DEPARTMENT OF ECONOMICS
AND BUSINESS ECONOMICS
AARHUS UNIVERSITY

CREATES
Center for Research in Econometric Analysis of Time Series

Temperature Anomalies, Long Memory, and Aggregation

J. Eduardo Vera-Valdés

CREATES Research Paper 2020-16

Department of Economics and Business Economics
Aarhus University
Fuglesangs Allé 4
DK-8210 Aarhus V
Denmark

Email: oekonomi@au.dk
Tel: +45 9352 1936

TEMPERATURE ANOMALIES, LONG MEMORY, AND AGGREGATION

J. Eduardo Vera-Valdés
Department of Mathematical Sciences
Aalborg University and CREATES
eduardo@math.aau.dk

November 1, 2020

ABSTRACT

Econometric studies for global heating have typically used regional or global temperature averages to show that they exhibit long memory properties. One typical explanation behind the long memory properties of temperature averages is cross-sectional aggregation. Nonetheless, the formal analysis regarding the effect that aggregation has on the long memory dynamics of temperature data has been missing. Thus, this paper studies the long memory properties of individual grid temperatures and compares them against the long memory dynamics of global and regional averages. Our results show that the long memory parameters in individual grid observations are smaller than the ones from regional averages. Global and regional long memory estimates are found to be greatly affected by temperature measurements at the Tropics, where the data is less reliable. Thus, this paper supports the notion that aggregation may be exacerbating the long memory estimated in regional and global temperature data. The results are robust to the bandwidth parameter, limit for station radius of influence, and sampling frequency.

***K*eywords** Global Heating · Regional Temperature · Climate Econometrics · Long Memory · Aggregation

JEL Classification Q54; C22; C43; C14.

1 Introduction

From its inception, long memory models have been associated with the analysis of climate data. One of the first works on long memory is due to Hurst (1956). The author studied the long-term capacity of reservoirs for the Nile and recommended to increase the height of a dam to be built given his observations on cycles of highs at the river. His analysis showed that a dam built based on a short memory model would be more prone to overflow, hence increasing the risk of a catastrophic event, that one built based on a long memory model. That is, incorporating long memory properties leads to a more accurate characterisation of climate data. Moreover, Hurst's observations served as inspiration for Mandelbrot's work on fractional Brownian motion; see Mandelbrot (1967); Mandelbrot and Van Ness (1968). The authors labelled this phenomenon the Joseph effect in reference to the account of "seven years of great abundance" followed by "seven years of famine".

In temperature data, Bloomfield (1992), and Bloomfield and Nychka (1992) are among the first to use long memory models in the analysis. They noted that long memory should be incorporated in trend estimations for temperature series. More recently, Baillie and Chung (2002), and Mills (2007) use fractionally differenced models to estimate the long memory parameter, while Gil-Alana (2005) and Mangat and Reschenhofer (2020) use semiparametric estimators in the frequency domain. They all conclude that temperature data possess long memory. Nonetheless, they differ as to the degree of memory. The distinction is relevant given that a large degree of memory implies a nonstationary process or even a process that does not revert to the mean.

Most articles have focused on analysing regional or global average temperature anomalies data. In this regard, several authors have argued that aggregation may be the reason behind the presence of long memory in temperature

data; see Baillie and Chung (2002); Mills (2007); Gil-Alana (2005). Nonetheless, to the best of our knowledge, there is no formal analysis on whether aggregation may explain the long memory in the temperature data. Thus, this paper looks to determine the role that aggregation has on the presence of long memory in temperature data. The analysis relies on estimating the long memory parameter in the individual grid temperature series and compare them against the estimates in regional and global averages.

We use data from the NASA Goddard Institute for Space Studies, the GISS Surface Temperature Analysis (GISTEMP, hereinafter). The data is updated monthly and combines data from land and ocean surface temperatures to estimate global temperature change; see Lenssen et al. (2019); GISTEMP Team (2020). A total of 16200 grid observations are reported in the data, along with several regional temperature averages. The data is publicly available at data.giss.nasa.gov/gistemp/

The long memory parameters are obtained using semiparametric estimators in the frequency domain given the results by Haldrup and Vera-Valdés (2017) on long memory by cross-sectional aggregation. The authors show that cross-sectionally aggregated data does not imply the fractional difference operator. Thus, parametric estimators based on the fractional difference operator are misspecified for aggregated data.

Our results show that the long memory parameters in individual grid observations are lower than for the regional averages. Moreover, the long memory dynamics for the aggregated data seems to be greatly affected by the measurements around the Tropics, where data is less reliable. Thus, our results support the notion that aggregation may be exacerbating the long memory estimated in aggregated temperature data.

The results are robust to the bandwidth parameter, limit for station radius of influence, and sampling frequency.

This paper proceeds as follows: Section 2 presents the data used in the analysis. Section 3 discusses long memory modelling, aggregation as a theoretical motivation for the long memory in the data, and semiparametric estimators for long memory. Our main results are presented in Section 4, while Section 5 discusses limitations and further work.

2 Temperature Anomalies

Figure 1 shows the temperature anomalies against the 1951-1980 baseline for August 2020.

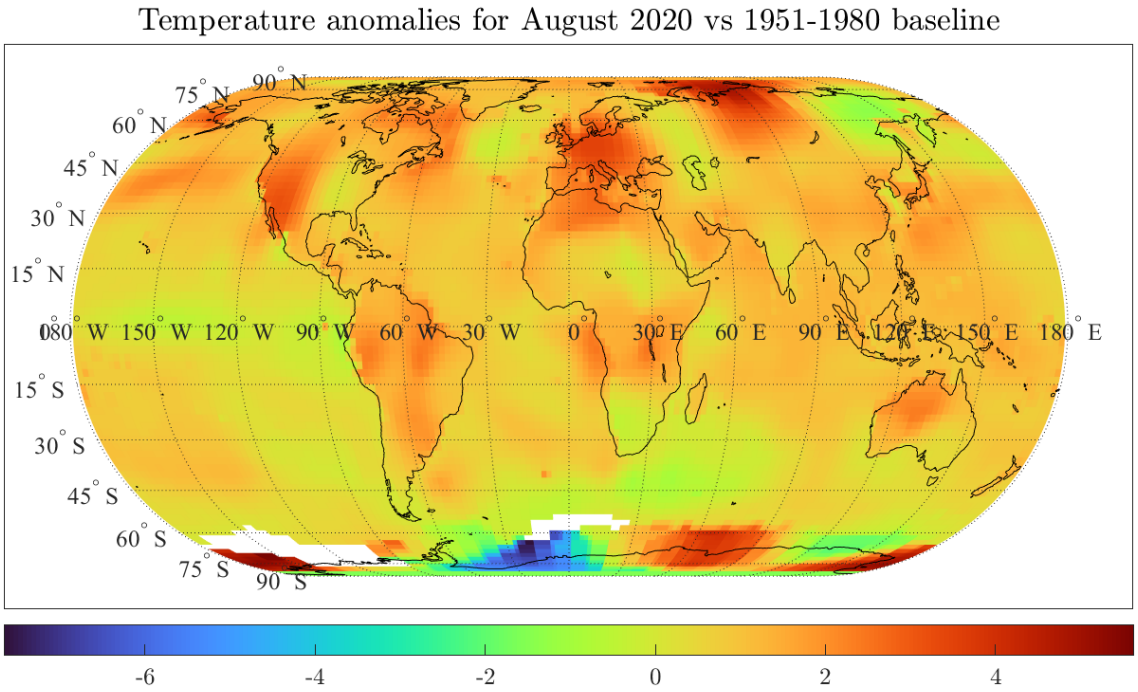

Figure 1: Author's own plot with data from GISTEMP. Smoothing radius of 1200km.

The figure shows the temperature anomalies at individual grids across the globe considering a 1200km smoothing radius; that is, each grid considers stations located within 1200km from the specified grid point. A similar map

can be plotted for each month since 1850. Furthermore, GISTEMP also presents data using a 250 smoothing radius, which we will use in our analysis as part of our robustness exercises.

Given the vast amount of information contained in GISTEMP, regional and global temperature averages are typically considered in climate econometrics research; see Section 1. Figure 2 shows the Global, Northern Hemisphere, and Southern Hemisphere temperature anomalies data reported by GISTEMP. Moreover, the figure shows the temperature anomaly series for the grid containing London, $(51°,-1°)$. Also shown are their autocorrelation functions. This paper estimates the long memory parameter for the temperature data at each grid shown in Figure 1 and compares them against the long memory estimates in the regional temperature data shown in Figure 2.

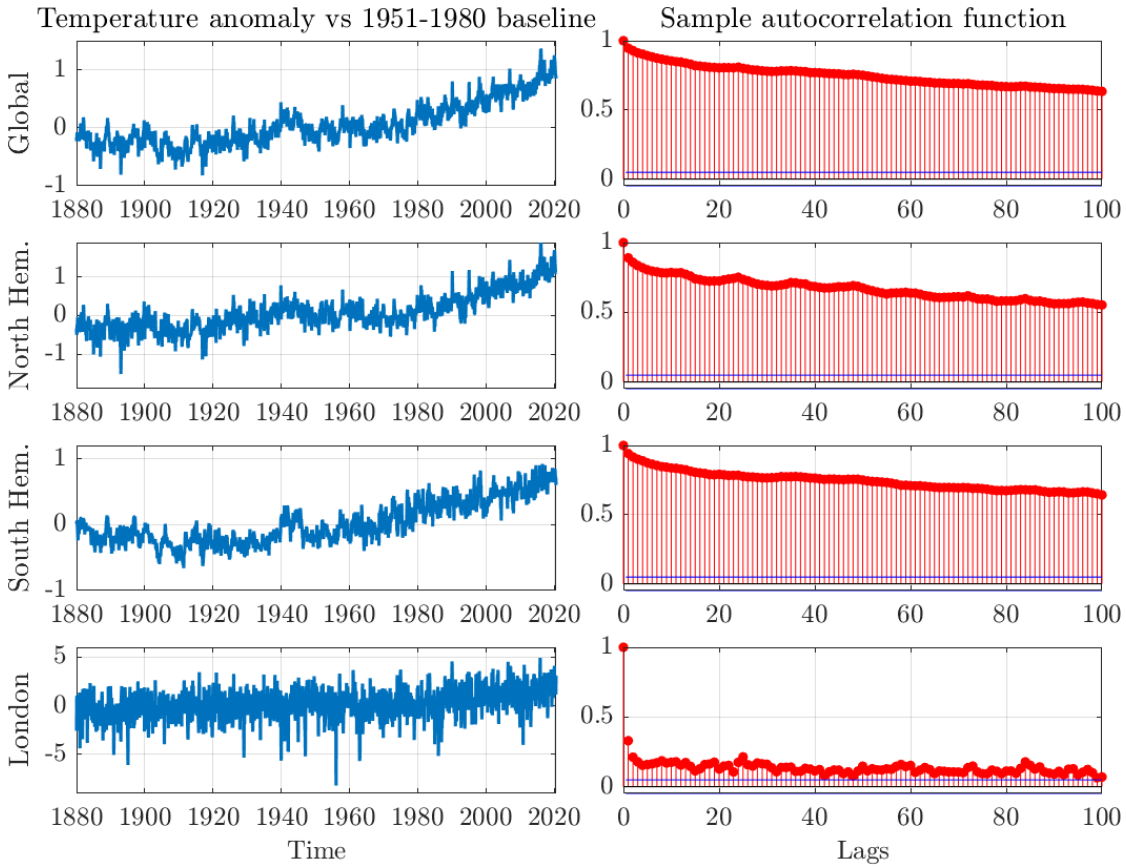

Figure 2: Temperature anomalies and sample autocorrelation functions for monthly regional temperatures presented by GISTEMP. The temperature reported near London $(51°,-1°)$ and its autocorrelation function is also shown.

One obstacle of global temperature datasets is the number of missing observations. Almost all meteorological stations undergo relocation or changes in instruments that result in missing observations; see Xu et al. (2013); Yan et al. (2010). One way to deal with missing data is to impute them using previous observations and measurements at nearby stations. Nonetheless, imputation intrinsically introduces further uncertainty into the analysis. Thus, our analysis focuses on the most recent subsample without missing observations for each grid. We discard grids with less than 50 observations in the last subsample without missing observations. A total of 219 grids out of 16200 are removed from the analysis.

3 Methods

In this section, we detail the long memory property and its different definitions in the literature. Moreover, we explain the aggregation argument behind the presence of long memory in real data and several semiparametric estimators.

3.1 Long Memory

The study of long memory in econometrics goes back to Granger's (1966) research on the shape of the spectrum of economic series near the origin. Granger found that the spectrum diverges to infinity as the frequency goes to zero for many financial and economic series, what the author called "the typical shape". This kind of behaviour has led to several definitions of long memory. In this paper, we consider three of the most common definitions of long memory; they are presented in the following definition.

Definition 1 *Let x_t be a time series with autocovariance function $\gamma_x(k)$, and spectral density function $f_x(\lambda)$, and let $d \in \mathbb{R}$, then x_t has long memory:*

1. *in the **spectral sense** if $f_x(\lambda) \sim C_f \lambda^{-2d}$ as $\lambda \to 0$ with C_f a constant;*
2. *in the **self-similar sense** if $m^{1-2d}\mathrm{Cov}\left(x_t^{(m)}, x_{t+k}^{(m)}\right) \sim C_m k^{2d-1}$ as $k, m \to \infty$ where $x_t^{(m)} = (x_{tm-m+1} + \cdots + x_{tm})/m$, with $m \in \mathbb{N}$, $m/k \to 0$, and C_m is a constant;*
3. *in the **covariance sense** if $\gamma_x(k) \sim C_x k^{2d-1}$ as $k \to \infty$ with C_x a constant.*

Above, $\mathrm{Cov}(x, y)$ is the covariance function between x and y, $g(x) \sim h(x)$ as $x \to x_0$ means that $g(x)/h(x)$ converges to 1 as x tends to x_0, and d is called the long memory parameter. Furthermore, the process is shown to be stationary if $d < 1/2$, and it reverts to its mean if $d < 1$.

Definition 1 is the feature discussed by Granger (1966) in his study of the typical spectral shape for economic variables. The definition is based on the property that the spectral density for a long memory process has a pole at the origin. The behaviour of the spectrum near the origin is also used in the construction of semi-parametric estimators in the frequency domain, see Section 3.3.

Definition 2 is based on the work on fractals and self-similarity by Mandelbrot and Van Ness (1968). Self-similarity implies that the degree of memory is asymptotically equivalent for different levels of temporal aggregation. Thus, asymptotically, the long memory parameter is statistically the same whether we estimate it at different sampling frequencies. This property is relevant for the study of temperature data given that the literature has used both monthly and yearly data to test for long memory properties. Under self-similarity, the long memory parameter estimates are statistically equivalent in both sampling frequencies. Nonetheless, the shortened sample size must be taken into consideration when estimating the long memory parameter on yearly data.

Definition 3 is concerned with the behaviour of the autocorrelation function for large lags. It was one of the motivations behind the ARFIMA class of models due to Granger and Joyeux (1980), and Hosking (1981). They proposed to use the fractional difference operator to induce long memory properties. Fractionally differenced processes are given by:

$$(1 - L)^d x_t = \epsilon_t, \tag{1}$$

where ϵ_t is a white noise process, and $d \in \mathbb{R}$. Using the standard binomial expansion, we can decompose the fractional difference operator, $(1 - L)^d$, in a series with coefficients $\phi_j = \Gamma(j + d)/(\Gamma(d)\Gamma(j + 1))$ for $j \in \mathbb{N}$, and where $\Gamma(j)$ is the gamma function. Using Stirling's approximation, it can be shown that these coefficients decay at a hyperbolic rate ($\phi_j \sim j^{d-1}$ as $j \to \infty$), which in turn translates to slowly decaying autocorrelations. Figure 2 shows the sample autocorrelation functions for global and regional temperature anomalies, as well as for the temperature anomaly at the London grid ($51°$,$-1°$). Note that all autocorrelation functions show hyperbolic decay. In particular, the autocorrelations are statistically significant even after 100 lags. That is, temperature data seems to exhibit long memory in the covariance sense. Furthermore, the degree of memory seems to be larger for the regional and global averages than for the individual temperature grid. This analysis relies on properly estimating these degrees of memory.

The fractional difference operator remains to be the most popular mechanism to model and generate long memory even though, to the best of our knowledge, there are no theoretical arguments for its occurrence in real data. In this regard, the next section presents cross-sectional aggregation, one of the most predominant theoretical explanations for the presence of long memory in real data.

3.2 Cross-sectional aggregation

Granger (1980) showed that long memory can result from the cross-sectional aggregation of short memory processes. The author considered the process defined by:

$$x_t = \frac{1}{\sqrt{N}} \sum_{j=1}^{N} x_{j,t}, \tag{2}$$

where $j = 1, \cdots, N \in \mathbb{N}$, and each $x_{j,t}$ is an independent autoregressive process with random coefficient given by:

$$x_{j,t} = \alpha_j x_{j,t-1} + \epsilon_{j,t} \tag{3}$$

where $\varepsilon_{j,t}$ is an independent identically distributed process with $E[\epsilon_{j,t}] = 0$, and $E[\epsilon_{j,t}^2] = \sigma^2$, $\forall t \in \mathbb{Z}$. Furthermore, α_j^2 is sampled from the Beta distribution, independent from $\varepsilon_{j,t}$, with the following density:

$$\mathcal{B}(\alpha; a, b) = \frac{1}{B(a,b)} \alpha^{a-1}(1-\alpha)^{b-1} \quad \text{for} \quad \alpha \in (0,1),$$

with $a, b > 0$, and where $B(a,b)$ is the Beta function.

Granger (1980) showed that taking a large number in the cross-sectional dimension, the resulting process will exhibit hyperbolic decay instead of the standard geometric one. Moreover, Haldrup and Vera-Valdés (2017) showed that cross-sectionally aggregated processes are long memory according to all definitions considered in this analysis.

The aggregation argument is illustrated in Figure 3. The figure shows the autocorrelation functions for four autoregressive processes with different autoregressive parameter, (3). Moreover, the figure shows the autocorrelation function for the aggregated process, (2).

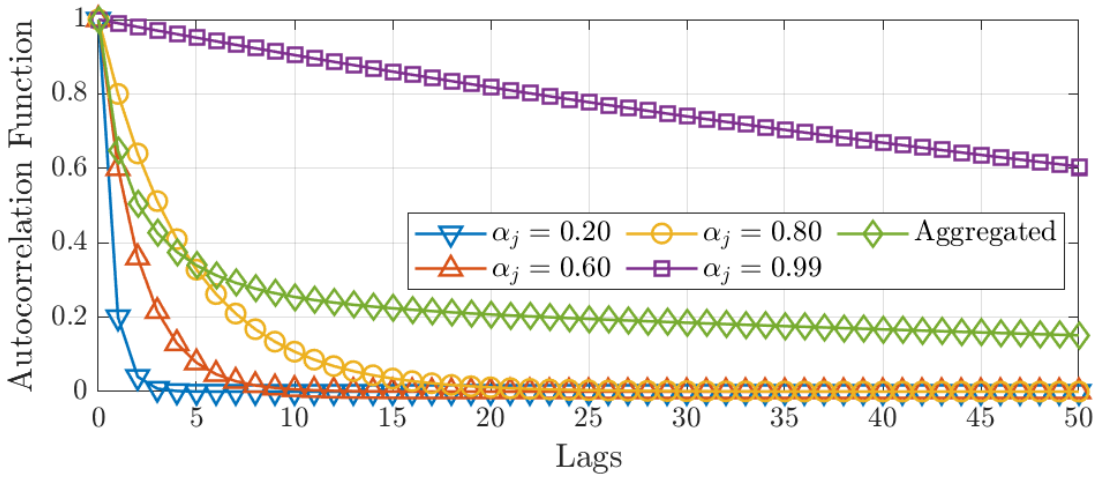

Figure 3: Autocorrelation functions for individual AR(1) processes with different autoregressive coefficient. Also shown, the autocorrelation function of the aggregated process.

As argued by Granger (1980), and later shown by Zaffaroni (2004), the degree of memory of the aggregated process is determined by the behaviour of the Beta density function near one. That is, the more probability mass around one, the greater the degree of memory. An example of this argument can be seen in Figure 3. Note that the long-term behaviour of the aggregated process is influenced by combining three autoregressive processes with coefficient away from one, with one autoregressive process with a near-unity coefficient. In this regard, the more autoregressive processes with coefficients near unity, the greater the degree of memory in the aggregate.

The cross-sectional aggregation result has been extended in several directions, including to allow for general $ARMA$ processes, as well as to other distributions; see Linden (1999), Oppenheim and Viano (2004), and Zaffaroni (2004), to name a few. Furthermore, Vera-Valdés (2020) has shown via simulations that the fractional difference operator produces well performing forecasts for aggregated processes.

In economic data, cross-sectional aggregation plays a significant role in the generation of long memory. For example, cross-sectional aggregation has been cited as the source of long-range dependence for inflation, output,

and volatility; see Balcilar (2004), Diebold and Rudebusch (1989), Altissimo et al. (2009), and Osterrieder et al. (2019). As pointed in Section 1, aggregation has been cited as the explanation behind the presence of long memory in temperature data given its reliance on temperature aggregates.

One important distinction regarding long memory generated by aggregation is that it does not belong to the class of processes generated using the fractional difference operator; see Haldrup and Vera-Valdés (2017). Thus, parametric estimators based on the fractional difference operator are misspecified for long memory processes by cross-sectional aggregation. Nonetheless, given that aggregated processes are long memory in the spectral sense, we can consistently estimate the long memory parameter in the frequency domain. The next section presents the semiparametric estimators of long memory in the frequency domain.

3.3 Semiparametric estimators of long memory

The semiparametric estimators in the frequency domain are based on the long memory definition in the spectral sense, see 1. The idea is to evaluate the periodogram of the time series, an estimator of the spectral density, in a vicinity of the origin, where the spectral density, $f_X(\lambda)$, is driven by the long memory parameter.

The semiparametric estimators are typically divided between the log-periodogram regression method suggested by Geweke and Porter-Hudak (1983), and the local Whittle approach developed by Künsch (1987).

On the one hand, the log-periodogram regression [GPH, henceforth] is given by:

$$\log(I(\lambda_k)) = c - 2d\log(\lambda_k) + u_k, \quad k = 1, \cdots, m, \tag{4}$$

where $I(\lambda_k)$ is the periodogram of x_t, $\lambda_k = e^{i2\pi k/T}$ are the Fourier frequencies, c is a constant, u_k is the error term, and m is a bandwidth parameter that grows with the sample size. Note that the zero frequency is not included in (4), making the estimator robust to the specification of the mean.

The consistency and asymptotic normality of the log-periodogram regression have been proved by Robinson (1995b); Velasco (1999b, 2000). Denote $\hat{d}_{GPH}$ to the estimate of the long memory parameter via the log-periodogram regression, the authors show that:

$$\sqrt{m}(\hat{d}_{GPH} - d) \underset{d}{\to} N\left(0, \frac{\pi^2}{24}\right), \tag{5}$$

where m is the bandwidth as before, and $\underset{d}{\to}$ denotes convergence in distribution.

Andrews and Guggenberger (2003) [AG, hereinafter] proposed to replace the constant in (4) with a polynomial in λ_k^2 to reduce the bias. Note that the polynomial only considers even degrees given that odd degrees do not help in reducing the bias. In classical bias-variance trade-off, the reduction in bias comes at the cost of an increase in the variance, which depends on the degree of the polynomial used for estimation. For our analysis, we add one polynomial term in (4), which results in the variance of the estimate of 2.25 times the one shown in (5).

On the other hand, Künsch (1987) used a likelihood approach in the so-called local Whittle estimator of the long memory parameter. The author proposed to estimate the parameter as the minimiser of the local Whittle likelihood function given by:

$$R(d) = \log\left(\frac{1}{m}\sum_{k=1}^{m}\lambda_k^{2d} I(\lambda_k)\right) - \frac{2d}{m}\sum_{k=1}^{m}\log(\lambda_k), \tag{6}$$

where $I(\lambda_k)$ is the periodogram of x_t, $\lambda_k = e^{i2\pi k/T}$ are the Fourier frequencies, and m is the bandwidth parameter.

The consistency and asymptotic normality of the local Whittle estimator have been proved for regions of d that are empirically relevant in most application by Robinson (1995a); Velasco (1999a); Phillips and Shimotsu (2004). Denote $\hat{d}_{LW}$ to the estimate of the long memory parameter via the local Whittle estimator, the authors show that:

$$\sqrt{m}(\hat{d}_{LW} - d) \underset{d}{\to} N\left(0, \frac{1}{4}\right), \tag{7}$$

where m is the bandwidth as before. In particular, note that the local Whittle estimator has less variance than both estimators by log-periodogram regression above.

A further refinement to the local Whittle approach was suggested by Shimotsu and Phillips (2005). The authors proposed the exact local Whittle estimator as the minimiser of the function given by:

$$R(d) = \log\left(\frac{1}{m}\sum_{k=1}^{m} I_{\Delta^d}(\lambda_k)\right) - \frac{2d}{m}\sum_{k=1}^{m}\log(\lambda_k),$$

where $I_{\Delta^d}(\lambda_k)$ is the periodogram of $(1-L)^d x_t$, $\lambda_k = e^{i2\pi k/T}$ are the Fourier frequencies, and m is the bandwidth parameter. The estimator was later extended to allow for nonzero means and time trends in the so-called feasible exact local Whittle (FELW, hereinafter); see Shimotsu (2010). Furthermore, the author shows that the consistency and asymptotic normality of the FELW are the same as for the local Whittle estimator.

Our preferred method uses the mean-squared error optimal bandwidth of $m = T^{4/5}$, where T is the sample size, for all the estimators. The optimal bandwidth was obtained by Hurvich et al. (1998). Nonetheless, we consider the commonly used bandwidth of $m = T^{1/2}$ to allow for easy comparisons to other research, and as a robustness exercise.

Furthermore, note that the variances of the semiparametric estimators of long memory only depend on the bandwidth, which directly depends on the sample size; see (5) and (7). Thus, estimates using the same bandwidth and approximately the same sample size will exhibit similar uncertainty. Moreover, analysis using yearly observations may suffer from increased uncertainty in the estimates, given the reduced sample size. However, following the self-similar property of long memory processes, they are asymptotically equivalent, see Definition 2.

4 Results

In this section, we show the results from estimating the long memory parameter using semiparametric estimators in temperature anomalies data. We consider temperature anomalies in individual grids and compare them against regional and global averages. Our preferred method uses the optimal bandwidth and the larger monthly datasets. As robustness exercises, we consider yearly data, different smoothing radius, and values for the bandwidth parameter.

4.1 Monthly data with 1200km smoothing radius and optimal bandwidth

Table 1 presents the long memory estimates from the monthly regional and global temperature averages computed by GISS, as well as for temperature at the London grid for the 1200km smoothing radius dataset; that is, the data presented in Figure 3. We use the optimal bandwidth of $T^{4/5}$, with T the sample size. The table shows that the long memory parameter is statistically larger for the regional and global averages than in the London grid. In particular, all estimates from the regional and global averages are more than four standard deviations apart from the estimates at the London grid. That is, the confidence intervals do not intersect, and are thus statistically different.

Moreover, long memory estimates for regional and global averages imply that they follow a nonstationary process; that is, $d > 0.5$. Nonetheless, all long memory estimates imply processes that revert to the mean, $d < 1$.

Series	GPH	AG	LW	FELW	Sample Size
Global	0.6332	0.7475	0.6682	0.6304	1688
North Hem.	0.5429	0.6247	0.5765	0.5584	1688
South Hem.	0.6407	0.6904	0.6430	0.6159	1688
London ($51°$,$-1°$)	0.2346	0.2388	0.2288	0.2354	1688
Std. Dev.	0.0328	0.0492	0.0256	0.0256	

Table 1: Long memory estimates for monthly regional temperature anomalies using the optimal bandwidth. For the London grid, we use the 1200km smoothing radius dataset.

As previously mentioned, GISTEMP contains individual temperature data for 16200 grids. It will be unwieldy to present a table with long memory estimates for the temperature at all grids. Thus, Figure 4 presents the results for all individual grid points in a global map for each estimator considered. The figure presents some interesting findings.

First, note that the degree of memory seems to be smaller for land than for ocean surface temperature. Moreover, the degree of memory decreases as we move away from the Equator. There are two possible explanations for this result. The small number of long stations records in Africa and South America could make the results less reliable there; see (Hansen et al., 2010). Moreover, temperature around the Tropics is influenced by El Niño cycle of ocean temperature.

Second, long memory estimates at individual land grids seem to be in the stationary range of $d < 0.5$, particularly for grids away from the Tropics. Thus, long memory estimates at near all individual land grids seem to be significantly smaller than the global and regional averages.

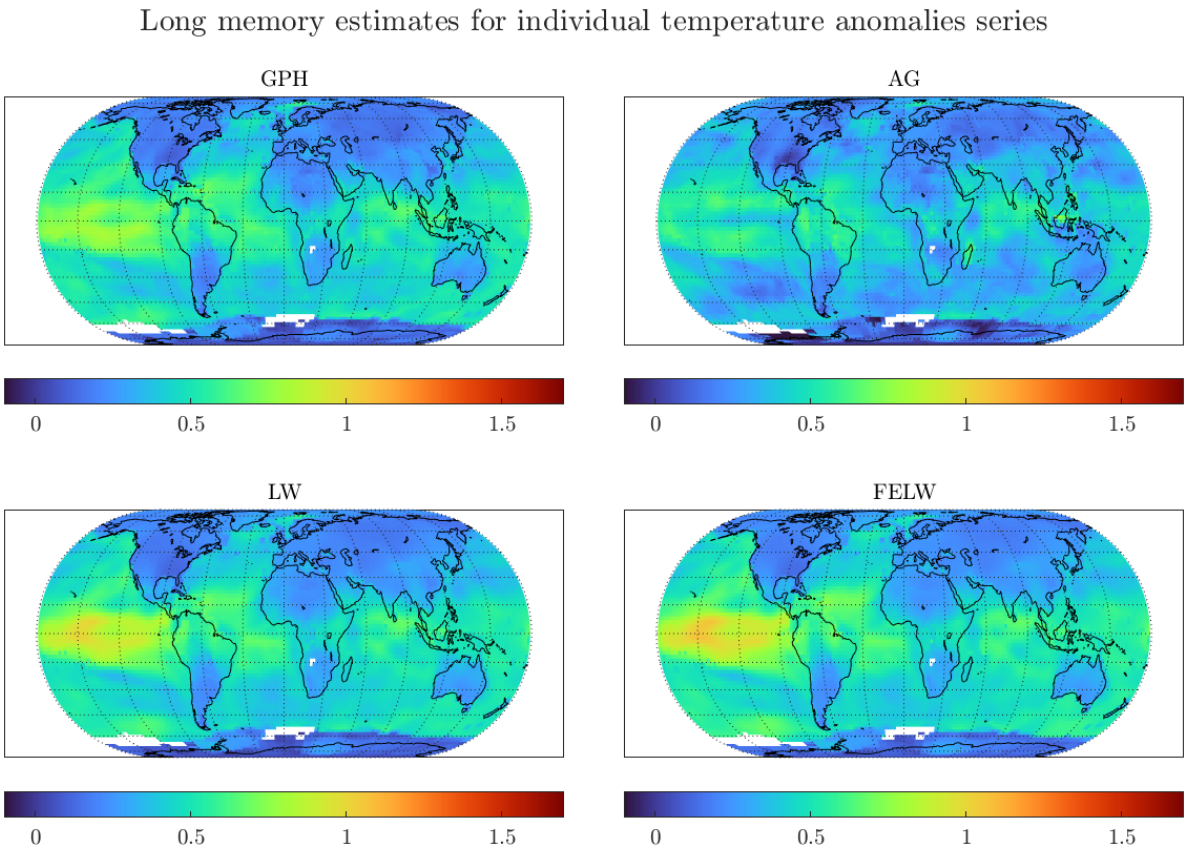

Figure 4: Long memory estimates for monthly individual grid observations of temperature anomalies. We use the optimal bandwidth and the 1200km smoothing radius dataset.

Third, besides a few ocean grid temperatures using the LW and FELW estimators, all long memory estimates are in the mean-reverting range, $d < 1$.

Overall, the results from Figure 4 suggest that the long memory dynamics for the global and regional averages are greatly affected by aggregation, particularly due to the influence of ocean temperature at the Tropics. The estimates from the less reliable temperature measurements at the Tropics, and particularly at the Oceans, propagate due to aggregation to the global and regional long memory estimates.

To shed light on the influence that the temperature data at the Tropics has on the long memory estimates, Table 2 presents averages from the long memory estimates in individual grid stations for selected regional subsamples. The table shows mean values for the long memory estimates for all grids, the Global estimate; and for grids in the Northern and Southern Hemispheres. Furthermore, the table shows long memory estimate averages for grids located in the Tropics, the Arctic, and the Antarctic; that is, those between latitudes $-30°$ and $30°$, $67°$ and $90°$, and $-67°$ and $-90°$, respectively.

Series	GPH	AG	LW	FELW	Sample Size
Global	0.3427	0.2951	0.3540	0.3738	1463
North Hem.	0.3236	0.3082	0.3389	0.3539	1560
South Hem.	0.3624	0.2817	0.3695	0.3942	1364
Tropics	0.4855	0.4103	0.5019	0.5216	1642
Arctic	0.2164	0.2588	0.2404	0.2525	1286
Antarctic	0.1174	0.0874	0.1084	0.1401	681

Table 2: Regional long memory averages for monthly temperature anomalies. We use the optimal bandwidth and the 1200km smoothing radius dataset. The bar on top denotes averages.

The table broadly confirms the results from the individual grid temperatures. That is, the degree of memory seems to be smaller as we move away from the Tropics. Moreover, only the average degree of memory at the Tropics seems to fall into the nonstationary range, and only for the LW and FELW estimators. This suggests that it is the temperature at a few grids around the Tropics that seems to be increasing the degree of memory at the global and regional temperature averages.

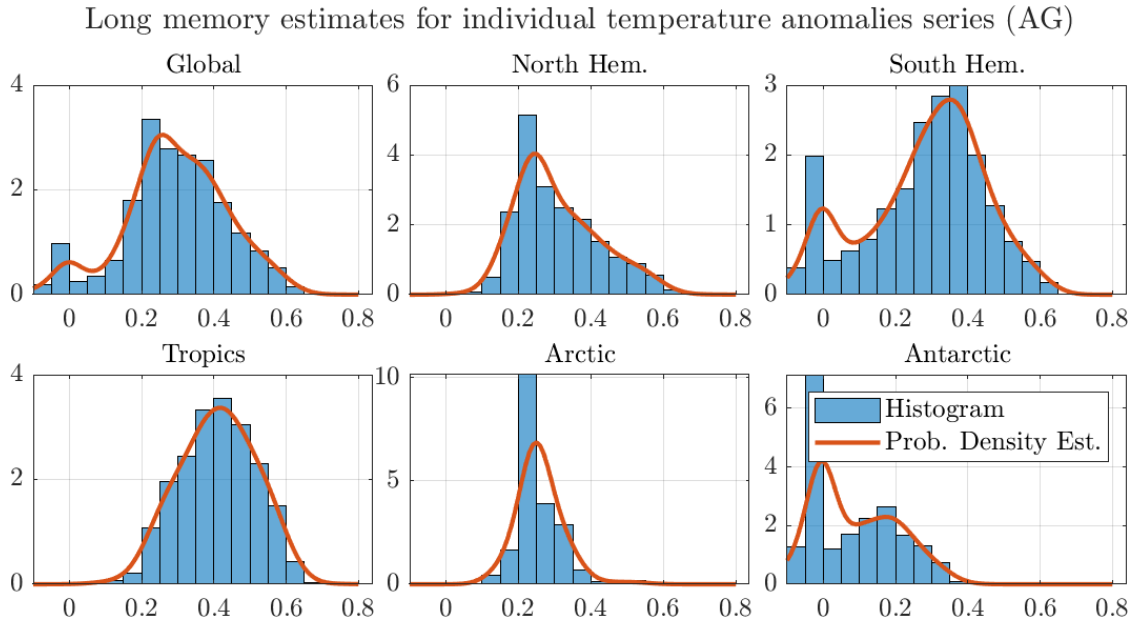

Figure 5: Histograms and probability density estimates of long memory estimates for monthly individual grid observations of temperature anomalies. We use the optimal bandwidth, the 1200km smoothing radius dataset, and the AG estimator.

The effect that the memory estimates around the Tropics have on the regional aggregates is illustrated in Figure 5. The figure shows histograms and estimated probability densities for the estimates of the memory in several regions. Note that the density for the estimates around the Tropics puts more weight to the nonstationary range than the densities for the other regions. Using an analogous argument to the one in Figure 3, this larger weight to nonstationary values results in a larger degree of memory for the regional average. Furthermore, the same effect can be observed in all regions. In particular, both the Arctic and Antarctic regional estimates assign almost no weight to the nonstationary region, which results in an average degree of memory well inside the stationary range.

The results above suggest that aggregation plays a role in the larger degree of long memory estimated for the global and regional temperature anomalies. They show that the region around the Tropics is associated with temperature anomalies with larger degrees of memory. These larger degrees of memory are then propagated to the global and regional averages through aggregation.

Given the increased uncertainty in ocean temperature measurements, and as a robustness exercise, next section analyses the effect that the smoothing radius has on the results. That is, next section analyses the dataset for the 250km smoothing radius.

4.2 Monthly data with 250km smoothing radius and optimal bandwidth

Figure 6 presents the results for all individual grid points for the 250km smoothing radius dataset using the optimal bandwidth, while Figure 7 shows their histograms and probability density estimates. Note that, given the smaller smoothing radius, Figure 6 shows several missing observations around the Tropics, particularly in the Oceans.

Figures 6 and 7 show that the long memory estimates are smaller for the smaller smoothing radius. In particular, almost all long memory estimates are in the stationary range. This once again points to the effect that aggregation has on the long memory estimates.

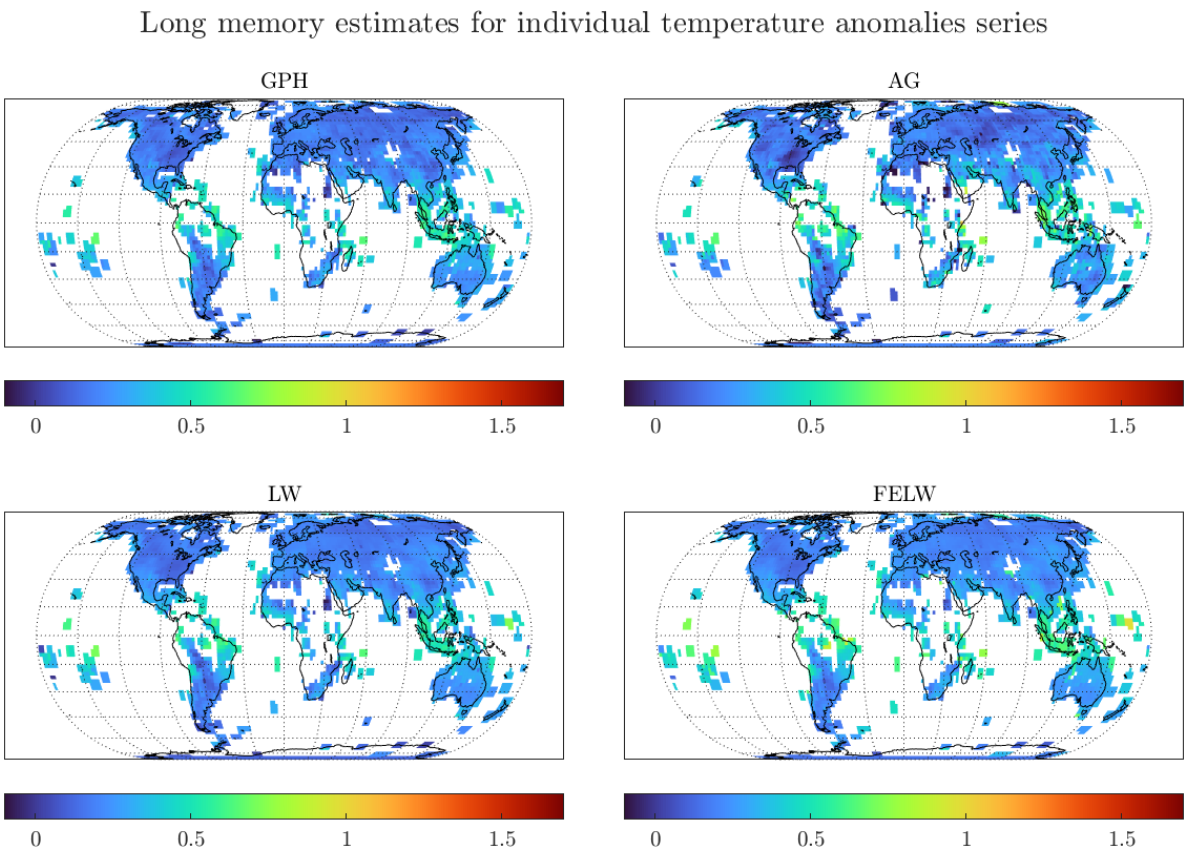

Figure 6: Long memory estimates for monthly individual grid observations of temperature anomalies. We use the optimal bandwidth and the 250km smoothing radius dataset.

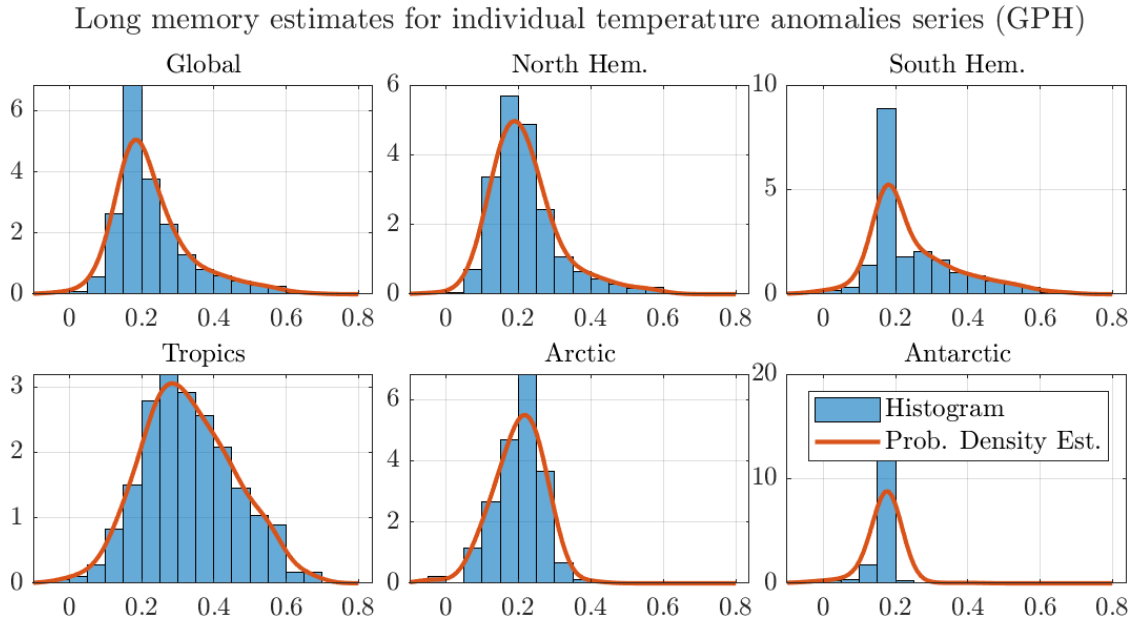

Figure 7: Histograms and probability density estimates of long memory estimates for monthly individual grid observations of temperature anomalies. We use the optimal bandwidth, the 250km smoothing radius dataset, and the GPH estimator.

Furthermore, Figure 7 shows that long memory estimate averages for regions away from the Tropics put no weight to the nonstationary range. Following the aggregation argument, this smaller weight translates into smaller degrees of memory for all regional averages, particularly for the Arctic and Antarctic, as can be seen in Table 3.

Series	GPH	AG	LW	FELW	Sample Size
Global	0.2279	0.2406	0.2277	0.2626	895
North Hem.	0.2179	0.2286	0.2315	0.2596	1106
South Hem.	0.2452	0.2615	0.2209	0.2678	528
Tropics	0.3308	0.3406	0.3315	0.3786	732
Arctic	0.2009	0.2259	0.2248	0.2678	845
Antarctic	0.1680	0.2005	0.1205	0.1633	284

Table 3: Regional long memory averages for monthly temperature anomalies. We use the optimal bandwidth and the 250km smoothing radius dataset. The bar on top denotes averages.

4.3 Robustness exercises

This section presents the results from some of the robustness exercises considered in this study. In particular, Tables 4 and 5 show results using yearly data for the 250km smoothing radius dataset. Table 4 presents the results using the optimal bandwidth, while Table 5 uses the commonly used bandwidth given by $T^{1/2}$, where T is the sample size.

Series	GPH	AG	LW	FELW	Sample Size
Global	0.4051	0.4278	0.3458	0.5573	78
North Hem.	0.3784	0.5079	0.3621	0.5218	94
South Hem.	0.4542	0.2802	0.3159	0.6225	47
Tropics	0.4621	0.5088	0.4074	0.6608	66
Arctic	0.3166	0.4880	0.3083	0.4939	77
Antarctic	0.4682	0.1177	0.2578	0.6449	24

Table 4: Regional long memory averages for yearly temperature anomalies. We use the optimal bandwidth and the 250km smoothing radius dataset. The bar on top denotes averages.

Series	GPH	AG	LW	FELW	Sample Size
Global	0.5947	0.8443	0.5343	0.6429	78
North Hem.	0.6276	1.0159	0.5650	0.6364	94
South Hem.	0.5342	0.5284	0.4777	0.6549	47
Tropics	0.6718	0.8727	0.6115	0.6725	66
Arctic	0.5528	1.2723	0.4883	0.5924	77
Antarctic	0.4133	0.4156	0.3766	0.6872	24

Table 5: Regional long memory averages for yearly temperature anomalies. The bandwidth is given by $m = T^{1/2}$, with T the sample size. We use the 250km smoothing radius dataset. The bar on top denotes averages.

The results from the robustness exercise broadly agree with the results from our preferred method. That is, the degree of memory decreases as we move away from the Tropics. Nonetheless, note that the reduced sample size increases the variance of all estimates, making them less reliable.

Moreover, note from (5) and (7) that decreasing the bandwidth also increases the variance of the estimates. In this regard, the degrees of memory estimated in Table 5 are even more uncertain. However, as argued before, they point to the same conclusion as the one from the larger dataset.

Finally, note that the results from Table 5 are in line with the ones from Mangat and Reschenhofer (2020). Using the same sampling frequency and bandwidth, the authors find that global and regional temperature series fall into the nonstationary range. Furthermore, the authors find some of the long memory estimates for regional temperature series to imply processes that do not revert to the mean; that is, $d \geq 1$. Our results suggest that these results may come from the small bandwidth and sample size.

Additional robustness exercises considering all combinations of the sampling frequency, smoothing radius, and bandwidth point broadly in the same direction, and they are available upon request.

5 Discussion

As recognised by Hurst in one of the first studies on long memory, assessing the long memory properties of natural phenomena can have major repercussions in policy design. Under the current climate emergency, assessing the long memory properties of temperature data can help to properly evaluate the pace of global heating and its likely repercussions. In a recent analysis, Calel et al. (2020) show that uncertainty plays a major role in computing the economic costs of climate change. The authors show that an analysis that does not incorporate all uncertainties can vastly underestimate the economic costs. Furthermore, they argue that long memory properties in temperature data can result in estimates of even greater economic damages; thus, they appeal for a proper characterisation of the long-term dynamics of temperature data.

In the econometric literature, most climate econometrics studies have focused on analysing regional or global average temperature data given the vast amount of information in temperature datasets. In this regard, they have shown that global and regional temperature anomalies possess long memory properties. While several authors have pointed to aggregation as a possible explanation behind the long memory in the data, to the best of our knowledge, no analysis has focused on determining if aggregation can indeed be the long memory generating mechanism for temperature data.

Thus, this paper looks to determine if aggregation is behind the presence of long memory in regional temperature data. The analysis relies on estimating the long memory parameter in the individual grid temperature series and compare them against the estimates in regional and global averages. We consider several regional averages in addition to the ones commonly used in the literature. Furthermore, the long memory parameters are obtained using semiparametric estimators in the frequency domain, given that the fractional difference operator is misspecified for aggregated processes.

Our results show that the long memory parameters at individual grid observations are lower than the regional averages. Furthermore, the degrees of memory of the aggregates seem to be greatly influenced by the long term dynamics of temperature anomalies in the Tropics, where data is less reliable. Thus, our results support the notion that aggregation may be exacerbating the long memory estimated in temperature data. Our results are robust to the frequency sampling, smoothing radius, and bandwidth parameter.

Nonetheless, our analysis relies on temperature anomalies at individual grids. Global datasets for temperature analysis already incorporate some level of aggregation at all levels. In light of the results of this paper, an additional analysis should be considered for temperature data at individual stations where the data may be obtained.

Furthermore, our analysis showed that temperature at the Tropics plays a major role in determining the long memory properties of global and regional temperature data. A robust analysis of temperature data around the Tropics that considers measurement uncertainties and El Niño effect should be considered.

Funding: This research received no external funding.

Conflicts of interest: The author declares no conflict of interest.

Data availability statement: The data that support the findings of this study are publicly available at data.giss.nasa.gov/gistemp/.

Abbreviations The following abbreviations are used in this manuscript:

GISTEMP	NASA Goddard Institute for Space Studies Surface Temperature Analysis
GPH	Geweke and Porter-Hudak log-periodogram regression
AG	Andrews and Guggenberger bias reduced log-periodogram estimator
LW	Local Whittle estimator
FELW	Feasible Exact Local Whittle estimator
Std. Dev.	Standard deviation
Prob. Density Est.	Probability density estimate

References

Altissimo, F., Mojon, B., and Zaffaroni, P. (2009). Can Aggregation Explain the Persistence of Inflation? *Journal of Monetary Economics*, 56(2):231–241.

Andrews, D. W. and Guggenberger, P. (2003). A Bias-Reduced Log-Periodogram Regression Estimator For The Long-Memory Parameter. *Econometrica*, 71(2):675–712.

Baillie, R. T. and Chung, S. K. (2002). Modeling and forecasting from trend-stationary long memory models with applications to climatology. *International Journal of Forecasting*, 18(2):215–226.

Balcilar, M. (2004). Persistence in Inflation: Does Aggregation Cause Long Memory? *Emerging Markets Finance and Trade*, 40(5):25–56.

Bloomfield, P. (1992). Trends in global temperature. *Climatic Change*, 21(1):1–16.

Bloomfield, P. and Nychka, D. (1992). Climate spectra and detecting climate change. *Climatic Change*, 21(3):275–287.

Calel, R., Chapman, S. C., Stainforth, D. A., and Watkins, N. W. (2020). Temperature variability implies greater economic damages from climate change. *Nature Communications*, 11(1):1–5.

Diebold, F. X. and Rudebusch, G. D. (1989). Long Memory and Persistence in Agregate Output. *Journal of Monetary Economics*, 24(2):189–209.

Geweke, J. and Porter-Hudak, S. (1983). The Estimation and Application of Long Memory Time Series Models. *Journal of Time Series Analysis*, 4(4):221–238.

Gil-Alana, L. A. (2005). Statistical modeling of the temperatures in the Northern Hemisphere using fractional integration techniques. *Journal of Climate*, 18(24):5357–5369.

GISTEMP Team (2020). GISS Surface Temperature Analysis (GISTEMP), version 4. NASA Goddard Institute for Space Studies. https://data.giss.nasa.gov/gistemp/. Accessed: 10/10/2020.

Granger, C. W. (1966). The Typical Spectral Shape of an Economic Variable. *Econometrica*, 34(1):150–161.

Granger, C. W. (1980). Long Memory Relationships and the Aggregation of Dynamic Models. *Journal of Econometrics*, 14(2):227–238.

Granger, C. W. and Joyeux, R. (1980). An Introduction to Long Memory Time Series Models and Fractional Differencing. *Journal of Time Series Analysis*, 1(1):15–29.

Haldrup, N. and Vera-Valdés, J. E. (2017). Long Memory, Fractional Integration, and Cross-Sectional Aggregation. *Journal of Econometrics*, 199(1):1–11.

Hansen, J., Ruedy, R., Sato, M., and Lo, K. (2010). Global surface temperature change. *Reviews of Geophysics*, 48(4):29.

Hosking, J. R. M. (1981). Fractional Differencing. *Biometrika*, 68(1):165–176.

Hurst, H. E. (1956). The Problem of Long-Term Storage in Reservoirs. *International Association of Scientific Hydrology. Bulletin*, 1(3):13–27.

Hurvich, C. M., Deo, R., and Brodsky, J. (1998). The Mean Squared Error of Geweke and Porter-Hudak's Estimator of the Memory Parameter of a Long-Memory Time Series. *Journal of Time Series Analysis*, 19(1):19–46.

Künsch, H. (1987). Statistical Aspects of Self-Similar Processes. *Bernouli*, 1:67–74.

Lenssen, N. J., Schmidt, G. A., Hansen, J. E., Menne, M. J., Persin, A., Ruedy, R., and Zyss, D. (2019). Improvements in the GISTEMP Uncertainty Model. *Journal of Geophysical Research: Atmospheres*, 124(12):6307–6326.

Linden, M. (1999). Time Series Properties of Aggregated AR(1) Processes with Uniformly Distributed Coefficients. *Economics Letters*, 64(1):31–36.

Mandelbrot, B. B. (1967). How Long Is the Coast of Britain? Statistical Self-Similarity and Fractional Dimension. *Science*, 156(3775):636–638.

Mandelbrot, B. B. and Van Ness, J. W. (1968). Fractional Brownian Motions, Fractional Noises and Applications. *SIAM Review*, 10(4):422–437.

Mangat, M. K. and Reschenhofer, E. (2020). Frequency-Domain Evidence for Climate Change. *Econometrics*, 8(3):28.

Mills, T. C. (2007). Time series modelling of two millennia of northern hemisphere temperatures: Long memory or shifting trends? *Journal of the Royal Statistical Society. Series A: Statistics in Society*, 170(1):83–94.

Oppenheim, G. and Viano, M. C. (2004). Aggregation of Random Parameters Ornstein-Uhlenbeck or AR Processes: Some Convergence Results. *Journal of Time Series Analysis*, 25(3):335–350.

Osterrieder, D., Ventosa-Santaulària, D., and Vera-Valdés, J. E. (2019). The VIX, the Variance Premium, and Expected Returns*. *Journal of Financial Econometrics*, 17(4):517–558.

Phillips, P. C. and Shimotsu, K. (2004). Local whittle estimation in nonstationary and unit root cases. *Annals of Statistics*, 32(2):656–692.

Robinson, P. M. (1995a). Gaussian Semiparametric Estimation of Long Range Dependence. *The Annals of Statistics*, 23(5):1630–1661.

Robinson, P. M. (1995b). Log-Periodogram Regression of Time Series with Long Range Dependence. *The Annals of Statistics*, 23(3):1048–1072.

Shimotsu, K. (2010). Exact local Whittle estimation of fractional integration with unknown mean and time trend. *Econometric Theory*, 26(2):501–540.

Shimotsu, K. and Phillips, P. C. (2005). Exact Local Whittle Estimation of Fractional Integration. *The Annals of Statistics*, 33(4):1890–1933.

Velasco, C. (1999a). Gaussian semiparametric estimation of non-stationary time series. *Journal of Time Series Analysis*, 20(1):87–127.

Velasco, C. (1999b). Non-stationary log-periodogram regression. *Journal of Econometrics*, 91(2):325–371.

Velasco, C. (2000). Non-Gaussian Log-Periodogram Regression. *Econometric Theory*, 16(1):44–79.

Vera-Valdés, J. E. (2020). On long memory origins and forecast horizons. *Journal of Forecasting*, 39(5):811–826.

Xu, W., Li, Q., Wang, X., Yang, S., Cao, L., and Feng, Y. (2013). Homogenization of Chinese daily surface air temperatures and analysis of trends in the extreme temperature indices. *Journal of Geophysical Research*, 118.

Yan, Z., Li, Z., Li, Q., and Jones, P. D. (2010). Effects of site change and urbanisation in the Beijing temperature series 1977-2006. *International Journal of Climatology*, 30(8):1226–1234.

Zaffaroni, P. (2004). Contemporaneous Aggregation of Linear Dynamic Models in Large Economies. *Journal of Econometrics*, 120(1):75–102.

Research Papers 2020

2020-01: Mikkel Bennedsen: Designing a sequential testing procedure for verifying global CO2 emissions

2020-02: Juan Carlos Parra-Alvarez, Hamza Polattimur and Olaf Posch: Risk Matters: Breaking Certainty Equivalence

2020-03: Daniel Borup, Bent Jesper Christensen, Nicolaj N. Mühlbach and Mikkel S. Nielsen: Targeting predictors in random forest regression

2020-04: Nicolaj N. Mühlbach: Tree-based Synthetic Control Methods: Consequences of moving the US Embassy

2020-05: Juan Carlos Parra-Alvarez, Olaf Posch and Mu-Chun Wang: Estimation of heterogeneous agent models: A likelihood approach

2020-06: James G. MacKinnon, Morten Ørregaard Nielsen and Matthew D. Webb: Wild Bootstrap and Asymptotic Inference with Multiway Clustering

2020-07: Javier Hualde and Morten Ørregaard Nielsen: Truncated sum of squares estimation of fractional time series models with deterministic trends

2020-08: Giuseppe Cavaliere, Morten Ørregaard Nielsen and Robert Taylor: Adaptive Inference in Heteroskedastic Fractional Time Series Models

2020-09: Daniel Borup, Jonas N. Eriksen, Mads M. Kjær and Martin Thyrsgaard: Predicting bond return predictability

2020-10: Alfonso A. Irarrazabal, Lin Ma and Juan Carlos Parra-Alvarez: Optimal Asset Allocation for Commodity Sovereign Wealth Funds

2020-11: Bent Jesper Christensen, Juan Carlos Parra-Alvarez and Rafael Serrano: Optimal control of investment, premium and deductible for a non-life insurance company

2020-12: Anine E. Bolko, Kim Christensen, Mikko S. Pakkanen and Bezirgen Veliyev: Roughness in spot variance? A GMM approach for estimation of fractional log-normal stochastic volatility models using realized measures

2020-13: Morten Ørregaard Nielsen and Antoine L. Noël: To infinity and beyond: Efficient computation of ARCH(∞) models

2020-14: Charlotte Christiansen, Ran Xing and Yue Xu: Origins of Mutual Fund Skill: Market versus Accounting Based Asset Pricing Anomalies

2020-15: Carlos Vladimir Rodríguez-Caballero and J. Eduardo Vera-Valdés: Air pollution and mobility in the Mexico City Metropolitan Area, what drives the COVID-19 death toll?

2020-16: J. Eduardo Vera-Valdés: Temperature Anomalies, Long Memory, and Aggregation