq k \leq n\}$ *the space of strictly increasing time sequences of length n,*

(ii) $\mathbb{T} = \cup_{n=1}^{\infty} \mathbb{T}_n$ *the space of strictly increasing time sequences,*

(iii) $\mathbb{R}^n$ *the space of observation values for n observations,*

(iv) $\mathcal{T}_n = \mathbb{T}_n \times \mathbb{R}^n$ *the space of real-valued, unevenly spaced time series of length n, and*

(v) $\mathcal{T} = \cup_{n=1}^{\infty} \mathcal{T}_n$ *the space of (real-valued) (unevenly spaced) time series.*

Definition 2.2 *For a time series $X \in \mathcal{T}$, we denote*

(i) $N(X)$ *to be the number of observations of X, particularly so that $X \in \mathcal{T}_{N(X)}$;*

(ii) $T(X) = (t_1, \ldots, t_{N(X)})$ *to be the sequence of observation times (of X); and*

(iii) $V(X) = (X_1, \ldots, X_{N(X)})$ *to be the sequence of observation values (of X).*

We will frequently use the informal but compact notation $((t_n, X_n) : 1 \leq n \leq N(X))$ and $(X_{t_n} : 1 \leq n \leq N(X))$ to denote a time series $X \in \mathcal{T}$ with observation times $(t_1, \ldots, t_{N(X)})$ and observation values $(X_1, \ldots, X_{N(X)})$.[1] Now that we have defined unevenly spaced time series, we will introduce methods for extracting basic information from them.

Definition 2.3 *For a time series $X \in \mathcal{T}$ and point in time $t \in \mathbb{R}$ (not necessarily an observation time), the most recent observation time is*

$$\text{Prev}^X(t) \equiv \text{Prev}(T(X), t) = \begin{cases} \max(s : s \leq t, s \in T(X)) & \text{if } t \geq \min(T(X)), \\ \min(T(X)) & \text{otherwise,} \end{cases}$$

and the next available observation time is

$$\text{Next}^X(t) \equiv \text{Next}(T(X), t) = \begin{cases} \min(s : s \geq t, s \in T(X)) & \text{if } t \leq \max(T(X)) \\ +\infty & \text{otherwise.} \end{cases}$$

For $\min T(X) \leq t \leq \max T(X)$, $\text{Prev}^X(t) < t < \text{Next}^X(t)$; unless $t \in T(X)$, in which case t is both the most recent and the next available observation time.

[1] For equally spaced time series, the reader may be used to using language like "the third observation" of a time series X. For unevenly spaced time series, it is often necessary to distinguish between the third observation *value*, X_{t_3}, and the third observation tuple, or simply the third observation, (t_3, X_3), of a time series.

Definition 2.4 (Sampling) *For $X \in \mathcal{T}$ and $t \in \mathbb{R}$, $X[t]_{\text{last}} = X_{\text{Prev}(X,t)}$ is the sampled value of X at time t, $X[t]_{\text{next}} = X_{\text{Next}^X(t)}$ is the next value of X at time t, and $X[t]_{\text{linear}} = (1 - \omega^X(t))X_{\text{Prev}(X,t)} + \omega^X(t)X_{\text{Next}^X(t)}$ where*

$$\omega^X(t) = \omega(T(X), t) = \begin{cases} \frac{t - \text{Prev}^X(t)}{\text{Next}^X(t) - \text{Prev}^X(t)} & \text{if } 0 < \text{Next}^X(t) - \text{Prev}^X(t) < \infty, \\ 0 & \text{otherwise,} \end{cases}$$

is the linearly interpolated value of X at time t. These sampling schemes are called last-point, next-point, and linear interpolation, respectively.

Note that the most recently available observation time *before* the first observation is taken to be the first observation time. As a consequence, the sampled value of a time series X *before* the first observation is equal to the first observation value. While potentially inappropriate for some applications, this convention greatly simplifies notation and avoids the treatment of a multitude of special cases in the exposition below.[2]

Remark 2.5 *Fix a time series $X \in \mathcal{T}$. Then,*

(i) $X[t]_{\text{last}} = X[t]_{\text{next}} = X[t]_{\text{linear}} = X_t$ *for* $t \in T(X)$*;*

(ii) $X[t]_{\text{last}}$ *and* $X[t]_{\text{next}}$*, as functions of t, are right-continuous piecewise-constant functions with finite number of discontinuities; and*

(iii) $X[t]_{\text{linear}}$ *as a function of t is a continuous piecewise-linear function.*

These observations suggest an alternative way of defining unevenly spaced time series; namely, as either piecewise-constant or piecewise-linear functions $X : \mathbb{R} \to \mathbb{R}$. Such a representation, however, cannot capture the occurrence of identical consecutive observations, and therefore ignores potentially important time series information. Moreover, such a framework does not naturally lend itself to interpreting an unevenly spaced time series as a discretely-observed continuous-time stochastic process, thereby ruling out a large class of data-generating processes.

Definition 2.6 (Simultaneous Sampling) *Fix an observation time sequence $T_X \in \mathbb{T}$.*

(i) For a time series $X \in \mathcal{T}$ and sampling scheme $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$, we call

$$X_\sigma[T_X] = ((t_i, X_\sigma[t_i]_{\text{last}}) : t_i \in T_X)$$

the sampled time series of X (using sampling times T).

(ii) For a continuous-time stochastic process X^c, we call

$$X^c[T] = ((t_i, X^c_{t_i}) : t_i \in T_X)$$

the observation time series of X^c (at observation times T_X).

[2] A software implementation might instead use a special symbol to denote a value that is not available. For example, `R` (www.r-project.org) uses the constant `NA`, which propagates through all steps of an analysis because the result of any calculation involving `NA`s is also `NA`.

In particular, $X[T(X)]_\sigma = X$ for all time series $X \in \mathcal{T}$ and sampling schemes σ.

Definition 2.7 *For $X \in \mathcal{T}$, we call*

(i) $\Delta t(X) = ((t_{n+1}, t_{n+1} - t_n) : 1 \le n \le N(X) - 1)$ *the time series of tick (or observation time) spacings (of X);*

(ii) $X\{s,t\} = ((t_n, X_n) : s \le t_n \le t, 1 \le n \le N(X))$ *for $s \le t$ the subperiod time series (of X) in $[s,t]$;*

(iii) $\mathrm{B}(X) = ((t_{n+1}, X_n) : 1 \le n \le N(X) - 1)$ *the backshifted time series (of X), and* B *the backshift operator;*

(iv) $\mathrm{L}(X, \tau) = ((t_n + \tau, X_n) : 1 \le n \le N(X))$ *the lagged time series (of X) with lag $\tau \in \mathbb{R}$, and* L *the lag operator; and*

(v) $\mathrm{D}(X, \tau) = ((t_n, X[t_n - \tau]_{\text{last}}) : 1 \le n \le N(X)) = \mathrm{L}(X[T(X) - \tau]_{\text{last}}, -\tau)$ *the delayed time series (of X) with delay $\tau \in \mathbb{R}$, and* D *the delay operator.*

A time series X is equally spaced, if the observation values of $\Delta t(X)$ are all equal to a constant $c > 0$. For such a time series and for $\tau = c$, the backshift operator is identical to the lag operator (apart from the last observation) and to the delay operator (apart from the first observation). In particular, the backshift, delay, and lag operator are identical for an equally spaced time series with time index $t \in \mathbb{Z}$. These transformations are genuinely different for unevenly spaced data, however, because the backshift operator shifts observation *values*, while the lag operator shifts observation *times*.[3]

Example 2.8 *Let X be a time series of length three with observation times $T(X) = (0, 2, 5)$ and observation values $V(X) = (1, -1, 2.5)$. Then*

$$X[t]_{\text{last}} = \begin{cases} 1 & \text{for } t < 2, \\ -1 & \text{for } 2 \le t < 5, \\ 2.5 & \text{for } t \ge 5, \end{cases} \qquad X[t]_{\text{next}} = \begin{cases} 1 & \text{for } t \le 0, \\ -1 & \text{for } 0 < t \le 2, \\ 2.5 & \text{for } t > 2, \end{cases}$$

$$X[t]_{\text{linear}} = \begin{cases} 1 & \text{for } t < 0 \\ 1 - t & \text{for } 0 \le t < 2 \\ \frac{7}{6}t - \frac{10}{3} & \text{for } 2 \le t < 5, \\ 2.5 & \text{for } t \ge 5. \end{cases}$$

These three graphs are shown in Figure 1. Moreover,

$$\Delta t(X)[s]_{\text{last}} = \begin{cases} 2 & \text{for } s < 5, \\ 3 & \text{for } s \ge 5, \end{cases}$$

because $T(\Delta t(X)) = (2, 5)$ and $V(\Delta t(X)) = (2, 3)$, and

$$\mathrm{B}(X)[t]_{\text{last}} = \begin{cases} 1 & \text{for } t < 5, \\ -1 & \text{for } t \ge 5, \end{cases} \qquad \mathrm{L}(X, 1)[t]_{\text{last}} = \begin{cases} 1 & \text{for } t < 3, \\ -1 & \text{for } 3 \le t < 6, \\ 2.5 & \text{for } t \ge 6. \end{cases}$$

[3] The difference between the lag and delay operator is that the former shifts the information filtration of observation *times and values*, while the latter shifts only the information filtration of observation *values*.

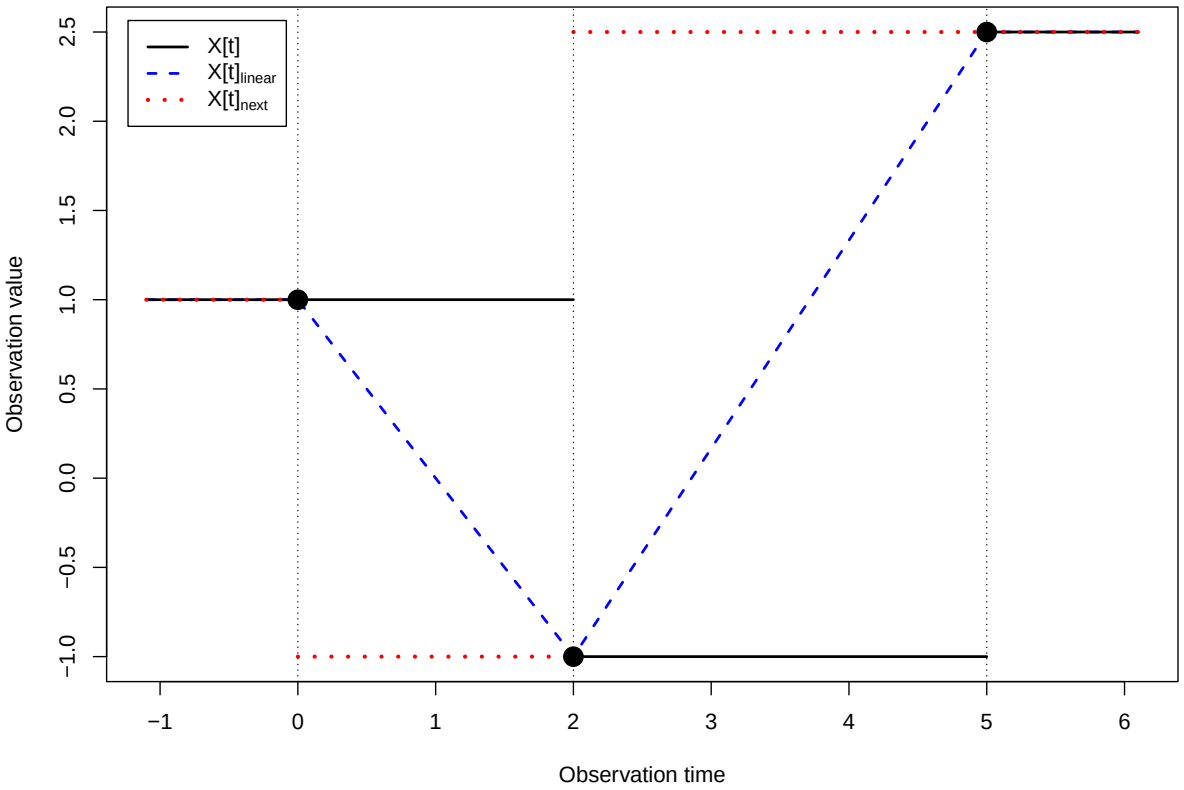

Figure 1: The sampling-scheme-dependent graph of the time series X with $T(X) = (0, 2, 5)$ and $V(X) = (1, -1, 2.5)$. The three observation time-value pairs are denoted by solid black circles.

The following result, which elaborates on the relationship between the lag operator L and the sampling operator, will be heavily used throughout the paper.

Lemma 2.9 *For $X \in \mathcal{T}$ and $\tau \in \mathbb{R}$,*

(i) $T(\mathrm{L}(X, \tau)) = T(X) + \tau$,

(ii) $\mathrm{L}(X, \tau)_{t+\tau} = X_t$ *for* $t \in T(X)$,

(iii) $\mathrm{L}(X, \tau)_t = X_{t-\tau}$ *for* $t \in T(\mathrm{L}(X, \tau))$, *and*

(iv) $\mathrm{L}(X, \tau)[t]_\sigma = X[t - \tau]_\sigma$ *for* $t \in \mathbb{R}$ *and sampling scheme* σ. *In other words, depending on the sign of* τ, *the lag operator shifts the sample path of a time series either backward or forward in time.*

Proof. Relationships (i) and (ii) follow directly from the definition of the lag operator, while (iii) follows by combining (i) and (ii). For (iv), we note that

$$\mathrm{Prev}^{\mathrm{L}(X,\tau)}(t) = \mathrm{Prev}(T(\mathrm{L}(X, \tau)), t) = \mathrm{Prev}(T(X) + \tau, t) = \mathrm{Prev}(T(X), t - \tau) + \tau$$

where the second equality follows from (i). Hence,

$$\begin{aligned}\mathrm{L}(X,\tau)[t]_{\text{last}} &= \mathrm{L}(X,\tau)_{\mathrm{Prev}^{\mathrm{L}(X,\tau)}(t)} \\ &= \mathrm{L}(X,\tau)_{\mathrm{Prev}(T(X),t-\tau)+\tau} \\ &= X_{\mathrm{Prev}(T(X),t-\tau)} \\ &= X[t-\tau]_{\text{last}},\end{aligned}$$

where the third equality follows from (ii). The proofs for the other two sampling schemes are similar. ■

At this point, the reader might want to check her understanding by verifying the identity $X = \mathrm{L}(\mathrm{L}(X,\tau)[T(X)+\tau]_{\text{last}}, -\tau)$ for all time series $X \in \mathcal{T}$ and $\tau \in \mathbb{R}$.

3 Time Series Operators

Time series operators (also called "systems" in the signal processing literature) take a time series as input and leave as output a transformed time series. We already encountered a few operators; such as the backshift (B), subperiod ({}), and tick spacing operator (Δ); in the previous section.

The key difference between time series operators for evenly and unevenly spaced observations is that, for the latter, the observation values in the transformed series can depend on the spacing of observation times. This interaction between observation values and observation times calls for a careful analysis and classification of the structure of such operators.

Definition 3.1 *A time series operator is a mapping* $\mathrm{O} : \mathcal{T} \to \mathcal{T}$*; or equivalently, a pair of mappings* $(\mathrm{O}_T, \mathrm{O}_V)$*; where* $\mathrm{O}_T : \mathcal{T} \to \mathbb{T}$ *is the transformation of observation times,* $\mathrm{O}_V : \mathcal{T} \to \cup_{n\geq 1}\mathbb{R}^n$ *is the transformation of observation values, and* $|\,\mathrm{O}_T(X)| = |\,\mathrm{O}_V(X)|$ *for all* $X \in \mathcal{T}$*.*

The constraint at the end of the definition ensures that the number of observation values and observation times for the transformed time series agree. In particular, using this notation, we have $T(\mathrm{O}(X)) = \mathrm{O}_T(X)$ and $V(\mathrm{O}(X)) = \mathrm{O}_V(X)$ for time series $X \in \mathcal{T}$ and a time series operator O.

Example 3.2 *Fix a sampling scheme* $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$ *and observation time sequence* $T \in \mathbb{T}$*. The mapping that assigns each time series* $X \in \mathcal{T}$ *the sampled time series* $X[T]_\sigma$ *(see Definition 2.6) is a time series operator.*

The above definition of a time series operator is completely general. In practice, most operators share one or more of the structural features listed below.

3.1 Tick Invariance

In many cases, the observation times of the transformed time series are identical to those of the original time series.

Definition 3.3 *A time series operator* O *is tick invariant if* $T(\mathrm{O}(X)) = T(X)$ *for all* $X \in \mathcal{T}$*.*

Such operators cover the vast majority of the cases in presented in this paper. Notable exceptions are the lag operator L and various resampling schemes. For some results, a weaker property than tick-invariance is sufficient.

Definition 3.4 *A time series operator* O *is lag free if* $\max T(\mathrm{O}(X)) \leq \max T(X)$ *for all* $X \in \mathcal{T}$.

In other words, for lag-free operators, when one has finished observing an input time series, the entire output time series is also observable at that point in time.

3.2 Causality

Definition 3.5 *A time series operator* O *is causal (or adapted) if*

$$\mathrm{O}(X)\{-\infty, t\} = \mathrm{O}(X\{-\infty, t\})\{-\infty, t\} \tag{3.1}$$

for all $X \in \mathcal{T}$ *and* $t \in \mathbb{R}$.

In other words, a time series operator is causal if the output up to each time point t depends only on the input up to that time.[4] Equivalently, if two time series are identical up to some time, then a causal transformation produces transformed time series that are likewise identical up to the same time.

In many cases, the following convenient characterization is useful:

Lemma 3.6 *A time series operator* O *is causal and lag free if and only if the order of* O *and the subperiod operator is interchangeable; that is*

$$\mathrm{O}(X)\{-\infty, t\} = \mathrm{O}(X\{-\infty, t\}) \tag{3.2}$$

for all $X \in \mathcal{T}$ *and* $t \in \mathbb{R}$.

Proof.

$\Longrightarrow$ If O is lag-free, then

$$\max T(\mathrm{O}(X\{-\infty, t\})) \leq \max T(X\{-\infty, t\}) \leq t \tag{3.3}$$

which implies that $\mathrm{O}(X\{-\infty, t\})\{-\infty, t\} = \mathrm{O}(X\{-\infty, t\})$ and combined with (3.1) shows (3.2).

$\Longleftarrow$ Applying $\max T(.)$ to both sides of (3.2) gives

$$\max T(\mathrm{O}(X\{-\infty, t\})) \leq t$$

for all $X \in \mathcal{T}$ and $t \in \mathbb{R}$. Setting $t = \max T(X)$ shows that O is lag-free. Hence (3.1) follows from (3.2) and (3.3).

■

In particular, the order of a causal, tick-invariant operator and the subperiod operator is interchangeable.

[4] If X was a stochastic process instead of a fixed time series, we would call an operator O adapted if $\mathrm{O}(X)$ was adapted to the filtration generated by X.

3.3 Shift or Time Invariance

For many data sets, only the relative, but not the absolute, positions of observation times are of interest; and this property should be reflected in the analysis that is carried out.

Definition 3.7 *A time series operator* O *is shift invariant (or time invariant), if the order of* O *and the lag operator* L *is interchangeable. In other words, for all* $X \in \mathcal{T}$ *and* $\tau \in \mathbb{R}$,

$$\mathrm{O}(\mathrm{L}(X,\tau)) = \mathrm{L}(\mathrm{O}(X),\tau); \tag{3.4}$$

or equivalently,

$$\mathrm{O}_T(\mathrm{L}(X,\tau)) = \mathrm{O}_T(X) + \tau, \text{ and} \tag{3.5}$$
$$\mathrm{O}_V(\mathrm{L}(X,\tau)) = \mathrm{O}_V(X). \tag{3.6}$$

Lemma 3.8 *A time series operator* O *is shift invariant if and only if*

$$\mathrm{O}(X)_t = \mathrm{O}(\mathrm{L}(X,-t))_0 \tag{3.7}$$

for all $X \in \mathcal{T}$ *and* $t \in T(\mathrm{O}(X))$.

Proof.

$\Longrightarrow$ Setting $\tau = -t$ in (3.4) and sampling at time zero gives

$$\begin{aligned}\mathrm{O}(\mathrm{L}(X,-t))_0 &= \mathrm{L}(\mathrm{O}(X),-t)_0\\ &= \mathrm{O}(X)_t,\end{aligned}$$

where the second equality follows from Lemma 2.9 (ii).

$\Longleftarrow$ Fix $X \in \mathcal{T}$, $t \in T(\mathrm{O}(\mathrm{L}(X,\tau)))$, and $\tau \in \mathbb{R}$. If (3.7) holds, then it particularly holds for $\mathrm{L}(X,t)$ and $\mathrm{L}(X,t-\tau)$, which gives

$$\begin{aligned}\mathrm{O}(\mathrm{L}(X,\tau))_t &= \mathrm{O}(\mathrm{L}(\mathrm{L}(X,\tau),-t))_0\\ &= \mathrm{O}(\mathrm{L}(X,\tau-t))_0\\ &= \mathrm{O}(\mathrm{L}(X,-(t-\tau)))_0\\ &= \mathrm{O}(X)_{t-\tau}\\ &= \mathrm{L}(\mathrm{O}(X,\tau))_t,\end{aligned}$$

where we use Lemma 2.9 (iii) for the last equality.

■

In particular, a (by definition $\mathcal{T}$-valued) time series operator that is shift invariant can be represented by a single real-valued function. Specifically, $\mathrm{O}(X)_t = g(X,t)$ for $t \in T(\mathrm{O}(X))$ and $g : \mathcal{T} \times \mathbb{R} \to \mathbb{R}$ given by $g(Y,t) = \mathrm{O}(\mathrm{L}(Y,-t))_0$.

3.4 Time Scale Invariance

The measurement unit of the time scale is generally not of interest in an analysis of time series data. Therefore, we usually focus on (families of) time series operators that are invariant under a rescaling of the time axis.

Definition 3.9 *For $X \in \mathcal{T}$ and $a > 0$, the time-scaling operator S_a is defined as*

$$\mathrm{S}_a(X) = \mathrm{S}(X, a) = ((at_n, X_n) : 1 \leq n \leq N(X)).$$

Lemma 3.10 *For $X \in \mathcal{T}$ and $a > 0$,*

(i) $T(S_a(X)) = aT(X)$,

(ii) $\mathrm{S}_a(X)_{at} = X_t$ *for* $t \in T(X)$,

(iii) $\mathrm{S}_a(X)_t = X_{t/a}$ *for* $t \in T(\mathrm{S}_a(X))$, *and*

(iv) $\mathrm{S}_a(X)[t]_\sigma = X[t/a]_\sigma$ *for* $t \in \mathbb{R}$ *and sampling scheme* σ. *In other words, depending on the sign of* $a - 1$, *the time-scaling operator either compresses or stretches the sample path of a time series.*

Proof. The proof is very similar to that of Lemma 2.9. ■

Certain sets of time series operators are naturally indexed by one or more parameters. For example, simple moving averages (see Section 8.1) can be indexed by the length of the moving average time window.

Definition 3.11 *A family of time series operators $\{\mathrm{O}_\tau : \tau \in (0, \infty)\}$ is timescale invariant if*

$$\mathrm{O}_\tau(X) = \mathrm{S}_{1/a}(\mathrm{O}_{\tau a}(\mathrm{S}_a(X)))$$

for all $X \in \mathcal{T}$, $a > 0$, and $\tau \in (0, \infty)$.

The families of rolling return operators (Section 5.2), simple moving averages (Section 8.1), and exponential moving averages (Section 8.2) are timescale invariant.

3.5 Homogeneity

Scaling properties are of interest not only for the time scale, but also for the observation value scale.

Definition 3.12 *A time series operator O is homogeneous of degree $d \geq 0$ if $\mathrm{O}(aX) = a^d\,\mathrm{O}(X)$ for all $X \in \mathcal{T}$ and $a > 0$.*[5]

The tick-spacing operator Δt is homogeneous of degree $d = 0$, moving averages are homogeneous of degree $d = 1$, and the operator that calculates the integrated p-variation of a time series is homogeneous of degree $d = p$.

Remark 3.13 *The sampling operator of Example 3.2 is linear (see Section 6), homogeneous of degree $d = 1$, and causal iff we use last-point sampling; and is neither tick invariant, shift invariant, nor timescale invariant.*

[5] For a time series X and scalar $a \in \mathbb{R}$, aX denotes the time series that results by multiplying each observation value of X by a. See Section 5.1 for a systematic treatment of time series arithmetic.

4 Convolution Operators

Convolution operators are a class of tick-invariant, causal, shift-invariant (and often homogeneous) time series operators that are particularly tractable. To this end, recall that a signed measure on a measurable space (Ω, Σ) is a measurable function μ that satisfies (i) $\mu(\varnothing) = 0$ and if $A = +_i A_i$ is a countable disjoint union of sets in Σ with either $\sum_i \mu(A_i)^- < \infty$ or $\sum_i \mu(A_i)^+ < \infty$, then also (ii) $\mu(A) = \sum_i \mu(A_i)$.

Proposition 4.1 *If μ is a signed measure on (Ω, Σ) that is absolutely continuous with respect to a σ-finite measure ν, then there exists a function (density) $f : \Omega \to \mathbb{R}$ such that*

$$\mu(A) = \int_A f(x)\, \mathrm{d}\nu(x)$$

for all $A \in \Sigma$.

This result is a consequence of the Jordan decomposition and the Radon Nikodym theorem. See, for example, Section 32 in Billingsley (1995) or Appendix A.8 in Durrett (2005) for details.

Definition 4.2 *A (univariate) time series kernel μ is a signed measure on $(\mathbb{R} \times \mathbb{R}_+, \mathcal{B} \otimes \mathcal{B}_+)$*[6] *with*

$$\int_0^\infty |\, \mathrm{d}\mu(f(s), s)| < \infty$$

for all bounded piecewise linear functions $f : \mathbb{R}_+ \to \mathbb{R}$, where the integration is over the time variable s.

The integrability condition ensures that the value of a convolution is well-defined and finite. In particular, the condition is satisfied if $\mathrm{d}\mu(x, s) = g(x)\, \mathrm{d}\mu_T(s)$ for some real function g and finite signed measure μ_T on $\mathbb{R}_+$, which is indeed the case throughout this paper.

Definition 4.3 (Convolution Operator) *For a time series $X \in \mathcal{T}$, kernel μ, and sampling scheme $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$, the convolution $*^\mu_\sigma(X) = X *_\sigma \mu$ is a time series with*

$$T(X *_\sigma \mu) = T(X), \tag{4.8}$$

$$(X *_\sigma \mu)_t = \int_0^\infty \mathrm{d}\mu(X[t-s]_\sigma, s), \qquad t \in T(X *_\sigma \mu). \tag{4.9}$$

If μ is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R} \times \mathbb{R}_+$, then (4.9) can be written as

$$(X *_\sigma \mu)_t = \int_0^\infty f(X[t-s]_\sigma, s)\, \mathrm{d}s, \qquad t \in T(X *_\sigma \mu), \tag{4.10}$$

where f is the density function of μ.

[6] $\mathcal{B} \otimes \mathcal{B}_+$ is the Borel σ-algebra on $\mathbb{R} \times \mathbb{R}_+$.

Remark 4.4 (Discrete Time Analog) *The discrete-time analogs to convolution operators are additive (but generally nonlinear), causal, time-invariant filters of the form*

$$Y_n = \sum_{k=0}^{\infty} f_k(X_{n-k}),$$

where $(X_n : n \in \mathbb{Z})$ is a stationary time series, and $f_0, f_1, \ldots$ are real-valued functions subject to certain integrability conditions. As discussed in Section 6, the representation simplifies further in the case of linearity.

In the remainder of the paper, we primarily focus on last-point sampling and therefore often omit the "σ" from the notation of a convolution operator. All of our results also hold for the other two sampling schemes, but we do not always provide separate derivations.

Proposition 4.5 *The convolution operator $*^{\mu}$, associated with a (univariate) time series kernel μ, is tick invariant, causal, and shift invariant.*

Proof. Tick-invariance and causality follow immediately from the definition of a convolution operator. Shift-invariance can be shown either using Lemma 3.8 or directly, which we do here for illustrative purposes. To this end, let $X \in \mathcal{T}$ and $\tau > 0$. Using (4.8) twice we get

$$(*^{\mu})_T(\mathrm{L}(X,\tau)) = T(\mathrm{L}(X,\tau)) = T(X) + \tau = T(X * \mu) + \tau = (*^{\mu})_T(X) + \tau,$$

showing that $*^{\mu}$ satisfies (3.5). On the other hand, for $t \in T(*^{\mu}(\mathrm{L}(X,\tau))) = T(X) + \tau$,

$$\begin{aligned}(X * \mu)_{t-\tau} &= \int_0^{\infty} \mathrm{d}\mu(X[(t-\tau) - s]_{\mathrm{last}}, s) \\ &= \int_0^{\infty} \mathrm{d}\mu(\mathrm{L}(X,\tau)[t-s]_{\mathrm{last}}, s) \\ &= (\mathrm{L}(X,\tau) * \mu)_t,\end{aligned}$$

where we use Lemma 2.9 (iv) for the second equality. Hence, we also have $(*^{\mu})_V(X) = (*^{\mu})_V(\mathrm{L}(X,\tau))$ and $*^{\mu}$ is therefore shift invariant. ■

Not all time series operators that are tick invariant, causal, and shift invariant can, however, be expressed as convolution operators. For example, the operator that calculates the smallest observation value in a rolling time window (see Section 5.3) is not of this form.

As the following example illustrates, when a time series operator can be expressed as a convolution operator, the associated kernel is not unique.

Example 4.6 *Let* O *be the operator that sets all observation values of a time series equal to zero, that is, $T(\mathrm{O}(X)) = T(X)$ and $\mathrm{O}(X)_t = 0$ for all $t \in T(\mathrm{O}(X))$ and $X \in \mathcal{T}$. Then $\mathrm{O}(X) = *^{\mu}(X)$ for all $X \in \mathcal{T}$ for the following kernels:*

(i) $\mu(x,s) = 0$,

(ii) $\mu(x,s) = \mathbf{1}_{\{s \in N\}}$ *where $N \in \mathcal{B}_+$ is a Lebesgue null set, and*

(iii) $\mu(x,s) = \mathbf{1}_{\{f(s)=x\}}$ *where $f : \mathbb{R}_+ \to \mathbb{R}$ satisfies $\lambda(f^{-1}(\{x\})) = 0$ for all $x \in \mathbb{R}$.*[7]

When we say that the kernel of a convolution operator has a certain structure, we imply that one of the equivalent kernels is of that structure.

[7] In other words, f is a function that assumes each value in $\mathbb{R}$ at most on a null set of time points. Examples are $f(s) = \sin(s)$, $f(s) = \exp(s)$, $f(s) = s^2$.

5 Examples

This section gives examples of univariate time series operators that can be expressed either as convolution operators or as simple transformations of them.

5.1 Time Series Arithmetic

It is straightforward to extend the four basic arithmetic operations of mathematics - addition, subtraction, multiplication and division - to unevenly spaced time series.

Definition 5.1 (Arithmetic Operations) *For a time series $X \in \mathcal{T}$ and $c \in \mathbb{R}$, we call*

(i) $c + X$ *(or* $X + c$*) with* $T(c + X) = T(X)$ *and* $V(c + X) = (c + X_1, \ldots, c + X_{N(X)})$ *"the sum of* c *and* X*,"*

(ii) cX *(or* Xc*) with* $T(cX) = T(X)$ *and* $V(cX) = (cX_1, \ldots, cX_{N(X)})$ *"the product of* c *and* X*," and*

(iii) $1/X$ *with* $T(1/X) = T(X)$ *and* $V(1/X) = (1/X_1, \ldots, 1/X_{N(X)})$ *"the inverse of* X*," provided that all observation values of* X *are non-zero.*

Definition 5.2 (Arithmetic Time Series Operations) *For time series $X, Y \in \mathcal{T}$ and sampling scheme $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$, we call*

(i) $X +_\sigma Y$ *with* $T(X +_\sigma Y) = T(X) \cup T(Y)$ *and* $(X +_\sigma Y)_t = X[t]_\sigma + Y[t]_\sigma$ *for* $t \in T(X +_\sigma Y)$ *"the sum of* X *and* Y*," and*

(ii) $X_\sigma Y$ *with* $T(X_\sigma Y) = T(X) \cup T(Y)$ *and* $(X_\sigma Y)_t = X[t]_\sigma Y[t]_\sigma$ *for* $t \in T(X_\sigma Y)$ *"the product of* X *and* Y*,"*

where $T_X \cup T_Y$ *for* $T_X, T_Y \in \mathbb{T}$ *denotes the sorted union of* T_X *and* T_Y*.*

Proposition 5.3 *The arithmetic operators in Definition 5.1 are convolution operators with kernel* $\mu(x, s) = (c + x)\delta_0(s)$*,* $\mu(x, s) = cx\delta_0(s)$*, and* $\mu(x, s) = \delta_0(s)/x$*, respectively.*

Proof. For $X \in \mathcal{T}$, $c \in \mathbb{R}$, and $\mu(x, s) = (c + x)\delta_0(s)$, by definition, $T(X * \mu) = T(X) = T(c + X)$. For $t \in T(X * \mu)$,

$$\begin{aligned}(X * \mu)_t &= \int_0^\infty \mathrm{d}\mu(X[t - s]_{\text{last}}, s) \\ &= \int_0^\infty (c + X[t - s]_{\text{last}})\delta_0(s)\,\mathrm{d}s \\ &= c + X_t,\end{aligned}$$

and, therefore, also $V(X * \mu) = V(c + X)$. The reasoning for the other two kernels is similar. ■

Proposition 5.4 *The sampling operator is linear, that is,*

$$(aX +_\sigma bY)[t]_\sigma = aX[t]_\sigma + bY[t]_\sigma$$

for time series $X, Y \in \mathcal{T}$*, sampling scheme* σ*, and* $a, b \in \mathbb{R}$*.*

Proof. The result follows directly from the definition of the corresponding arithmetic operations. For example, in the case of last-point sampling,

$$\begin{aligned}
(aX + bY)[t]_{\text{last}} &= (aX + bY)_{\text{Prev}(T(aX+bY),t)} \\
&= ((aX + (bY))_{\text{Prev}(T(X)\cup T(X),t)} \\
&= (aX)[\text{Prev}(T(X) \cup T(X), t)]_{\text{last}} + (bY)[\text{Prev}(T(X) \cup T(X), t)]_{\text{last}} \\
&= aX[\text{Prev}(T(X) \cup T(X), t)]_{\text{last}} + bY[\text{Prev}(T(X) \cup T(X), t)]_{\text{last}} \\
&= aX_{\text{Prev}^X(t)} + bY_{\text{Prev}^Y(t)} \\
&= aX[t]_{\text{last}} + bY[t]_{\text{last}},
\end{aligned}$$

where we used

$$\text{Prev}^X(\text{Prev}(T(X) \cup T(X), t)) = \text{Prev}^X(t)$$

for the second to last equality. ■

Note that in general, $1/X[t]_{\text{linear}}$ does not equal $(1/X)[t]_{\text{linear}}$ for $X \in \mathcal{T}$ and $t \notin T(X)$, although equality holds for the other two sampling schemes. Similarly, $(X_\sigma Y)[t]_{\text{linear}}$ usually does not equal $X[t]_{\text{linear}}Y[t]_{\text{linear}}$ for $X, Y \in \mathcal{T}$ and $t \notin T(X)$, but equality holds for the other two sampling schemes.

What does "the sum of X and Y" actually mean for two time series X and Y that are not synchronized; or what would we want it to mean? If X and Y are discretely observed realizations of continuous-time stochastic processes X^c and Y^c, respectively, we might want $(X+Y)[t]_{\text{last}}$ to be a "good guess" of $(X^c+Y^c)_t$, given all available information. The following example describes two settings where this desirable property holds.

Example 5.5 *Let X^c and Y^c be two independent continuous-time stochastic processes, $T_X, T_Y \in \mathbb{T}$ be fixed observation time sequences, and $X = X^c[T_X]$ and $Y = Y^c[T_Y]$ be the corresponding observation time series of X^c and Y^c, respectively. Furthermore, let $\mathcal{F}_t = \sigma(X_s : s \le t)$ and $\mathcal{G}_t = \sigma(Y_s : s \le t)$ denote the filtration generated by X and Y, respectively, and $\mathcal{H}_t = \mathcal{F}_t \cup \mathcal{G}_t$.*

(i) If X and Y are martingales, then

$$\text{E}((X^c + Y^c)_t|\mathcal{H}_t) = (X + Y)[t]_{\text{last}} \tag{5.11}$$

for all $t \ge \max(\min T_X, \min T_Y)$, that is, for time points for which both X and Y have at least one past observation.

(ii) If X^c and Y^c are Lévy processes (but not necessarily martingales),[8] *then*

$$\text{E}((X^c + Y^c)_t|\mathcal{H}_\infty) = X[t]_{\text{linear}} + Y[t]_{\text{linear}} = (X +_{\text{linear}} Y)[t] \tag{5.12}$$

for all $t \in \mathbb{R}$.[9]

[8] A Lévy process has independent and stationary increments, and is continuous in probability. Brownian motion and homogenous Poisson processes are special cases. See Sato (1999) and Protter (2005) for details.

[9] If X^c and Y^c are martingales without stationary increments, then hardly anything can be said about the conditional expectation on the left-hand side of (5.12) without additional information about X^c and Y^c. The LHS could even be larger than the largest observation value of $X +_{\text{linear}} Y$. Details are available from the author upon request.

Proof. If X^c and Y^c are independent martingales, then

$$\begin{aligned}
\mathrm{E}((X^c + Y^c)_t|\mathcal{H}_t) &= \mathrm{E}(X^c_t|\mathcal{H}_t) + \mathrm{E}(Y^c_t|\mathcal{H}_t) \\
&= \mathrm{E}(X^c_t|\mathcal{F}_t) + \mathrm{E}(Y^c_t|\mathcal{G}_t) \\
&= \mathrm{E}(X^c_t|\mathcal{F}_{\mathrm{Prev}^X(t)}) + \mathrm{E}(Y^c_t|\mathcal{G}_{\mathrm{Prev}^Y(t)}) \\
&= X^c_{\mathrm{Prev}^X(t)} + Y^c_{\mathrm{Prev}^Y(t)} \\
&= X[t]_{\mathrm{last}} + Y[t]_{\mathrm{last}} \\
&= (X+Y)[t]_{\mathrm{last}},
\end{aligned}$$

showing (5.11). For a Lévy process Z and times $s \leq t \leq r$, $Z_r - Z_s$ has an infinitely divisible distribution, and the conditional expectation $\mathrm{E}(Z_t|Z_s, Z_r)$ is therefore the linear interpolation of (s, Z_s) and (r, Z_r) evaluated at time t. Hence, (5.12) follows from similar reasoning as (5.11). ■

Apart from the four basic arithmetic operations, other scalar mathematical transformations can also be extended to unevenly spaced time series. For $X \in \mathcal{T}$ and function $f : \mathbb{R} \to \mathbb{R}$, we call $f(X)$ the time series that results when applying the function f to each observation value of X. For example,

$$\exp(X) = ((t_n, \exp(X_n)) : 1 \leq n \leq N(X)).$$

Elementwise time series operators are convolution operators with kernel $\mu(x, s) = f(x)\delta_0(s)$.

5.2 Return Calculations

In many applications, it is of interest to either analyze or report the change of a time series over a fixed time horizon; such as one day, month, or year. This section examines how to calculate such time series returns.

Definition 5.6 *For time series $X, Y \in \mathcal{T}$, we call*

(i) $\Delta^k X = \Delta(\Delta^{k-1}X)$ *for* $k \in \mathbb{N}$ *with* $\Delta X = ((t_n, X_n - X_{n-1}) : 1 \leq n \leq N(X) - 1)$ *for* $k > 0$ *and* $\Delta^0 X = X$ *the k-th order difference time series (of X), and*

(ii) $\mathrm{diff}_\gamma(X, Y)$ *with scale* $\gamma \in \{\mathrm{abs}, \mathrm{rel}, \log\}$ *the absolute/relative/log difference between X and Y, where*

$$\mathrm{diff}_\gamma(X, Y) = \begin{cases} X - Y & \text{if } \gamma = \mathrm{abs} \\ \frac{X}{Y} - 1 & \text{if } \gamma = \mathrm{rel} \\ \log\left(\frac{X}{Y}\right) & \text{if } \gamma = \log, \end{cases}$$

provided that X and Y are positive-valued time series for $\gamma \in \{\mathrm{rel}, \log\}$.[10]

Definition 5.7 (Returns) *For a time series $X \in \mathcal{T}$, time horizon $\tau > 0$, and return scale* $\gamma \in \{\mathrm{abs}, \mathrm{rel}, \log\}$, *we call*

(i) $\mathrm{ret}^{\mathrm{roll}}_\gamma(X, \tau) = \mathrm{diff}_\gamma(X, \mathrm{D}(X, \tau))$ *the rolling absolute/relative/log return time series (of X over horizon τ),*

[10] Again, we use an analogous definition for other sampling schemes. For example, $\mathrm{diff}_{\mathrm{abs},\mathrm{lin}}(X, Y) = X -_{\mathrm{linear}} Y$ denotes the absolute, linearly interpolated difference between X and Y.

(ii) $\mathrm{ret}_\gamma^{\mathrm{obs}}(X) = \mathrm{diff}_\gamma(X, \mathrm{B}(X))$ *the absolute/relative/log observation (or tick) return time series (of* X*),*

(iii) $\mathrm{ret}_\gamma^{\mathrm{SMA}_{\mathrm{last}}}(X, \tau) = \mathrm{diff}_\gamma(X, \mathrm{SMA}_{\mathrm{last}}(X, \tau))$ *the absolute/relative/log simple moving average (SMA) return time series (of* X *over horizon* τ*), and*

(iv) $\mathrm{ret}_\gamma^{\mathrm{EMA}_{\mathrm{last}}}(X, \tau) = \mathrm{diff}_\gamma(X, \mathrm{EMA}_{\mathrm{last}}(X, \tau))$ *the absolute/relative/log exponential moving average (EMA) return time series (of* X *over horizon* τ*),*

provided that X *is a positive-valued time series for* $\gamma \in \{\mathrm{rel}, \log\}$*. See Section 8 for details regarding the moving average operators* $\mathrm{SMA}_{\mathrm{last}}$ *and* $\mathrm{EMA}_{\mathrm{last}}$*.*

The interpretation of the first two return definitions is immediately clear, while Section 8.3 provides a motivation for the last two definitions.

Proposition 5.8 *The return operators* $\mathrm{ret}^{\mathrm{roll}}$*,* $\mathrm{ret}^{\mathrm{SMA}_{\mathrm{last}}}$*, and* $\mathrm{ret}^{\mathrm{EMA}_{\mathrm{last}}}$ *are either convolution operators or convolution operators combined with simple transformations.*

Proof. For a time series $X \in \mathcal{T}$ and time horizon $\tau > 0$, it is easy to see that

$$\mathrm{diff}_\gamma(X, \mathrm{D}(X, \tau)) = \begin{cases} X * \mu & \text{if } \gamma \in \{\mathrm{abs}, \log\}, \\ \exp(X * \mu) - 1 & \text{if } \gamma = \mathrm{rel}, \end{cases}$$

with

$$\mu(x, s) = \begin{cases} x(\delta_0(s) - \delta_\tau(s)) & \text{if } \gamma = \mathrm{abs}, \\ \log(x)(\delta_0(s) - \delta_\tau(s)) & \text{if } \gamma \in \{\mathrm{rel}, \log\}, \end{cases}$$

provided that X is a positive-valued time series for $\gamma \in \{\mathrm{rel}, \log\}$. Using their respective definitions in Section 8, the proofs for the other two return operators are similar. ■

5.3 Rolling Time Series Functions

In many cases, it is desirable to extract a certain piece of local information about a time series. Rolling time series functions allow us to do just that.

Definition 5.9 (Rolling Time Series Functions) *Assume that we are given a function* $f : \mathcal{T} \to \mathbb{R}$ *that is shift invariant in the sense that* $f(X) = f(\mathrm{L}(X, \eta))$ *for all* $\eta \in \mathbb{R}$ *and* $X \in \mathcal{T}$*. For a time series* $X \in \mathcal{T}$ *and time horizon* $\tau \in \mathbb{R}_+ \cup \{\infty\}$*, the "rolling function* f *of* X *over horizon* τ*," denoted by* $\mathrm{roll}(X, f, \tau)$*, is the time series with*

$$\begin{aligned} T(\mathrm{roll}(X, f, \tau)) &= T(X), \\ \mathrm{roll}(X, f, \tau)_t &= f(X\{t - \tau, t\}), \qquad t \in T(\mathrm{roll}(X, f, \tau)). \end{aligned}$$

Proposition 5.10 *The class of rolling time series operators is identical to the class of causal, shift-invariant, and tick-invariant operators.*

Proof.

$\Longrightarrow$ By construction, every rolling time series function is causal and tick-invariant. Furthermore, shift-invariance of f implies

$$\begin{aligned}\operatorname{roll}(X, f, \tau)_t &= f(X\{t-\tau, t\}) \\ &= f(\mathrm{L}(X, \eta)\{t+\eta-\tau, t+\eta\}) \\ &= \operatorname{roll}(\mathrm{L}(X, \eta), f, \tau)_{t+\eta},\end{aligned}$$

or equivalently, $V(\operatorname{roll}(X, f, \tau)) = V(\operatorname{roll}(\mathrm{L}(X, \eta), f, \tau))$ for all $X \in \mathcal{T}$ and $\eta \in \mathbb{R}$. Hence, $\operatorname{roll}(., f, \tau)$ is also shift invariant.

$\Longleftarrow$ If O is a causal, shift-invariant, and tick-invariant time series operator, then

$$\mathrm{O}(X)_t = \mathrm{O}(X\{-\infty, t\})_t.$$

Hence, if we define $f : \mathcal{T} \to \mathbb{R}$ as $f(X) = \mathrm{O}(X)_{\max T(X)}$, then we have $\mathrm{O}(X)_t = \operatorname{roll}(X, f, -\infty)_t$ for all $X \in \mathcal{T}$ and $t \in T(X)$. We are left to show that f is shift-invariant. Using Lemma 2.9 (ii) applied to $\mathrm{O}(X)$ and the shift-invariance of O,

$$\begin{aligned}f(X) &= \mathrm{O}(X)_{\max T(X)} \\ &= \mathrm{L}(\mathrm{O}(X), \eta)_{\max T(X)+\eta} \\ &= \mathrm{O}(\mathrm{L}(X, \eta))_{\max T(X)+\eta} \\ &= \mathrm{O}(\mathrm{L}(X, \eta))_{\max T(\mathrm{L}(X,\eta))} \\ &= f(\mathrm{L}(X, \eta))\end{aligned}$$

for all $\eta \in \mathbb{R}$ and $X \in \mathcal{T}$.

■

In particular, rolling time series functions include convolution operators. Many operators, however, that cannot be expressed as convolution operators are included as well:

Example 5.11 *For a time series $X \in \mathcal{T}$, horizon $\tau \in \mathbb{R}_+ \cup \{\infty\}$, and function $f : \mathcal{T} \to \mathbb{R}$, we call* $\operatorname{roll}(X, f, \tau)$ *with*

(i) $f(Y) = |V(Y)| = N(Y)$ *the rolling number of observations,*

(ii) $f(Y) = \sum_{i=1}^{N(Y)} V(Y)_i$ *the rolling sum,*

(iii) $f(Y) = \frac{1}{N(Y)} \sum_{i=1}^{N(Y)} V(Y)_i$ *the rolling average (or mean),*

(iv) $f(Y) = \max V(Y)$ *the rolling maximum (also denoted by* $\operatorname{rollmax}(X, \tau)$*),*

(v) $f(Y) = \min V(Y)$ *the rolling minimum (also denoted by* $\operatorname{rollmin}(X, \tau)$*),*

(vi) $f(Y) = \max V(Y) - \min V(Y)$ *the rolling range (also denoted by* $\operatorname{range}(X, \tau)$*), and*

(vii) $f(Y) = \frac{1}{N(Y)} \left|\{i : 1 \le i \le N(Y), V(Y)_i \le V(Y)_{N(Y)}\}\right|$ *the rolling quantile,*

of X over horizon $\tau > 0$.

6 Linear Time Series Operators

Linear operators (also called "linear systems'in the signal processing literature) are pervasive when working with vector spaces, and the space of unevenly spaced time series is no exception. To fully appreciate the structure of such operators, we need to take a closer look at the properties of time series sample paths.

6.1 Sample Paths

Definition 6.1 (Sample Path) *For a time series $X \in \mathcal{T}$ and sampling scheme $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$, the function $\text{SP}_\sigma(X) : \mathbb{R} \to \mathbb{R}$ with $\text{SP}_\sigma(X)(t) = X[t]_\sigma$ for $t \in \mathbb{R}$ is called the sample path of X (for sampling scheme σ). Furthermore,*

$$\text{SP}_\sigma = \{\text{SP}_\sigma(X) : X \in \mathcal{T}\}$$

is the space of sample paths, and $\text{SP}_\sigma(t)$ is the space of sample paths that are constant after time $t \in \mathbb{R}$.

In particular, $\text{SP}_\sigma(X) \in \text{SP}_\sigma(\max T(X))$ for $X \in \mathcal{T}$, because the sampled value of a time series is constant after its last observation.

Lemma 6.2 *The mapping of a time series to its sample path is linear, that is,*

$$\text{SP}_\sigma(aX +_\sigma bY) = a\,\text{SP}_\sigma(X) + b\,\text{SP}_\sigma(Y)$$

for all $X, Y \in \mathcal{T}$, each sampling scheme σ, and $a, b \in \mathbb{R}$.

Proof. Using the sample path definition and Proposition 5.4,

$$\begin{aligned}\text{SP}_\sigma(aX +_\sigma bY)(t) &= (aX +_\sigma bY)[t]_\sigma \\ &= aX[t]_\sigma + bY[t]_\sigma \\ &= a\,\text{SP}_\sigma(X)(t) + b\,\text{SP}_\sigma(Y)(t)\end{aligned}$$

for all $t \in \mathbb{R}$. ■

Lemma 6.3 *Fix a sampling scheme σ. Two time series $X, Y \in \mathcal{T}$ have the same sample path if and only if the observation values of $X -_\sigma Y$ are identical to zero.*

Proof. The result follows from Lemma 6.2 with $a = 1$ and $b = -1$. ■

The space of sample paths SP_σ (or $\text{SP}_\sigma(t)$ for fixed $t \in \mathbb{R}$) can be turned into a normed vector space, see Kreyszig (1989) or Kolmogorov and Fomin (1999). Specifically, given two elements $x, y \in \text{SP}_\sigma$ and number $a \in \mathbb{R}$, define $x + y$ and ax as the real functions with $(X + Y)(t) = x(t) + y(t)$ and $(aX)(t) = ax(t)$, respectively, for $t \in \mathbb{R}$. It is straightforward to verify that

$$\|x\|_{\text{SP}} = \max_{t \in \mathbb{R}} \|x(t)\|$$

for $x \in \text{SP}_\sigma$ defines a norm on SP_σ and $(\text{SP}_\sigma, \|\ \|_{\text{SP}})$ is therefore a normed vector space.

It is easy to see that

$$\max_{t\in T(X)} \|X_t\| = \|\operatorname{SP}_\sigma(X)\|_{\mathrm{SP}}$$

for all $X \in \mathcal{T}$ and each sampling scheme σ. Hence,

$$\|X\|_{\mathcal{T}} = \max_{t\in T(X)} |X_t|$$

defines a norm on $\mathcal{T}$, making $(\mathcal{T}, \|\ \|_{\mathcal{T}})$ a normed vector space also.[11]

Corollary 6.4 *For each sampling scheme σ, the mapping $X \to \operatorname{SP}_\sigma(X)$ is an isometry between $(\mathcal{T}, \|\ \|_{\mathcal{T}})$ and $(\operatorname{SP}_\sigma, \|\ \|_{\mathrm{SP}})$. In other words,*

$$\|X\|_{\mathcal{T}} = \|\operatorname{SP}_\sigma(X)\|_{\mathrm{SP}}$$

for all $X \in \mathcal{T}$.

Lemma 6.5 *The order of the lag operator* L *and the mapping of a time series to its sample path is interchangeable in the sense that*

$$\operatorname{SP}_\sigma(X)(t-\tau) = \operatorname{SP}_\sigma(\mathrm{L}(X,\tau))(t)$$

for all $X \in \mathcal{T}$ and $t, \tau \in \mathbb{R}$, and each sampling scheme σ.
Proof. *The result follows from Lemma 2.9 (iv).* ■

6.2 Bounded Linear Operators

Definition 6.6 *A time series operator* O *is linear for a sampling scheme σ if $\mathrm{O}(aX +_\sigma bY) = a\,\mathrm{O}(X) +_\sigma b\,\mathrm{O}(Y)$ for all $X, Y \in \mathcal{T}$ and $a, b \in \mathbb{R}$.*

In particular, linear time series operators are homogeneous of degree one, although the reverse is generally not true. For example, the operator that calculates the smallest observation value in a rolling time window (see Example 5.11) is homogeneous of degree one but not linear.

Definition 6.7 *A time series operator* O *is*

(i) bounded if there exists a constant $M < \infty$ such that

$$\|\mathrm{O}(X)\|_{\mathcal{T}} \le M\|X\|_{\mathcal{T}}$$

for all $X \in \mathcal{T}$, and

(ii) continuous (for sampling scheme σ) if for all $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$\|X -_\sigma Y\|_{\mathcal{T}} < \delta$$

for $X, Y \in \mathcal{T}$ implies

$$\|\mathrm{O}(X) -_\sigma \mathrm{O}(Y)\|_{\mathcal{T}} < \varepsilon.$$

[11] Strictly speaking, $\mathcal{T}$ is not a vector space because it has no unique zero element. If we consider two time series X and Y to be equivalent if their sample paths are identical, however, then the space of equivalence classes in $\mathcal{T}$ is a well-defined vector space. For our discussion, this distinction is not important, see Lemma 6.11, and we therefore do not distinguish between $\mathcal{T}$ and the space of equivalence classes in $\mathcal{T}$.

Proposition 6.8 *A linear time series operator* O *is bounded if and only if it is continuous.*

Proof. The equivalence of boundedness and continuity holds for any linear operator between two normed vector spaces; see Kreyszig (1989), Chapter 2.7 or Kolmogorov and Fomin (1999), §29. ■

For the remainder of this paper, we shall exclusively focus on bounded, and therefore continuous, linear operators.

Theorem 6.9 *If a time series kernel* μ *is of the form*

$$\mu(x, s) = x\mu_T(s), \tag{6.13}$$

where μ_T *is a finite signed measure on* $(\mathbb{R}_+, \mathcal{B}_+)$*, then for each sampling scheme* σ *the associated convolution operator* $*_\sigma^\mu$ *is a bounded linear operator.*

Proof. First note

$$\begin{aligned}
T((aX +_\sigma bY) *_\sigma \mu) &= T(aX +_\sigma bY) \\
&= T(aX) \cup T(bY) \\
&= T(X) \cup T(X) \\
&= T(X *_\sigma \mu) \cup T(Y *_\sigma \mu) \\
&= T(a(X *_\sigma \mu)) \cup T(b(Y *_\sigma \mu)) \\
&= T(a(X *_\sigma \mu) +_\sigma b(Y *_\sigma \mu))
\end{aligned}$$

because convolution and arithmetic operators are tick invariant. For $t \in T((aX +_\sigma bY) *_\sigma \mu)$,

$$\begin{aligned}
((aX +_\sigma bY) *_\sigma \mu)_t &= \int_0^\infty \mathrm{d}\mu((aX +_\sigma bY)[t - s]_\sigma, s) \\
&= \int_0^\infty \mathrm{d}\mu(aX[t - s]_\sigma + bY[t - s]_\sigma, s) \\
&= \int_0^\infty (aX[t - s]_\sigma + bY[t - s]_\sigma)\, \mathrm{d}\mu_T(s) \\
&= a\int_0^\infty X[t - s]_\sigma\, \mathrm{d}\mu_T(s) + b\int_0^\infty Y[t - s]_\sigma\, \mathrm{d}\mu_T(s) \\
&= a(X *_\sigma \mu)_t + b(Y *_\sigma \mu)_t,
\end{aligned}$$

showing that $*_\sigma^\mu$ is indeed linear. Furthermore,

$$\begin{aligned}
|(X *_\sigma \mu)_t| &= \left|\int_0^\infty X[t - s]_\sigma\, \mathrm{d}\mu_T(s)\right| \\
&\le \int_0^\infty |X[t - s]_\sigma|\, |\mathrm{d}\mu_T(s)| \\
&\le \|X\|_{\mathcal{T}} \int_0^\infty |\,\mathrm{d}\mu_T(s)| \\
&= \|X\|_{\mathcal{T}} \|\mu_T\|_{\mathrm{TV}},
\end{aligned} \tag{6.14}$$

for all $t \in T(X)$, where $\|\mu_T\|_{\mathrm{TV}}$ is the total variation of the signed measure μ_T, which is finite by assumption. Taking the maximum over all $t \in T(X *_\sigma \mu)$ on the left-hand side of (6.14) gives

$$\| *_\sigma^\mu (X)\|_{\mathcal{T}} \leq \|X\|_{\mathcal{T}} \cdot \|\mu_T\|_{\mathrm{TV}},$$

which shows the boundedness of $*_\sigma^\mu$. ■

The next subsection shows that, under reasonable conditions, the reverse is also true. In other words, convolution operators with linear kernel of the form (6.13) are the only interesting bounded linear operators. To show this result, we need to take a closer look at how individual time series observations are used by a linear convolution operator of the form (6.13).

Remark 6.10 *Assume we are given a time series $X \in \mathcal{T}$ and a linear convolution operator of form (6.13). Define $t_0 = -\infty$ and $X_{t_0} = X_{t_1}$. For each observation time $t = t_n \in T(X)$,*

$$(X * \mu)_{t_n} = \mu_T(\{0\})X_{t_n} + \sum_{k=1}^{n} \mu_T((t_n - t_k, t_n - t_{k-1}])X_{t_{k-1}}$$

for last-point sampling,

$$(X *_{\mathrm{next}} \mu)_{t_n} = \sum_{k=1}^{n} \mu_T([t_n - t_k, t_n - t_{k-1}))X_{t_k}$$

for next-point sampling, and

$$(X *_{\mathrm{linear}} \mu)_{t_n} = \sum_{k=1}^{n} a_{k,n} X_{t_k}$$

for sampling with linear interpolation, where the coefficients $a_{k,n}$ for $1 \leq k \leq n$ depend only on μ_T and the observation time spacings.

6.3 Linear Operators as Convolution Operators

Clearly, a time series contains all of the information about its sample path. On the other hand, the sample path of a time series contains only a subset of the time series information, because the observation times are not uniquely determined by the sample path alone. The following result shows that linear time series operators use only this reduced amount of information about a time series.

Lemma 6.11 *Let O be a linear time series operator for sampling scheme σ. There exists a function $g_\sigma : \mathrm{SP}_\sigma \to \mathrm{SP}_\sigma$ such that $\mathrm{SP}_\sigma(\mathrm{O}(X)) = g_\sigma(\mathrm{SP}_\sigma(X))$ for all $X \in \mathcal{T}$. In other words, for linear operators the sample path of the output time series depends only on the sample path of the input time series.*

Proof. For each sample path $x \in \mathrm{SP}_\sigma$, we choose[12] one (of the infinitely many) time series $X \in \mathcal{T}$ with $\mathrm{SP}_\sigma(X) = x$ and define

$$g_\sigma(x) = \mathrm{SP}_\sigma(\mathrm{O}(X)).$$

[12] Given a sample path, a matching time series can be constructed by looking at the kinks (for linear interpolation) or jumps (for the other sampling schemes) of the sample path.

We need to show that g_σ is uniquely defined for each $x \in \mathrm{SP}_\sigma$. Assume there exist two time series $X, Y \in \mathcal{T}$ with $\mathrm{SP}_\sigma(X) = \mathrm{SP}_\sigma(Y)$, but $\mathrm{SP}_\sigma(\mathrm{O}(X)) \neq \mathrm{SP}_\sigma(\mathrm{O}(Y))$. Because O is linear,

$$a\,\mathrm{O}(X -_\sigma Y) = \mathrm{O}(a(X -_\sigma Y)) \tag{6.15}$$

for all $a \in \mathbb{R}$. Lemma 6.3 implies that the observation values of $X -_\sigma Y$ (and therefore also $a(X -_\sigma Y)$) are identical to zero. Hence, (6.15) can be satisfied for all $a \in \mathbb{R}$ only if the observation values of $\mathrm{O}(X -_\sigma Y)$ (and therefore also $\mathrm{O}(X) -_\sigma \mathrm{O}(Y)$) are identical to zero. Applying Lemma 6.3 once more shows that $\mathrm{O}(X)$ and $\mathrm{O}(Y)$ have the same sample path, in contradiction to our assumption. ■

A lot more can be said about linear operators that satisfy additional, quite general, properties.

Theorem 6.12 *Let* O *be a bounded, causal, shift- and tick-invariant time series operator that is linear for sampling scheme* σ*. Then* O *is a convolution operator with kernel of the form* $\mu(x, s) = x\mu_T(s)$*, where* μ_T *is a finite signed measure on* $(\mathbb{R}_+, \mathcal{B}_+)$*.*

Proof. See Appendix A. ■

Combining Proposition 4.5, and Theorems 6.9 and 6.12 yields the main result of this section:

Theorem 6.13 *The class of bounded, linear, causal, shift-and tick-invariant time series operators coincides with the class of convolution operators with kernels of the form* $\mu(x, s) = x\mu_T(s)$ *where* μ_T *is a finite signed measure on* $(\mathbb{R}_+, \mathcal{B}_+)$*.*

Remark 6.14 *The class of bounded, linear, shift- and tick-invariant (but not necessarily causal) time series operators coincides with the class of convolution operators with kernel of the form* $\mu(x, s) = x\mu_T(s)$ *where* μ_T *is a finite signed measure on* $(\mathbb{R}, \mathcal{B})$*.*

Using these results, the discrete-time analogs of convolution operators derived in Remark 4.4 can be further simplified for linear operators.

Remark 6.15 (Discrete Time Analog - continued) *The discrete-time analog of linear convolution operators are causal, linear, shift-invariant filters of the form*

$$Y_n = \sum_{k=0}^{\infty} a_k X_{n-k},$$

with $\sum_{k=0}^{\infty} |a_k| < \infty$*, where* $(X_n : n \in \mathbb{Z})$ *is a stationary time series. The condition that the filter coefficients are absolutely convergent implies that the transformed time series is stationary if the input time series is stationary; see Brockwell and Davis (1991), Proposition 3.2.1. The coefficients of the discrete-time filter are related to the time series kernel* (6.13) *via the equations given in Remark 6.10.*

6.4 An Application

Definition 6.16 *We say that a time series operator* O *has finite memory if there exists a constant* $C < \infty$ *such that*

$$\mathrm{O}(X)_t = \mathrm{O}(X\{\mathrm{Prev}^X(t - C), +\infty\})_t$$

for all $X \in \mathcal{T}$ *and* $t \in T(\mathrm{O}(X))$*.*

Corollary 6.17 *A rolling time series function over horizon $\tau < \infty$ is a time series operator with finite memory.*

Proof. The result follows immediately from Definition 5.9. ■

The following theorem shows that for linear convolution operators and a quite general class of data-generating processes, the operator applied to the corresponding *observation* time series provides the "best guess" of the operator applied to the unobserved, data-generating process.

Theorem 6.18 *Let $(X_t^c : t \geq 0)$ be a Lévy process, $T_X \in \mathbb{T}$ with $\min T_X \geq 0$ be a fixed observation time sequence, and $X = X^c[T_X]$ be the corresponding observation time series. Let* O *be a linear convolution operator for sampling scheme $\sigma = \text{linear}$. If* O *has finite memory, then*

(i) the associated kernel $\mu(x, s) = x\mu_T(s)$ satisfies $\mu_T \equiv 0$ on $(C, +\infty)$ for some constant $C < \infty$,

(ii) and

$$\mathrm{E}((X^c * \mu)_t | X) = (\mathrm{SP}_{\text{linear}}(X) * \mu_T)(t)$$

for $C + \min T(X) \leq t \leq \max T(X)$. In particular

$$\mathrm{E}((X^c * \mu))_t | X) = \mathrm{O}(X)_t \tag{6.16}$$

for all $t \in T(X)$ with $t \geq C + \min T(X)$.

Proof.

(i) By the definition of a linear convolution operator,

$$\begin{aligned}
\int_0^\infty X[t-s]_{\text{linear}} \, \mathrm{d}\mu_T(s) &= \mathrm{O}(X)_t \\
&= \mathrm{O}(X\{\mathrm{Prev}^X(t-C), +\infty\})_t \\
&= \int_0^{t-\mathrm{Prev}^X(t-C)} X[t-s]_{\text{linear}} \, \mathrm{d}\mu_T(s) \\
&= \int_{[0,C]} X[t-s]_{\text{linear}} \, \mathrm{d}\mu_T(s) + \int_{(C,t-\mathrm{Prev}^X(t-C)]} X[t-s]_{\text{linear}} \, \mathrm{d}\mu_T(s)
\end{aligned}$$

for all $X \in \mathcal{T}$ and $t \in T(X)$, where we used the finite memory property for the second equation. Equivalently,

$$\int_{(C,\infty)} X[t-s]_{\text{linear}} \, \mathrm{d}\mu_T(s) = \int_{(C,t-\mathrm{Prev}^X(t-C)]} X[t-s]_{\text{linear}} \, \mathrm{d}\mu_T(s) \tag{6.17}$$

for all $X \in \mathcal{T}$ and $t \in T(X)$. For each time series $X \in \mathcal{T}$ and time point $t \in T(X)$, there exists a time series Y with identical sample path and $t - C \in T(X)$, which implies that $t - \mathrm{Prev}^Y(t-C) = C$. Hence, the right-hand side of (6.17) is always equal to zero, which can only be the case if μ_T is identical to zero on $(C, +\infty)$.

(ii) For a Lévy process Z, the conditional expectation $\mathrm{E}(Z_t|Z_s, Z_r)$ for $s \le t \le r$ is the linear interpolation of (s, Z_s) and (r, Z_r) evaluated at time t. Hence, for $C + \min T(X) \le t \le \max T(X)$,

$$\begin{aligned}
\mathrm{E}((X^c * \mu))_t|X) &= \mathrm{E}\left(\int_0^\infty \mathrm{d}\mu(X^c_{t-s}, s)\Big|X\right) \\
&= \mathrm{E}\left(\int_0^C X^c_{t-s}\,\mathrm{d}\mu_T(s)\,\mathrm{d}s\Big|X\right) \\
&= \int_0^C \mathrm{E}(X^c_{t-s}|X)\,\mathrm{d}\mu_T(s) \\
&= \int_0^C X[t-s]_{\text{linear}}\,\mathrm{d}\mu_T(s) \\
&= (\mathrm{SP}_{\text{linear}}(X) * \mu_T)(t)
\end{aligned} \tag{6.18}$$

where Fubini's theorem was used to change the order of integration and the conditional expectation. Additionally, if $t \in T(X)$, then (6.18) simplifies further to

$$\begin{aligned}
\mathrm{E}((X^c * \mu))_t|X) &= \int_0^C \mathrm{d}\mu(X[t-s]_{\text{linear}}, s) \\
&= \mathrm{O}(X)_t.
\end{aligned}$$

■

However, (6.16) generally does not hold for nonlinear time series operators such as volatility or correlation estimators. In these cases, the right-hand side of (6.16) requires correction terms that depend on the exact dynamics of the data-generating process.

7 Multivariate Time Series

Multivariate data sets often consist of time series with different frequencies, and are thus naturally treated as multivariate unevenly spaced time series, even if the observations of the individual time series are reported at regular intervals. For example, macroeconomic data; such as the gross domestic product (GDP), the rate of unemployment, and foreign exchange rates; is released in a nonsynchronous manner and at vastly different frequencies (quarterly, monthly, and essentially continuously, in case of the US). Moreover, the frequency of reporting may change over time.

Many multivariate time series operators are natural extensions of their univariate counterparts, so the concepts of the prior sections require only a few modifications. If one is primarily interested in the analysis of univariate time series, then this section may be skipped on a first reading without detriment to the understanding of the rest of the paper.

Definition 7.1 *For $K \ge 1$, a K-dimensional unevenly spaced time series X^K is a K-tuple $(X^K_k : 1 \le k \le K)$ of univariate unevenly spaced time series $X^K_k \in \mathcal{T}$ for $1 \le k \le K$. $\mathcal{T}^K$ is the space of (real-valued) K-dimensional time series.*

Definition 7.2 *For $K, M \ge 1$, a (K, M)-dimensional time series operator is a mapping $\mathrm{O} : \mathcal{T}^K \to \mathcal{T}^M$, or equivalently, an M-tuple of mappings $\mathrm{O}^m : \mathcal{T}^K \to \mathcal{T}$ for $1 \le m \le M$.*

The following two operators are helpful for extracting basic information from such objects.

Definition 7.3 (Subset Selection) *For a multivariate time series $X^K \in \mathcal{T}^K$ and indices $(i_1, \ldots, i_M)$ with $1 \leq i_j \leq K$ for $j = 1, \ldots, M$, we call*

$$X^K(i_1, \ldots, i_M) = \begin{cases} (X^K_{i_1}, \ldots, X^K_{i_M}) & \text{if } M \geq 2, \\ X^K_{i_1} & \text{if } M = 1, \end{cases}$$

the subset (vector) time series of X^K (for the indices $(i_1, \ldots, i_M)$).

Definition 7.4 (Multivariate Sampling) *For a multivariate time series $X^K \in \mathcal{T}^K$, time vector $t^K \in \mathbb{R}^K$ and sampling scheme $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$, we call*

$$X^K[t^K]_\sigma = (X^K_k[t^K_k]_\sigma : 1 \leq k \leq K) \tag{7.19}$$

the sampled value (vector) of X^K at time (vector) t^K.

Unless stated otherwise, we apply univariate time series operators element-wise to a multivariate time series X^K. In other words, we assume that a univariate time series operator O, when applied to a multivariate time series, is replaced by its natural multivariate extension. For example, we interpret $X^K[t]_{\text{last}}$ for $t \in \mathbb{R}$ as $(X^K_k[t]_{\text{last}} : 1 \leq k \leq K)$ and call it the sampled value (vector) of X^K at time t. Similarly, $\mathrm{L}(X^K, \tau)$ for $\tau \in \mathbb{R}$ equals $(\mathrm{L}(X^K_k, \tau) : 1 \leq k \leq K)$, and so on. Of course, whenever there is a risk of confusion, we must explicitly define the extension of a univariate time series operator to the multivariate case.

7.1 Structure

In Section 3 we already took a detailed look at the common structural features among univariate time series operators. With very minor modifications (not shown here), that analysis is still relevant to the multivariate case; however, additional properties are worth considering now.

7.1.1 Dimension Invariance

For many multivariate operators, the dimension of the output time series is either equal to one (typically for data aggregations) or equal to the dimension of the input time series (typically for data transformations).

Definition 7.5 *A (K, M)-dimensional time series operator O^K is dimension invariant if $K = M$.*

In particular, if a vector time series operator O^K is the natural multivariate extension of a univariate operator O, then the former is, by construction, dimension invariant.

7.1.2 Permutation Invariance

Many multivariate time series operators have a certain symmetry in the sense that they assume no natural ordering among the input time series.

Definition 7.6 *A (K, K)-dimensional time series operator* O *is permutation invariant if*

$$\mathrm{O}(p(X^K)) = p(\mathrm{O}(X^K))$$

for all permutations $p : \mathcal{T}^K \to \mathcal{T}^K$ *and time series* $X^K \in \mathcal{T}^K$.

In particular, if a vector time series operator O^K is the natural multivariate extension of a univariate operator O, then the former is, by construction, permutation invariant.

7.1.3 Example

We will now end the discussion of structural features with a brief example.

Definition 7.7 (Merging) *For* $X, Y \in \mathcal{T}$, $X \cup Y$ *denotes the merged time series of* X *and* Y, *where* $T(X \cup Y) = T(X) \cup T(X)$ *and*

$$(X \cup Y)_t = \begin{cases} X_t & \text{if } t \in T(X) \\ Y_t & \text{if } t \notin T(X) \end{cases}$$

for $t \in T(X \cup Y)$.

In particular, if both time series have an observation at the same time point, then the observation value of the first time series takes precedence. Therefore, $X \cup Y$ and $Y \cup X$ are generally not equal unless the observation times of X and Y are disjoint.

The operator that merges two time series is causal, timescale invariant, tick invariant (in the sense that $T(X \cup Y) = T(X) \cup T(X)$), shift invariant (in the sense that $\mathrm{L}(X \cup Y, \tau) = \mathrm{L}(X, \tau) \cup \mathrm{L}(Y, \tau)$), homogeneous of degree $d = 1$ (in the sense that $(aX) \cup (aY) = a(X \cup Y)$), and neither dimension nor permutation invariant.

7.2 Multivariate Convolution Operators

Definition 7.8 *A $(K, 1)$-dimensional (or K-dimensional) time series kernel* μ *is a signed measure on* $(\mathbb{R}^K \times \mathbb{R}_+, \mathcal{B}^K \otimes \mathcal{B}_+)$ *with*

$$\int_0^\infty |\,\mathrm{d}\mu(f(s), s)| < \infty$$

for all bounded piecewise linear functions $f : \mathbb{R}_+^K \to \mathbb{R}$.[13]

As in the univariate case, the integrability condition ensures that the value of a convolution is well-defined and finite. In particular, the condition is satisfied if $\mathrm{d}\mu(x, s) = g(x)\,\mathrm{d}\mu_T(s)$ for some real function $g : \mathbb{R}^K \to \mathbb{R}$ and finite signed measure μ_T on $\mathbb{R}_+$.

[13] A more general definition is possible, where μ is a signed measure on $(\mathbb{R}^K \times (\mathbb{R}_+)^K, \mathcal{B}^K \otimes (\mathcal{B}_+)^K)$. For our purposes, however, the simpler definition is sufficient.

Definition 7.9 (Multivariate Convolution Operator) *For a multivariate time series $X^K \in \mathcal{T}^K$, K-dimensional kernel μ and sampling scheme σ, the convolution $*_\sigma^\mu(X^K) = X^K *_\sigma \mu$ is a univariate time series with*

$$T(X^K *_\sigma \mu) = \bigcup_{k=1}^{K} T(X_k^K), \tag{7.20}$$

$$(X^K *_\sigma \mu)_t = \int_0^\infty \mathrm{d}\mu(X^K[t-s]_\sigma, s), \qquad t \in T(X^K *_\sigma \mu). \tag{7.21}$$

If μ is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^K \times \mathbb{R}_+$, then (7.21) can be written as

$$(X^K *_\sigma f)_t = \int_0^\infty f(X^K[t-s]_\sigma, s)\,\mathrm{d}s, \qquad t \in T(X^K *_\sigma \mu),$$

where f is the density function of μ.

A K-dimensional convolution operator is a mapping of $\mathcal{T}^K \to \mathcal{T}$. More generally, a (K, M)-dimensional convolution operator is an M-tuple of K-dimensional convolution operators $(*_\sigma^{\mu_1}, \ldots, *_\sigma^{\mu_M})$.

Note that a (K, K)-dimensional convolution operator is generally not equivalent to K one-dimensional convolution operators. In the former case, each output time series depends on *all* input time series; whereas in the latter case, each output time series depends only on the corresponding input time series. In particular, the observation times of the output time series of a (K, K)-dimensional convolution operator are the union of observation times of all input time series, see (7.20).

7.3 Examples

This section gives examples of multivariate time series operators that can be expressed as multivariate convolution operators.

Proposition 7.10 *The arithmetic operators $X+Y$ and XY in Definition 5.2 are multivariate (specifically $(2,1)$-dimensional) convolution operators with kernel $\mu(x, y, s) = (x + y)\delta_0(s)$ and $\mu(x, y, s) = xy\delta_0(s)$, respectively, for the vector time series $(X, Y) \in \mathcal{T}^2$.*

Proof. For $X, Y \in \mathcal{T}$ and $\mu(x, y, s) = (x + y)\delta_0(s)$, we have $T((X, Y) * \mu) = T(X) \cup T(X)$ by the definition of a multivariate convolution. For $t \in T((X, Y) * \mu)$,

$$\begin{aligned}((X, Y) * \mu)_t &= \int_0^\infty (X[t-s]_{\text{last}} + Y[t-s]_{\text{last}})\delta_0(s)\,\mathrm{d}s \\ &= X[t]_{\text{last}} + Y[t]_{\text{last}},\end{aligned}$$

and also, therefore, $V((X, Y) * \mu) = V(X + Y)$. The reasoning for the multiplication of two time series is similar. ■

Definition 7.11 (Cross-sectional Operators) *For a function $f : \mathbb{R}^K \to \mathbb{R}$ and sampling scheme σ, the cross-sectional or (contemporaneous) time series operator $\mathrm{C}_\sigma(., f) : \mathcal{T}^K \to \mathcal{T}$ is given by*

$$T(\mathrm{C}_\sigma(X^K, f)) = \bigcup_{k=1}^{K} T(X_k^K),$$
$$(\mathrm{C}_\sigma(X^K, f))_t = f(X^K[t]_\sigma), \qquad t \in T(\mathrm{C}_\sigma(X^K, f))$$

for $X^K \in \mathcal{T}^K$.

It is easy to see that a cross-sectional time series operator $\mathrm{C}_\sigma(., f)$ is a K-dimensional convolution operator with kernel $\mu^f(x^K, s) = f(x^K)\delta_0(s)$.

Example 7.12 *For a multivariate time series $X^K \in \mathcal{T}^K$, we call $\mathrm{C}_\sigma(X^K, f)$ with*

(i) $f(x^K) = \mathrm{sum}(x^K)$ *the cross-sectional sum of X^K (also written $\mathrm{sum}_{\mathrm{C},\sigma}(X^K)$),*

(ii) $f(x^K) = \mathrm{avg}(x^K)$ *the cross-sectional average of X^K (also written $\mathrm{avg}_{\mathrm{C},\sigma}(X^K)$),*

(iii) $f(x^K) = \min(x^K)$ *the cross-sectional minimum of X^K (also written $\min_{\mathrm{C},\sigma}(X^K)$),*

(iv) $f(x^K) = \max(x^K)$ *the cross-sectional maximum of X^K (also written $\max_{\mathrm{C},\sigma}(X^K)$),*

(v) $f(x^K) = \max(x^K) - \min(x^K)$ *the cross-sectional range of X^K (also written $\mathrm{range}_{\mathrm{C},\sigma}(X^K)$), and*

(vi) $f(x^K) = \mathrm{quant}(x^K, q)$ *the cross-sectional q-quantile of X^K (also written $\mathrm{quant}_{\mathrm{C},\sigma}(X^K, q)$).*

It is easy to see that arithmetic and cross-sectional time series operators are consistent with each other. For example, $\mathrm{sum}_{\mathrm{C},\sigma}(X^K)$ equals $X_1^K + \ldots + X_K^K$ for all $K \geq 1$ and $X^K \in \mathcal{T}^K$.

Contemporaneous time series operators, such as the ones in Example 7.12, are useful for summarizing the behavior of high-dimensional time series. For example, a common question among economists is how the *distribution* of household income within a certain country changes over time. If X^K denotes the time series of individual incomes from a panel data set, $\mathrm{quant}_{\mathrm{C}}(X^K, 0.8)/\mathrm{quant}_{\mathrm{C}}(X^K, 0.2)$ is the time series of the inter-quintile income ratio. Similarly, in a medical study, the dispersion of a certain measurement across subjects as a function of time might be of interest.

8 Moving Averages

Moving averages - with exponentially declining weights or otherwise - are used to extract the average value of a time series over a certain time horizon; for example, to smooth out noisy observation values. Moving averages are very similar to kernel smoothing methods, see Wand and Jones (1994) and **?**, except that they use only past observations and are therefore causal time series operators.

For equally spaced time series data there is only one way of calculating simple moving averages (SMAs) and exponential moving averages (EMAs), and the properties of such linear filters are well understood. For unevenly spaced data, however, there are multiple alternatives (for example, due to the choice of the sampling scheme), all of which may be sensible depending on the underlying data-generating process of a given time series and the desired application.

Definition 8.1 *A convolution operator $*_\sigma^\mu$, associated with a kernel μ and sampling scheme $\sigma \in \{\text{last}, \text{next}, \text{linear}\}$, is a moving average operator if*

$$\mu(x, s) = x\mu_T(s) = xdF(s), \tag{8.22}$$

where μ_T is a probability measure on $(\mathbb{R}_+, \mathcal{B}_+)$ and F its distribution function. We call μ_T (and sometimes μ itself) a moving average kernel.

Based on the remarks following Definition 2.4, we know that the first observation value of a moving average is equal to the first observation value of the input time series. Moreover, the moving average of a constant time series is identical to the input time series.

Theorem 8.2 *Fix a sampling scheme σ and restrict attention to the set of causal, shift- and tick-invariant time series operators. The class of moving average operators coincides with the class of linear time series operators in the aforementioned set with (i) $\mathrm{O}(X) \geq 0$ for all $X \in \mathcal{T}$ with $X \geq 0$, and (ii) $\mathrm{O}(X) = X$ for all constant time series $X \in \mathcal{T}$.*

Proof.

$\Longrightarrow$ It immediately follows from Definition 8.1 and Proposition 4.5 that a moving average operator satisfies the specified conditions.

$\Longleftarrow$ By Theorem 6.13, the kernel associated with such a time series operator O is of the form $\mu(x, s) = x\mu_T(s)$ for some finite signed measure μ_T on $(\mathbb{R}_+, \mathcal{B}_+)$. Because $\mathrm{O}(X) \geq 0$ for all $X \in \mathcal{T}$ with $X \geq 0$, it follows from a simple measure-theoretic argument that μ_T is a positive measure. Furthermore, because $\mathrm{O}(X) = X$ for all constant time series $X \in \mathcal{T}$, it follows that

$$\int_0^\infty \mathrm{d}\mu_T(s) = 1,$$

showing that μ_T is a probability measure on $(\mathbb{R}_+, \mathcal{B}_+)$ and, therefore, that O is a moving average operator.

■

For a moving average kernel μ_T, associated cumulative distribution function F, and $X \in \mathcal{T}$,

$$\begin{array}{lll} \mathrm{MA}_{\text{last}}(X, \mu_T) & = \mathrm{MA}_{\text{last}}(X, F) & = X * \mu, \\ \mathrm{MA}_{\text{next}}(X, \mu_T) & = \mathrm{MA}_{\text{next}}(X, F) & = X *_{\text{next}} \mu, \\ \mathrm{MA}_{\text{linear}}(X, \mu_T) & = \mathrm{MA}_{\text{linear}}(X, F) & = X *_{\text{linear}} \mu \end{array} \tag{8.23}$$

with $\mu(x, s) = x\mu_T(s) = xdF(s)$ are three different moving averages of X. If X is non-decreasing, then

$$\mathrm{MA}_{\text{next}}(X, F)_t \geq \mathrm{MA}_{\text{linear}}(X, F)_t \geq \mathrm{MA}_{\text{last}}(X, F)_t \tag{8.24}$$

for all $t \in T(X)$, because $X[t]_{\text{next}} \geq X[t]_{\text{linear}} \geq X[t]_{\text{last}}$ for all $t \in \mathbb{R}$ for a non-decreasing time series.

8.1 Simple Moving Averages

Definition 8.3 *For time series $X \in \mathcal{T}$, we define three versions of the simple moving average (SMA) of length $\tau > 0$. For $t \in T(X)$,*

(i) $\mathrm{SMA}_{\mathrm{last}}(X,\tau)_t = \frac{1}{\tau}\int_0^\tau X[t-s]_{\mathrm{last}}\,\mathrm{d}s$,

(ii) $\mathrm{SMA}_{\mathrm{next}}(X,\tau)_t = \frac{1}{\tau}\int_0^\tau X[t-s]_{\mathrm{next}}\,\mathrm{d}s$,

(iii) $\mathrm{SMA}_{\mathrm{linear}}(X,\tau)_t = \frac{1}{\tau}\int_0^\tau X[t-s]_{\mathrm{linear}}\,\mathrm{d}s$,

where in all cases the observation times of the input and output time series are identical.

The simple moving averages $\mathrm{SMA}_{\mathrm{last}}$, $\mathrm{SMA}_{\mathrm{next}}$, and $\mathrm{SMA}_{\mathrm{linear}}$ are moving average operators in the sense of Definition 8.1. Specifically,

$$\begin{aligned}\mathrm{SMA}_{\mathrm{last}}(X,\tau) &= \mathrm{MA}_{\mathrm{last}}(X,\mu_T)\\ \mathrm{SMA}_{\mathrm{next}}(X,\tau) &= \mathrm{MA}_{\mathrm{next}}(X,\mu_T)\\ \mathrm{SMA}_{\mathrm{linear}}(X,\tau) &= \mathrm{MA}_{\mathrm{linear}}(X,\mu_T)\end{aligned}$$

with $\mu_T(t) = \frac{1}{\tau}\mathbf{1}_{\{0\le t\le\tau\}}$.

The first SMA can be used to analyze discrete observation *values*; for example, to calculate the average FED funds target rate[14] over the past three years. In such a case, it is desirable to weight each observation value by the amount of time it remained unchanged. The rolling average is ideal for analyzing discrete *events*; for example, calculating the average number of casualties per deadly car accident over the past twelve months, or determining the average number of IBM common shares traded on the NYSE per executed order during the past 30 minutes. The $\mathrm{SMA}_{\mathrm{linear}}$ can be used to estimate the rolling average value of a discretely-observed continuous-time stochastic processes, with observation times that are independent of the observation values, see Theorem 6.18. Finally, the $\mathrm{SMA}_{\mathrm{next}}$ is useful for certain trend and return calculations, see Section 8.3.

That said, the values of all SMAs will generally be quite similar as long as the moving average time horizon τ is considerably larger than the spacing of observation times. The type of moving average used (for example, SMA vs. EMA) and the moving average time horizon will usually have a much greater influence on the outcome of a time series analysis.

Proposition 8.4 *For all $X \in \mathcal{T}$,*

$$\lim_{\tau\searrow 0}\mathrm{SMA}_{\mathrm{linear}}(X,\tau)_t = \lim_{\tau\searrow 0}\mathrm{SMA}_{\mathrm{next}}(X,\tau)_t = X_t$$

for all $t \in T(X)$, while

$$\lim_{\tau\searrow 0}\mathrm{SP}(\mathrm{SMA}_{\mathrm{last}}(X,\tau))(t) = \mathrm{SP}(\mathrm{B}(X))(t)$$

for all $t \in \mathbb{R}$.

[14] The FED funds target rate is the desired interest rate (by the Federal Reserve) at which depository institutions (such as a savings bank) lend balances held at the Federal Reserve to other depository institutions overnight. See www.federalreserve.gov/fomc/fundsrate.htm for details.

We will now end our discussion of SMAs by illustrating a connection to the corresponding operator for equally spaced data.

Proposition 8.5 (Equally spaced time series) *If $X \in \mathcal{T}$ is an equally spaced time series with observation time spacings equal to some constant $c > 0$, and if the moving average length τ is an integer multiple of c, then*

$$\text{rolling\_avg}(X, \tau)_t = \text{SMA}_{\text{next}}(X, \tau)_t$$

for all $t \in T(X)$ with $t \geq \min T(X) + \tau - c$. In other words, the rolling average and SMA_{next} are identical after an initial ramp-up period of length $\tau - c$.

Proof. Because $\tau = K\Delta t$ for some integer K, for all $t \in T(X)$ with $t \geq \min T(X) + \tau - c$ we have

$$\begin{aligned}
\text{SMA}_{\text{next}}(X, \tau)_t &= \frac{1}{\tau} \int_0^\tau X[t-s]_{\text{next}} \, \mathrm{d}s \\
&= \frac{X_t \Delta t(X)_t + X_{t-\Delta t} \Delta t(X)_{t-\Delta t} + \ldots + X_{t-\Delta t(K-1)} \Delta t(X)_{t-\Delta t(K-1)}}{\tau} \\
&= \frac{X_t + X_{t-\Delta t} + \ldots + X_{t-\Delta t(K-1)}}{K} \\
&= \text{avg}\{X_s : s \in [t, t-\tau) \cap T(X)\} \\
&= \text{rolling_avg}(X, \tau)_t.
\end{aligned}$$

■

See Eckner (2017) for an efficient $O(N(X))$ implementation in the programming language `C` of simple moving averages and various other time series operators for unevenly spaced data.

8.2 Exponential Moving Averages

This section discusses exponential moving averages, also known as exponentially weighted moving averages. We use the former name in this paper, because the associated weights for unevenly spaced time series are defined only implicitly, via a kernel, as opposed to explicitly, as they are for equally spaced time series.

Definition 8.6 *For a time series $X \in \mathcal{T}$, we define three versions of the exponential moving average (EMA) of length $\tau > 0$. For $t \in \{t_1, \ldots, t_{N(X)}\}$,*

(i) $\text{EMA}_{\text{last}}(X, \tau)_t = \frac{1}{\tau} \int_0^\infty X[t-s]_{\text{last}} e^{-s/\tau} \, \mathrm{d}s$,

(ii) $\text{EMA}_{\text{next}}(X, \tau)_t = \frac{1}{\tau} \int_0^\infty X[t-s]_{\text{next}} e^{-s/\tau} \, \mathrm{d}s$,

(iii) $\text{EMA}_{\text{linear}}(X, \tau)_t = \frac{1}{\tau} \int_0^\infty X[t-s]_{\text{linear}} e^{-s/\tau} \, \mathrm{d}s$,

where in all cases the observation times of the input and output time series are identical.

The exponential moving averages EMA_{last}, EMA_{next}, $\text{EMA}_{\text{linear}}$ are moving average operators in the sense of Definition 8.1. Specifically,

$$\begin{aligned}
\text{EMA}_{\text{last}}(X, \tau) &= \text{MA}_{\text{last}}(X, \mu_T) \\
\text{EMA}_{\text{next}}(X, \tau) &= \text{MA}_{\text{next}}(X, \mu_T) \\
\text{EMA}_{\text{linear}}(X, \tau) &= \text{MA}_{\text{linear}}(X, \mu_T)
\end{aligned}$$

with $\mu_T(s) = \frac{1}{\tau} e^{-s/\tau}$.

Proposition 8.7 *For all $X \in \mathcal{T}$,*

$$\lim_{\tau \searrow 0} \mathrm{EMA}_{\mathrm{linear}}(X,\tau)_t = \lim_{\tau \searrow 0} \mathrm{EMA}_{\mathrm{next}}(X,\tau)_t = X_t$$

for all $t \in T(X)$, while

$$\lim_{\tau \searrow 0} \mathrm{SP}(\mathrm{EMA}_{\mathrm{last}}(X,\tau))(t) = \mathrm{SP}(\mathrm{B}(X))(t)$$

for all $t \in \mathbb{R}$.

As the following result shows, applying the usual EMA for equally spaced data to unevenly spaced data implicitly applies next-point interpolation to the time series, something that, depending on the application, may not be desirable.

Proposition 8.8 *For all $X \in \mathcal{T}$ and $\tau > 0$,*

$$\mathrm{EMA}_{\mathrm{next}}(X,\tau)_t = \begin{cases} X_{t_1} & \text{if } t = t_1, \\ (1 - e^{-\Delta t_n/\tau}) X_{t_n} + e^{-\Delta t_n/\tau}\, \mathrm{EMA}_{\mathrm{next}}(X,\tau)_{t_{n-1}} & \text{if } t = t_n > t_1 \end{cases}$$

for $t \in T(X)$.

Proof. $\mathrm{EMA}_{\mathrm{next}}(X,\tau)_{t_1} = X_{t_1}$, because $X[t_1 - s]_{\mathrm{next}} = X_{t_1}$ for all $s > 0$, showing the result for $n = 1$. For $1 < n \leq N(X)$:

$$\begin{aligned} \mathrm{EMA}_{\mathrm{next}}(X,\tau)_{t_n} &= \frac{1}{\tau} \int_0^\infty X[t_n - s]_{\mathrm{next}} e^{-s/\tau}\, \mathrm{d}s \\ &= \frac{1}{\tau} \int_0^{\Delta t_n} X[t_n - s]_{\mathrm{next}} e^{-s/\tau}\, \mathrm{d}s + \frac{1}{\tau} \int_{\Delta t_n}^\infty X[t_n - s]_{\mathrm{next}} e^{-s/\tau}\, \mathrm{d}s \\ &= X_{t_n}(1 - e^{-\Delta t_n/\tau}) \\ &\quad + e^{-\Delta t_n/\tau} \int_{\Delta t_n}^\infty X[(t_n - \Delta t_n) - (s - \Delta t_n)]_{\mathrm{next}} e^{-(s-\Delta t_n)/\tau}\, \mathrm{d}s \\ &= X_{t_n}(1 - e^{-\Delta t_n/\tau}) + e^{-\Delta t_n/\tau}\, \mathrm{EMA}_{\mathrm{next}}(X,\tau)_{t_{n-1}}. \end{aligned}$$

■

It is easy to show that $\mathrm{EMA}_{\mathrm{last}}$ can also be calculated recursively. Specifically,

$$\mathrm{EMA}_{\mathrm{last}}(X,\tau)_t = \begin{cases} X_{t_1} & \text{if } t = t_1, \\ (1 - e^{-\Delta t_n/\tau}) X_{t_{n-1}} + e^{-\Delta t_n/\tau}\, \mathrm{EMA}_{\mathrm{last}}(X,\tau)_{t_{n-1}} & \text{if } t = t_n > t_1 \end{cases}$$

for $t \in T(X)$, which is the $\mathrm{EMA}_{\mathrm{next}}$ of the original time series with lagged observation values.

Finally, Müller (1991) showed that the $\mathrm{EMA}_{\mathrm{linear}}$ can also be calculated recursively:

Proposition 8.9 *For $X \in \mathcal{T}$ and $\tau > 0$,*

$$\begin{aligned} \mathrm{EMA}_{\mathrm{linear}}(X,\tau)_{t_1} &= X_{t_1}, \\ \mathrm{EMA}_{\mathrm{linear}}(X,\tau)_{t_n} &= e^{-\Delta t_n/\tau}\, \mathrm{EMA}_{\mathrm{linear}}(X,\tau)_{t_{n-1}} + X_{t_n}(1 - \omega(\tau, \Delta t_n)) \\ &\quad + X_{t_{n-1}}(\omega(\tau, \Delta t_n) - e^{-\Delta t_n/\tau}), \end{aligned}$$

for $t_n \in T(X)$ with $n \geq 2$, where

$$\omega(\tau, \Delta t_n) = \frac{\tau}{\Delta t_n}(1 - e^{-\Delta t_n/\tau}).$$

In particular, $\omega(\tau, \Delta t_n) \approx 0$ for $\tau \ll \Delta t_n$ in which case $\mathrm{EMA}_{\mathrm{linear}}(X,\tau)_{t_n} \approx X_{t_n}$.

See Eckner (2017) for an efficient $\mathrm{O}(N(X))$ implementation in the `C` and `R` programming language of exponential moving averages and various other time series operators for unevenly spaced data.

8.3 Continuous-Time Analog

Calculating moving averages in *discrete* time requires us to choose a of the sampling scheme. In contrast, there is no such choice when calculating in *continuous* time, a fact which allows succinct illustration of certain concepts.

Definition 8.10 *Let $\widetilde{X} : \mathbb{R} \to \mathbb{R}$ be a function satisfying suitable integrability conditions. For $\tau > 0$ we call the function $\mathrm{SMA}_{\mathrm{last}}(\widetilde{X}, \tau) : \mathbb{R} \to \mathbb{R}$ defined by*

$$\mathrm{SMA}_{\mathrm{last}}(\widetilde{X}, \tau)_t = \frac{1}{\tau} \int_0^\tau \widetilde{X}_{t-s} \, \mathrm{d}s, \qquad t \in \mathbb{R},$$

the simple moving average of $\widetilde{X}$ (with length τ), and $\mathrm{EMA}_{\mathrm{last}}(\widetilde{X}, \tau) : \mathbb{R} \to \mathbb{R}$ defined by

$$\mathrm{EMA}_{\mathrm{last}}(\widetilde{X}, \tau)_t = \frac{1}{\tau} \int_0^\infty \widetilde{X}_{t-s} e^{-s/\tau} \, \mathrm{d}s, \qquad t \in \mathbb{R}, \tag{8.25}$$

the exponential moving average of $\widetilde{X}$ (with length τ).

Unless stated otherwise, we apply a time series operator to a real-valued function using its natural analog. For example, $\mathrm{L}(\widetilde{X}, \tau)$ for $\tau \in \mathbb{R}$ is the real-valued function with $\mathrm{L}(\widetilde{X}, \tau)_t = \widetilde{X}_{t-\tau}$ for all $t \in \mathbb{R}$. Of course, whenever there is a risk of confusion, we must explicitly define this extension.

With the SMA and EMA in continuous time established, we are ready to show several fundamental relationships between the trend measures and returns of a time series.

Theorem 8.11 *Let $\widetilde{X} : \mathbb{R} \to \mathbb{R}$ be a differentiable function satisfying suitable integrability conditions. For $\tau > 0$ and $t \in \mathbb{R}$,*

$$\frac{\widetilde{X}_t - \widetilde{X}_{t-\tau}}{\tau} = \mathrm{SMA}_{\mathrm{last}}(\widetilde{X}', \tau)_t \tag{8.26}$$

and

$$\frac{1}{\tau} \log\left(\frac{\widetilde{X}_t}{\widetilde{X}_{t-\tau}}\right) = \mathrm{SMA}_{\mathrm{last}}\left(\log(\widetilde{X})', \tau\right)_t, \tag{8.27}$$

if $\widetilde{X} > 0$, where $\widetilde{X}'$ denotes the derivative of $\widetilde{X}$.

Proof. By the fundamental theorem of calculus

$$\begin{aligned} \widetilde{X}_t - \widetilde{X}_{t-s} &= \int_0^s \widetilde{X}'_{t-z} \, \mathrm{d}z \\ \log(\widetilde{X}_t) - \log(\widetilde{X}_{t-s}) &= \int_0^s \log(\widetilde{X})'_{t-z} \, \mathrm{d}z \end{aligned}$$

for $s \in \mathbb{R}$. Hence, (8.26) and (8.27) follow from the definition of the simple moving average. ■

For an unevenly spaced time series $X \in \mathcal{T}$ (in discrete time),

$$\frac{1}{\tau}\operatorname{ret}^{\text{roll}}_{\text{abs}}(X,\tau)_t = \frac{X_t - X[t-\tau]}{\tau} \approx \operatorname{SMA}_{\text{next}}\left(\frac{\Delta X}{\Delta t(X)},\tau\right)_t$$

and

$$\frac{1}{\tau}\operatorname{ret}^{\text{roll}}_{\log}(X,\tau)_t = \frac{1}{\tau}\log\left(\frac{X_t}{X[t-\tau]}\right) \approx \operatorname{SMA}_{\text{next}}\left(\frac{\Delta \log(X)}{\Delta t(X)},\tau\right)$$

for $t \in T(X)$ and $\tau > 0$ can be used as an approximation of (8.26) and (8.27), respectively. In some cases, the approximation holds exactly. For example, if $t - \tau \in T(X)$, then $t = t_n$ and $t - \tau = t_{n-k}$ for some $1 \le k < n \le N(X)$, so that

$$\begin{aligned}\operatorname{SMA}_{\text{next}}\left(\frac{\Delta X}{\Delta t(X)},\tau\right)_t &= \frac{1}{\tau}\left(\frac{\Delta X_{t_n}\Delta t_n}{\Delta t_n} + \ldots + \frac{\Delta X_{t_{n-k+1}}\Delta t_{n-k+1}}{\Delta t_{n-k+1}}\right) \\ &= \frac{1}{\tau}(\Delta X_{t_n} + \Delta X_{t_{n-1}} + \ldots + \Delta X_{t_{n-k+1}}) \\ &= \frac{X_{t_n} - X_{t_n-\tau}}{\tau} \\ &= \frac{X_t - X[t-\tau]}{\tau}.\end{aligned} \tag{8.28}$$

A similar result holds for exponential moving averages

Theorem 8.12 *Let $\widetilde{X} : \mathbb{R} \to \mathbb{R}$ be a differentiable function satisfying suitable integrability conditions. For $\tau > 0$ and $t \in \mathbb{R}$,*

$$\frac{\widetilde{X}_t - \operatorname{EMA}_{\text{last}}(\widetilde{X},\tau)_t}{\tau} = \operatorname{EMA}_{\text{last}}(\widetilde{X}',\tau)_t \tag{8.29}$$

where $\widetilde{X}'$ denotes the derivative of $\widetilde{X}$.

Proof. The result can again be shown using the fundamental theorem of calculus, but the calculation is tedious. Alternatively, partial integration gives

$$\begin{aligned}\operatorname{EMA}_{\text{last}}(\widetilde{X}',\tau)_t &= \frac{1}{\tau}\int_0^\infty \widetilde{X}'_{t-z}e^{-z/\tau}dz \\ &= -\frac{1}{\tau}\widetilde{X}_{t-z}e^{-z/\tau}\Big|_{z=0}^{z=\infty} - \frac{1}{\tau^2}\int_0^\infty \widetilde{X}_{t-z}e^{-z/\tau}\,\mathrm{d}z \\ &= \frac{1}{\tau}\widetilde{X}_t - \frac{1}{\tau}\operatorname{EMA}_{\text{last}}(\widetilde{X},\tau)_t.\end{aligned}$$

■

For an unevenly spaced time series $X \in \mathcal{T}$ (in discrete time),

$$\frac{X_t - \operatorname{EMA}_{\text{last}}(X,\tau)_t}{\tau} \approx \operatorname{EMA}_{\text{next}}\left(\frac{\Delta X}{\Delta t(X)},\tau\right)_t$$

can be used as an approximation of (8.29).

Appendices

A Proof of Theorem 6.12

We proceed by breaking down the proof into three separate results.

Lemma A.1 *Let* O *be a shift-invariant time series operator that is linear for sampling scheme* σ*. There exists a function* $h_\sigma : \mathrm{SP}_\sigma \to \mathbb{R}$ *(as opposed to* $\mathrm{SP}_\sigma \to \mathrm{SP}_\sigma$ *in Lemma 6.11) such that*

$$\mathrm{SP}_\sigma(\mathrm{O}(X))(t) = h_\sigma(\mathrm{SP}_\sigma(\mathrm{L}(X, -t))) \tag{A.30}$$

for all time series $X \in \mathcal{T}$ *and all* $t \in \mathbb{R}$*. In other words, the sample path of the output time series depends only on the sample path of the input time series, and this dependence is shift-invariant.*

Proof. By Lemma 6.5 and the shift-invariance of O,

$$\begin{aligned}\mathrm{SP}_\sigma(\mathrm{O}(X))(t-\tau) &= \mathrm{SP}_\sigma(\mathrm{L}(\mathrm{O}(X), \tau))(t) \\ &= \mathrm{SP}_\sigma(\mathrm{O}(\mathrm{L}(X, \tau)))(t)\end{aligned}$$

for all $X \in \mathcal{T}$ and $t, \tau \in \mathbb{R}$. Setting $t = 0$ and $\tau = -t$ gives

$$\mathrm{SP}_\sigma(\mathrm{O}(X))(t) = \mathrm{SP}_\sigma(\mathrm{O}(\mathrm{L}(X, -t)))(0) \tag{A.31}$$

for all $t \in \mathbb{R}$. According to Lemma 6.11 there exists a function $g_\sigma : \mathrm{SP}_\sigma \to \mathrm{SP}_\sigma$ such that

$$\mathrm{SP}_\sigma(\mathrm{O}(\mathrm{L}(X, -t))) = g_\sigma(\mathrm{SP}_\sigma(\mathrm{L}(X, -t))). \tag{A.32}$$

Combining (A.31) and (A.32) yields

$$\mathrm{SP}_\sigma(\mathrm{O}(X))(t) = \mathrm{SP}_\sigma(\mathrm{O}(\mathrm{L}(X, -t)))(0) = g_\sigma(\mathrm{SP}_\sigma(\mathrm{L}(X, -t)))(0).$$

Hence, the desired function h_σ is given by $h_\sigma(X) = g_\sigma(x)(0)$ for $x \in \mathrm{SP}_\sigma$. ■

For operators that are bounded, even more can be said.

Theorem A.2 *Let* O *be a bounded and shift-invariant time series operator that is linear for sampling scheme* σ*. Then there exists a finite signed measure* μ_T *on* $(\mathbb{R}, \mathcal{B})$ *such that*

$$\mathrm{SP}_\sigma(\mathrm{O}(X))(t) = \int_{-\infty}^{+\infty} \mathrm{SP}_\sigma(X)(t-s)\,\mathrm{d}\mu_T(s) \tag{A.33}$$

for all $t \in \mathbb{R}$*, or equivalently*

$$\mathrm{SP}_\sigma(\mathrm{O}(X)) = \mathrm{SP}_\sigma(X) * \mu_T,$$

where "$*$*" denotes the convolution of a real function with a Borel measure.*

Proof. By Lemma A.1 there exists a function $h_\sigma : \mathrm{SP}_\sigma \to \mathbb{R}$ such that

$$\begin{aligned}\mathrm{SP}_\sigma(\mathrm{O}(X))(t) &= h_\sigma(\mathrm{SP}_\sigma(\mathrm{L}(X,-t))) \\ &= h_\sigma(\mathrm{SP}_\sigma(\widetilde{X})) \\ &= \mathrm{SP}_\sigma(\mathrm{O}(\widetilde{X}))(0),\end{aligned}$$

where $\widetilde{X} = \mathrm{L}(X,-t)$. Furthermore, by Lemma 6.5 (setting $t=-s$ and $\tau=-t$)

$$\mathrm{SP}_\sigma(X)(t-s) = \mathrm{SP}_\sigma(\widetilde{X})(-s).$$

Hence, it suffices to show (A.33) for $t=0$, or equivalently, that there exists a finite signed measure μ_T on $(\mathbb{R},\mathcal{B})$ such that

$$h_\sigma(\mathrm{SP}_\sigma(X)) = \int_{-\infty}^{+\infty} \mathrm{SP}_\sigma(X)(-s)\,\mathrm{d}\mu_T(s) \tag{A.34}$$

for all $X \in \mathcal{T}$.
Let us for a moment consider only the last-point and next-point sampling scheme. Define the set of indicator time series $I_{a,b} \in \mathcal{T}$ for $-\infty \le a \le b < \infty$ by

$$T(I_{a,b}) = \begin{cases} (a-1,a,b) & \text{if } a<b, \\ (A) & \text{if } a=b, \end{cases}$$

and

$$V(I_{a,b}) = \begin{cases} (0,1,0) & \text{if } a<b, \\ (0) & \text{if } a=b. \end{cases}$$

The sample path of an indicator time series is given $\mathrm{SP}_\sigma(I_{a,b})(t) = \mathbf{1}_{[a,b)}(t)$ for $t \in \mathbb{R}$, which gives these time series their name. For disjoint intervals $[a_i,b_i)$ with $a_i \le b_i$ for $i \in \mathbb{N}$, we define

$$\mu_T\left(\bigcup_i [a_i,b_i)\right) = h_\sigma\left(\mathrm{SP}_\sigma\left(\sum_i I_{(-b_i,-a_i)}\right)\right).$$

Because O and therefore h_σ are bounded,

$$\begin{aligned}\mu_T\left(\bigcup_i [a_i,b_i)\right) &= h_\sigma\left(\mathrm{SP}_\sigma\left(\sum_i I_{(-b_i,-a_i)}\right)\right) \\ &\le M\left\|\mathrm{SP}_\sigma\left(\sum_i I_{(-b_i,-a_i)}\right)\right\|_{\mathrm{SP}} \\ &= M < \infty.\end{aligned}$$

Combined with the linearity of O we have that μ_T is a finite countably additive set function on $(\mathbb{R},\mathcal{A})$ where

$$\mathcal{A} = \left\{\bigcup_i [a_i,b_i) : -\infty \le a_i \le b_i < \infty \text{ for } i \in \mathbb{N}\right\}.$$

By Carathéodory's extension theorem, there exists a unique extension of μ_T (for simplicity, also called μ_T) to $(\mathbb{R}, \mathcal{B})$. We are left to show that μ_T satisfies (A.34). To this end, note that the sample path of each time series $X \in \mathcal{T}$ can be written as a weighted sum of the sample paths of indicator time series. For example, in the case of last-point sampling,

$$\mathrm{SP}(X) = \sum_{i=0}^{N(X)} X_{t_i} \, \mathrm{SP}_\sigma(I_{t_i,t_{i+1}})$$

where we define $t_0 = -\infty$, $X_{t_0} = X_{t_1}$, and $t_{N(X)+1} = +\infty$. Using the linearity of O, and therefore h_σ,

$$\begin{aligned} h_\sigma(\mathrm{SP}(X)) &= h_\sigma \left(\sum_{i=0}^{N(X)} X_{t_i} \, \mathrm{SP}(I_{t_i,t_{i+1}}) \right) \\ &= \sum_{i=0}^{N(X)} X_{t_i} h_\sigma(\mathrm{SP}(I_{t_i,t_{i+1}})) \\ &= \sum_{i=0}^{N(X)} X_{t_i} \mu_T([-t_{i+1}, -t_i)) \\ &= \sum_{i=0}^{N(X)} X_{t_i} \mu_T([0 - t_{i+1}, 0 - t_i)). \end{aligned}$$

The last expression equals

$$\int_{-\infty}^{+\infty} X[0 - s]_{\text{last}} \, \mathrm{d}\mu_T(s) = (\mathrm{SP}(X) * \mu_T)(0),$$

see Remark 6.10. Thus, we have shown the desired result for last-point sampling, and the final steps of the proof for next-point sampling are completely analogous.

For sampling with linear interpolation, the proof is complicated by the fact that the sample path $\mathrm{SP}_{\text{linear}}(I_{a,b})$ of an indicator function is not an indicator function but rather a triangular function. However, a given step function can be approximated arbitrarily close by a sequence of trapezoid functions (which are elements of $\mathrm{SP}_{\text{linear}}$), and the measure μ_T can be defined as the limiting value of h_σ when applied to this sequence of trapezoid functions. The rest of the proof then proceeds as above, but is notationally more involved. Alternatively, one can invoke a version of the Riesz representation theorem. Specifically, the sample path $\mathrm{SP}_{\text{linear}}(X)$ of each time series $X \in \mathcal{T}$ is constant outside an interval of finite length and the space $\mathrm{SP}_{\text{linear}}$ can therefore be embedded in the space of continuous real functions that vanish at infinity. A version of the Riesz representation theorem shows that bounded, linear functionals on the latter space can be written as integrals with respect to a finite, signed measure; see Arveson (1996), Theorem 5.2. ■

Proof of Theorem 6.12. By Theorem A.2 and tick-invariance, there exists a finite signed

measure μ_T on $(\mathbb{R}, \mathcal{B})$ such that

$$\begin{aligned}\mathrm{O}(X)_t &= \mathrm{SP}_\sigma(\mathrm{O}(X))(t) \\ &= \int_{-\infty}^{+\infty} \mathrm{SP}_\sigma(X)(t-s)\,\mathrm{d}\mu_T(s) \\ &= \int_{-\infty}^{+\infty} X[t-s]_\sigma\,\mathrm{d}\mu_T(s),\end{aligned}$$

for all $t \in T(X) = T(\mathrm{O}(X))$. Because O is causal,

$$\mathrm{O}(X)_t = \mathrm{O}(X+Y)_t$$

for all time series $Y \in \mathcal{T}$ with $\mathrm{SP}_\sigma(Y)(s) = 0$ for $s \leq t$, which implies

$$\begin{aligned}0 &= \int_{-\infty}^{+\infty} Y[t-s]_\sigma\,\mathrm{d}\mu_T(s) \\ &= 0 + \int_{-\infty}^{0} Y[t-s]_\sigma\,\mathrm{d}\mu_T(s)\end{aligned}$$

for all such time series. Hence, μ_T is identical to the zero measure on $(-\infty, 0)$, and μ_T can be restricted to $(\mathbb{R}_+, \mathcal{B}_+)$, which makes O a convolution operator. ■

B Frequently Used Notation

$\mathbf{0}_n$	the null vector of length n
$*_\sigma^\mu$	the convolution operator associated with signed measure μ and sampling scheme σ
B	the backshift operator, see Definition 2.7
$\mathcal{B}$	the Borel σ-algebra on $\mathbb{R}$
$\mathcal{B}_+$	the Borel σ-algebra on $\mathbb{R}_+$
L	the lag operator, see Definition 2.7
D	the delay operator, see Definition 2.7
$N(X)$	the number of observations of time series X, see Definition 2.2
$\mathbb{T}_n$	the space of strictly increasing time sequences of length n, see Definition 2.1
$\mathbb{T}$	the space of strictly increasing time sequences, see Definition 2.1
$\mathbb{R}^n$	n-dimensional Euclidean space
$\mathbb{R}_+$	the interval $[0, \infty)$
σ	one of three sampling schemes, see Definition 2.4 and 2.6
$\mathrm{SP}_\sigma(X)$	the sample path of a time series X with sampling scheme σ, see Definition 6.1
SP_σ	the space of time series sample paths for sampling scheme σ, see Definition 6.1
$\mathcal{T}_n$	the space of time series of length n, see Definition 2.1
$\mathcal{T}$	the space of unevenly spaced time series, see Definition 2.1
$\mathcal{T}^K$	the space of K-dimension time series, see Definition 7.1
$T(X)$	the vector of observation times of a time series X, see Definition 2.2
$V(X)$	the vector of observation values of a time series X, see Definition 2.2
$X[t]_\sigma$	the sampled value of time series X at time t with sampling scheme σ, see Def. 2.4
$X^c[T_X]$	the observation time series of a continuous-time stochastic process X^c, see Def. 2.6

References

Arveson, W. (1996). Notes on measure and integration in locally compact spaces. Lecture Notes, Department of Mathematics, University of California Berkeley.

Belcher, J., J. S. Hampton, and G. Tunnicliffe-Wilson (1994). Parameterization of continuous time autoregressive models for irregularly sampled time series data. *Journal of the Royal Statistical Society. Series B 56*, 141–155.

Billingsley, P. (1995). *Probability and Measure* (Third ed.). New York: John Wiley & Sons.

Bos, R., S. de Waele, and P. M. T. Broersen (2002). Autoregressive spectral estimation by application of the burg algorithm to irregularly sampled data. *IEEE Transactions on Instrumentation and Measurement 51*(6), 1289–1294.

Brockwell, P. J. (2008). Lévy-driven continuous-time arma processes. Working Paper, Department of Statistics, Colorado State University.

Brockwell, P. J. and R. A. Davis (1991). *Time Series: Theory and Methods* (Second ed.). Springer.

Broersen, P. M. T. (2008). Time series models for spectral analysis of irregular data far beyond the mean data rate. *Measurement Science and Technology 19*, 1–14.

Dacorogna, M. M., R. Gençay, U. A. Müller, R. B. Olsen, and O. V. Pictet (2001). *An Introduction to High-Frequency Finance*. Academic Press.

Durrett, R. (2005). *Probability: Theory and Examples* (Third ed.). Duxbury Advanced Series.

Eckner, A. (2017). Algorithms for unevenly-spaced time series: Moving averages and other rolling operators. Working Paper.

Gilles Zumbach, U. A. M. (2001). Operators on inhomogeneous time series. *International Journal of Theoretical and Applied Finance 4*, 147–178.

Jones, R. H. (1981). Fitting a continuous time autoregression to discrete data. In D. F. Findley (Ed.), *Applied Time Series Analysis II*, pp. 651–682. Academic Press, New York.

Jones, R. H. (1985). Time series analysis with unequally spaced data. In E. J. Hannan, P. R. Krishnaiah, and M. M. Rao (Eds.), *Time Series in the Time Domain, Handbook of Statistics*, Volume 5, pp. 157–178. North Holland, Amsterdam.

Kolmogorov, A. N. and S. V. Fomin (1999). *Elements of the Theory of Functions and Functional Analysis*. Dover Publications.

Kreyszig, E. (1989). *Introductory Functional Analysis with Applications*. Wiley.

Lomb, N. R. (1975). Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science 39*, 447–462.

Müller, U. A. (1991). Specially weighted moving averages with repeated application of the ema operator. Working Paper, Olsen and Associates, Zurich, Switzerland.

Protter, P. (2005). *Stochastic Integration and Differential Equations* (Second ed.). New York: Springer-Verlag.

Sato, K. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press.

Scargle, J. D. (1982). Studies in astronomical time series analysis. ii. statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal 263*, 835–853.

Thiebaut, C. and S. Roques (2005). Time-scale and time-frequency analyses of irregularly sampled astronomical time series. *EURASIP Journal on Applied Signal Processing 15*, 2486–2499.

Wand, M. and M. Jones (1994). *Kernel Smoothing*. Chapman and Hall/CRC.