



The Sun's Magnetic (Hale) Cycle and 27 Day Recurrences in the *aa* Geomagnetic Index

S. C. Chapman¹, S. W. McIntosh², R. J. Leamon^{3,4}, and N. W. Watkins^{1,5,6}¹ Centre for Fusion, Space and Astrophysics, Physics Department, University of Warwick, Coventry CV4 7AL, UK; S.C.Chapman@warwick.ac.uk² National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307, USA³ University of Maryland–Baltimore County, Goddard Planetary Heliophysics Institute, Baltimore, MD 21250, USA⁴ NASA Goddard Space Flight Center, Code 672, Greenbelt, MD 20771, USA⁵ Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, London WC2A 2AE, UK⁶ Faculty of Science, Technology, Engineering and Mathematics, The Open University, Milton Keynes MK7 6AA, UK

Received 2021 January 8; revised 2021 May 26; accepted 2021 May 27; published 2021 August 17

Abstract

We construct a new solar cycle phase clock which maps each of the last 18 solar cycles onto a single normalized epoch for the approximately 22 yr Hale (magnetic polarity) cycle, using the Hilbert transform of daily sunspot numbers (SSNs) since 1818. The occurrences of solar maxima show almost no discernible Hale cycle dependence, consistent with the clock being synchronized to polarity reversals. We reengineer the Sargent R27 index and combine it with our epoch analysis to obtain a high time resolution parameter for 27 day recurrence in *aa*, ($\langle acv(27) \rangle$). This reveals that the transition to recurrence, that is, to an ordered solar wind dominated by high-speed streams, is fast, with an upper bound of a few solar rotations. It resolves an extended late declining phase which is approximately twice as long on even Schwabe cycles as odd. Galactic cosmic ray flux rises in step with $\langle acv(27) \rangle$ but then stays high. Our analysis also identifies a slow-timescale trend in SSN that simply tracks the Gleissberg cycle. We find that this trend is in phase with the slow-timescale trend in the modulus of sunspot latitudes, and in antiphase with that of the R27 index.

Unified Astronomy Thesaurus concepts: Solar cycle (1487); Sunspot cycle (1650); Space weather (2037)

1. Introduction

Hale observed that the magnetic polarity of sunspots follows a roughly 22 yr cycle (Hale et al. 1919). This Hale (magnetic polarity) cycle spans two (approximately 11 yr long) Schwabe cycles of the sunspot number (Schwabe 1844). In the century since Hale's discovery there have been many works exploring the importance of the Hale cycle in developing a coherent picture of the climatology of solar activity that is intrinsically tied to our star's magnetic field (see, e.g., Hathaway 2015 and references therein).

The fact that the solar magnetic field undergoes a magnetic polarity reversal each Schwabe cycle is intrinsically linked to the concept of an “extended solar cycle” (Cliver 2014). In this extended picture, solar magnetic activity coherently originates at high latitudes ($\sim 55^\circ$) and migrates equatorward over nearly two decades (e.g., Wilson et al. 1988; McIntosh et al. 2021). About a decade after leaving high latitudes the pattern gives rise to the familiar “butterfly diagram” of sunspot latitudes (Maunder 1904). The pattern continues to move toward the equator where the oppositely polarized magnetic bands in each hemisphere eventually terminate (McIntosh & Leamon 2014; McIntosh et al. 2019). The Hale cycle is thus a fundamental mode of solar activity and is intrinsic to our understanding of the process that generates the magnetic field in the first place—the solar dynamo (see, e.g., McIntosh et al. 2021 and references therein).

At a coarse level the range of phenomenology exhibited as solar (and associated) activity approximately repeats every Schwabe cycle. A wealth of observations now exist for solar activity, but while they do suggest a multi-Schwabe cycle, and thus potentially Hale cycle climatology (e.g., Basu et al. 2012) they extend over only a few Schwabe cycles. However, there are a number of phenomena which are observed over a sufficient number of cycles to show a pronounced magnetic

polarity, or Hale cycle, dependence (see, e.g., Hathaway 2015). As well as sunspot magnetic polarity, these include Galactic cosmic ray flux (GCR; Jokipii et al. 1977; Jokipii & Thomas 1981) and how solar coronal activity translates into the level of disorder present in the solar wind and its space weather impact. There are records demonstrating the imprint of solar activity on solar wind structure and geomagnetic activity at Earth which reliably span a number of Hale cycles. Indeed, Sargent (1985) (Sargent 2021 has constructed a current version) used the longest continuous record of geomagnetic activity, the *aa* index (Mayaud 1972), to construct the “R27 index” of 27 day, solar rotation recurrence. This index peaks every Schwabe cycle but Sargent (1985; see also Cliver et al. 1996) identified a Hale cycle dependence in the shape of the peaks.

This paper focuses on observed quantities that span multiple Hale cycles and brings two innovations to bear that can extract climatology on that scale. First, the Sargent (1985) prescription for the R27 index was an algorithmic one; we present a reengineered version, built directly on the autocovariance of the *aa* index which is nowadays much more readily computed. Second, we recently showed (Chapman et al. 2020b) that the Hilbert transform of sunspot number time series can be used to map the irregular Schwabe cycle onto a regular normalized timebase, or solar cycle “clock.” Once constructed, this solar cycle clock cleanly organizes the cycle variation of solar flare statistics, the F10.7 index, and the measured geomagnetic activity at Earth. Applying this technique to the daily sunspot number (SSN) we construct a Hale cycle clock and use a superposed epoch analysis (Chree 1913) to investigate the organization of solar climatology on the Hale cycle.

In Section 2 we revisit Sargent's R27 index. R27 is an index based on a measure of autocorrelation of the *aa* index; however, since the underlying autocorrelation or autocovariance is highly variable, smoothing is required in order to discern underlying

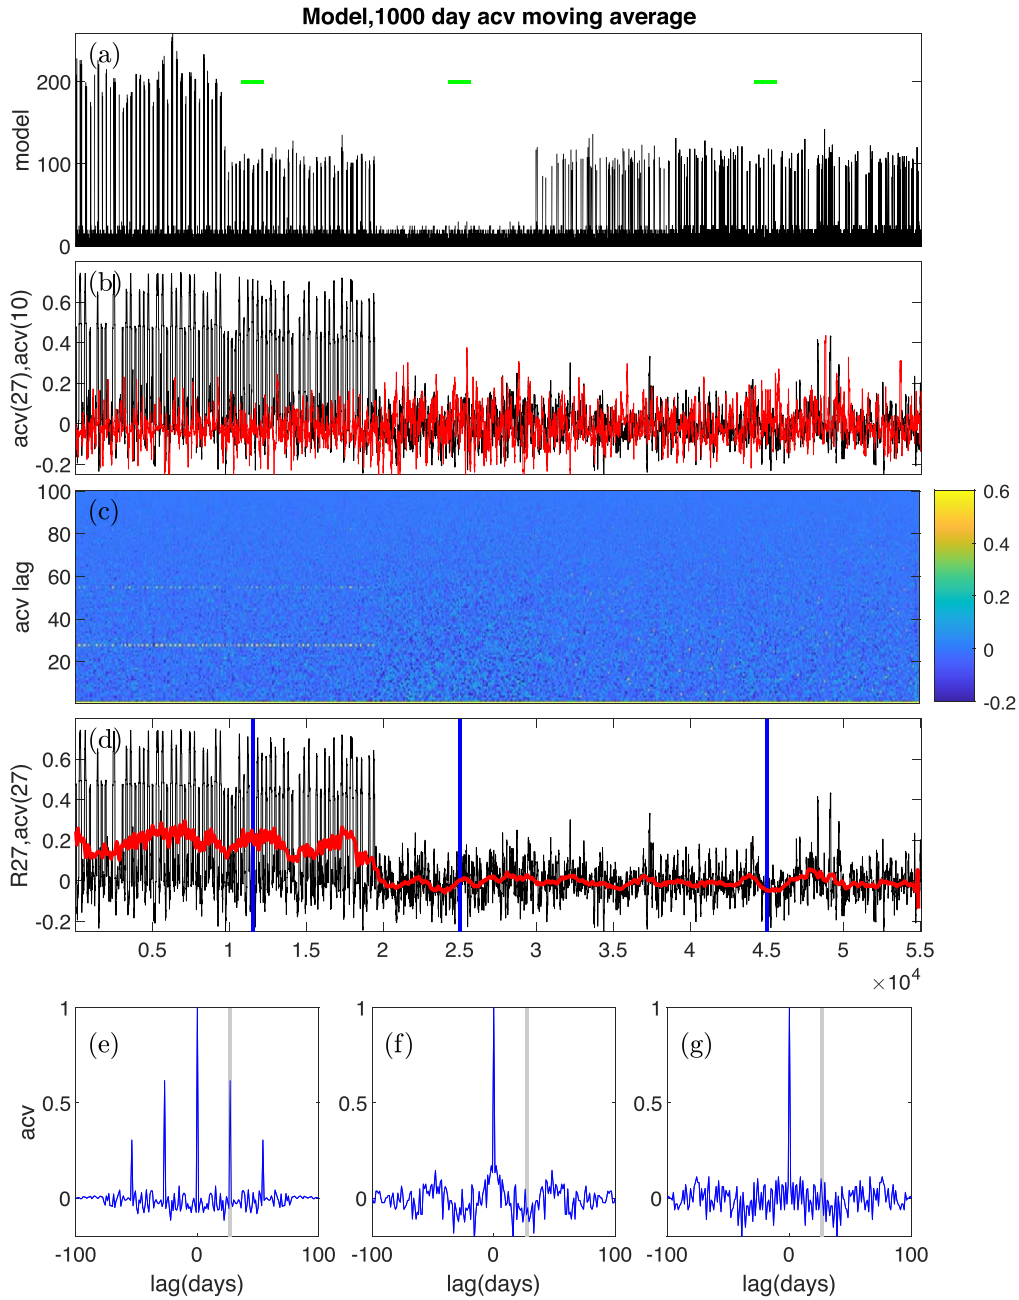


Figure 1. Model data test of the acv-R27 method. The panels plot (a) the modeled data, (b) the running 100 sample autocovariance at lag 27 (black) and at lag 10 (red), (c) the running sample autocovariance from lags $m = 5-100$ as a color map, (d) the running 100 sample autocovariance at lag 27 (black) and its 1000 point running mean (red); blue lines indicate the times for which samples (e)–(g) autocovariance are plotted. Green bars indicate 1500 point intervals which are shown enlarged in Figure 2.

trends in the time domain. In Section 3 we use the Hilbert transform of daily sunspot number to construct a Hale cycle clock. As discussed in Section 4, once the autocovariance of the aa index is mapped onto the Hale cycle clock, we can perform an averaging across successive Hale cycles rather than smoothing over time. This reveals that the transitions between disordered and 27 day recurrent structure in the aa index occur on a timescale that is less than or of order 2–3 solar rotations. The duration of this 27 day recurrent structure is almost twice as long during even cycles as odd ones, and we compare it with the Hale cycle variation of other multi-cycle observations.

The superposed epoch analysis of the revised R27 allows us to revisit the pioneering work of Mayaud (1980), Legrand & Simon (1981, 1989, 1991), and Simon & Legrand (1989).

This series of papers culminated in the deduction that the aa index is indicating that the solar cycle has two distinct components, (in the language of Legrand & Simon, “dipole” and “toroidal” components), which are out of phase relative to one another—the “toroidal” component exhibiting a 5–6 yr delay relative to the “dipole” component. In the discussion in Section 5 we consider their finding in the context of R27 and of the extended solar cycle. In this way we can start to unite the “light” (active regions, etc.) and “dark” (coronal holes, solar wind, etc.) sides of solar activity as natural and coupled manifestations of the Hale cycle. Such efforts may also go a long way to revealing *why* the aa index at solar minimum is an effective precursor of the upcoming maximum in SSN (Feynman 1982).




Figure 2. Enlargements of the three 1500 point long intervals shown in Figure 1 (green bars): (a) 27 day recurrent signal plus noise, (b) noise, (c) randomly occurring signal plus noise.

Finally, identification and subtraction of a slow-timescale trend in SSN is intrinsic to the Hilbert transform method (e.g., Chapman et al. 2020b and references therein). The slow-timescale trend in SSN that forms part of our analysis (presented in Section 6) simply tracks the Gleissberg cycle (Gleissberg 1967). We find that this trend is in phase with the slow-timescale trend in the modulus of sunspot latitudes, and in antiphase with the R27.

2. A Reengineered R27 Index

Sargent (1985) originally obtained the cross-correlation coefficients between successive 27 day intervals of the *aa* index (Mayaud 1972, 1980) and then performed truncation and smoothing to produce the original R27 index. We will base our analysis on the autocovariance of the *aa* index which is available from 1868 January 1. For a real-valued discrete signal x_i the raw (R_m) and normalized ($acv(m)$) covariance (Bendat & Piersol 2010) of a sequence with itself (i.e., the “autocovariance”) as a function of lag m is, for $m \geq 0$,

$$R_m = \sum_{n=0}^{W-m-1} \left(x_{n+m} - \frac{1}{W} \sum_{i=0}^{W-1} x_i \right) \times \left(x_n - \frac{1}{W} \sum_{i=0}^{W-1} x_i \right) \quad (1)$$

$$acv(m) = \frac{R_m}{R_0} \quad (2)$$

with the symmetry property that for $m < 0$, $acv(m) = acv(-m)$. In the above, the autocovariance is obtained for the sample

window $i = 1 \dots W$ of the x_i . We will consider two measures of the level of ordered structure in the solar wind, inferred from $acv(27)$, the strength of the autocovariance of the *aa* index at an $m = 27$ day autocovariance lag.

The first of these is a “reengineered” R27 index, hereafter “acv-R27,” that performs the same analysis as that of Sargent (1985, 2021) and in addition can be straightforwardly modeled and reproduced at each stage of its construction. Since, as we will demonstrate, the $acv(27)$ time series is highly variable, to extract any trend in it requires some averaging or smoothing in the time domain. We calculate the autocovariance of the daily average of the *aa* index at all lags m for a W day window centered on each day of the record. A running mean of length $S > W$ is then performed on the $m = 27$ day values, $acv(27)$, to obtain acv-R27. This procedure captures the essential operation of the original R27 index but has the advantage that, since all lags are calculated, we can (i) test that $acv(27)$ is indeed a maximum of the autocovariance and (ii) use one of the other lags, say $acv(10)$, as a check on the relative noise level in the time series; we can compare it to the standard error on the acv for white noise which is $\sim 1/\sqrt{W}$ (Box et al. 2016).

Our second measure will avoid smoothing in the time domain, desirable both to obtain higher time resolution and to avoid the possibility that smoothing can introduce spurious periodicities in finite-length, noisy time series (Slutsky 1937; Yule 1926). It will rely instead on averaging at the same Hilbert phase across successive Hale cycles. We will discuss how the Hilbert transform of daily SSN can be used to establish a Hale cycle epoch in more detail in Section 4.1; in this section we will model the process of averaging across epochs.

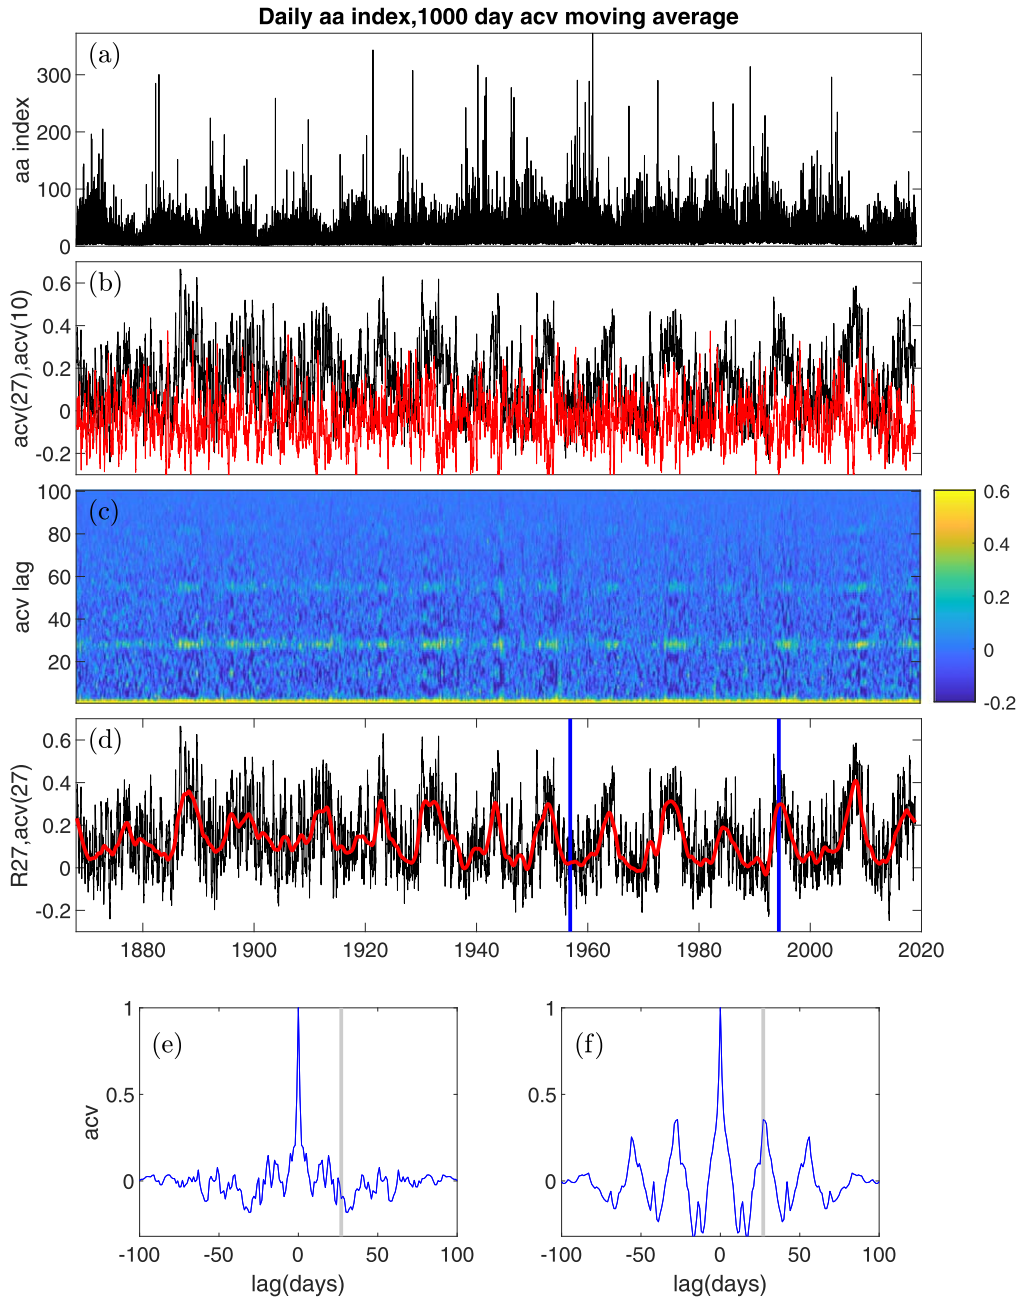


Figure 3. acv-R27 index based on the daily *aa* index. (a) Daily *aa* index, (b) daily autocovariance $acv(27)$ at lag 27 (black) and at lag 10 $acv(10)$ (red), (c) running sample daily autocovariance from lags $m = 5-100$ as a color map, (d) daily $acv(27)$ (black) and its 1000 point running mean $acv-R27$ (red), blue lines indicate times where individual daily autocovariance functions are plotted, for a time where the $acv-R27$ is relatively low (e) and high (f).

We will first use model data to test the properties of the autocovariance $acv(27)$, and the $acv-R27$ index and will then compute them using the *aa* index. The model is of a positive-definite sequence that incorporates (i) Poisson-distributed noise added to a sporadically occurring spike signal which occurs either (ii) in short 27 day recurrent runs or (iii) randomly. The model time series is plotted in panel (a) of Figure 1, and intervals indicated with green bars are shown expanded in Figure 2. We use a Poisson distribution with probability mass function $f_P(z|\lambda) = \frac{\lambda^z}{z!}e^{-\lambda}$, where z is an integer, to generate random samples $z_{k,\lambda}$ that are both discrete and positive definite, as is the *aa* index. In the model the noise is the sequence $5z_{k,1}$ with $\lambda=1$. The signal is the sequence $z_{k,100}$ with $\lambda=100$.

Each modeled “day” of data is then a randomly generated value for the noise, plus that of the signal if it has occurred on that day. Early in the model data sequence, the signal is organized into short runs of spikes with a 27 day recurrence as can be seen in panel (a) of Figure 2. These model recurrent co-rotating streams. The runs are randomly spaced, with a spacing $1 + z_{k,200}$ and $\lambda = 200$ and each run contains a random number of streams $1 + z_{k,2}$, $\lambda = 2$. The only non-random feature of these streams is their 27 day recurrence. For $t < 1000$ the amplitude of the signal is twice that in the rest of the sequence. Next in the model data there is an interval of noise only; see panel (b) of Figure 2. Finally for $t > 3000$ the signal is comprised of randomly occurring (uniformly distributed) spikes; see panel (c) of Figure 2.




Figure 4. Construction of the epoch-averaged $\langle acv(27) \rangle$ for model data with a sharp “switch on” of 27 day recurrences. The panels plot (a) one realization of the model data, (b) its running 100 sample autocovariance at lag 27 ($acv(27)$, black) and lag 10 (red), (c) the $acv(27)$ binned into non-overlapping 30 point averages, (d) an epoch average of $acv(27)$ obtained by averaging across seven realizations of the modeled data.

We now calculate $acv(27)$ and $acv\text{-R}27$ for these test data using an autocovariance window of $W=100$ and a running mean of length $S=1000$. Panel (a) of Figure 1 shows the test data. Panel (b) plots the running “daily” autocovariance for lag 27, $acv(27)$ (black) and lag 10, $acv(10)$ (red) and we see that $acv(27)$ discriminates the 27 day recurrent runs from the noise: $acv(27)$ is well above the noise estimated as $acv(10)$ which as expected has an amplitude at the standard error of $\sim 1/\sqrt{W} = 0.1$. The “daily” $acv(27)$ is also insensitive to the amplitude of the recurrent runs. Panel (c) plots the daily $acv(27)$ as a color map and we can see a peak at lag $m=27$ and more faintly at twice that value. Panel (d) plots $acv(27)$ (black) and its running mean (red) which is the $acv\text{-R}27$ index. We see that $acv\text{-R}27$ discriminates the 27 day recurrent runs with a smoothed value of ~ 0.2 . The bottom panels plot examples of the individual $acv(27)$ for all lags m , and $m=27$ is indicated in gray; it is clearly a maximum of the autocovariance where the modeled recurrent streams are present. Where there is an absence of recurrent streams we do not obtain a peak in the autocovariance except at zero lag as we would expect for δ -correlated random signals. This model data then confirms that $acv\text{-R}27$ has the essential properties of R27: it discriminates 27 day recurrences and is insensitive to amplitude. It has the

additional advantage that we can directly obtain a noise floor from the data, and can check that the $m=27$ lag autocovariance $acv(27)$ is indeed the dominant $m \neq 0$ peak.

We now use the same procedure to generate $acv\text{-R}27$ from the aa index. This is shown in Figure 3. We will start with daily aa values obtained by averaging the eight values of the three hourly aa record for each day, shown in panel (a). Using the same parameters as above, we obtain daily samples of $acv(27)$ shown in panel (b); we can see that the high values of the daily $acv(27)$ are significantly above that at $acv(10)$ which has an amplitude of order the standard error of $\sim 1/\sqrt{W} = 0.1$. The daily $acv(m)$ are plotted at all lags in panel (c) and we can see that the autocovariance has its main $m \neq 0$ peak at lag $m=27$, and a weaker peak at $m=2 \times 27$. Panel (d) plots the daily and smoothed $acv(27)$ which is the $acv\text{-R}27$ index (analogous to R27). Individual daily $acv(m)$ are plotted (panels (e), (f)) which again show high values of $acv\text{-R}27$ correspond to peaks in the autocovariance at $m=27$ whereas when $acv\text{-R}27$ takes a low value. There are no significant autocovariance peaks at $m \neq 0$; the acv fluctuates approximately at an amplitude of 0.1.

As well as constructing $acv\text{-R}27$ across the full time series we can perform an average of $acv(27)$ across Hale cycle epochs once these epochs are determined by Hilbert transform of daily

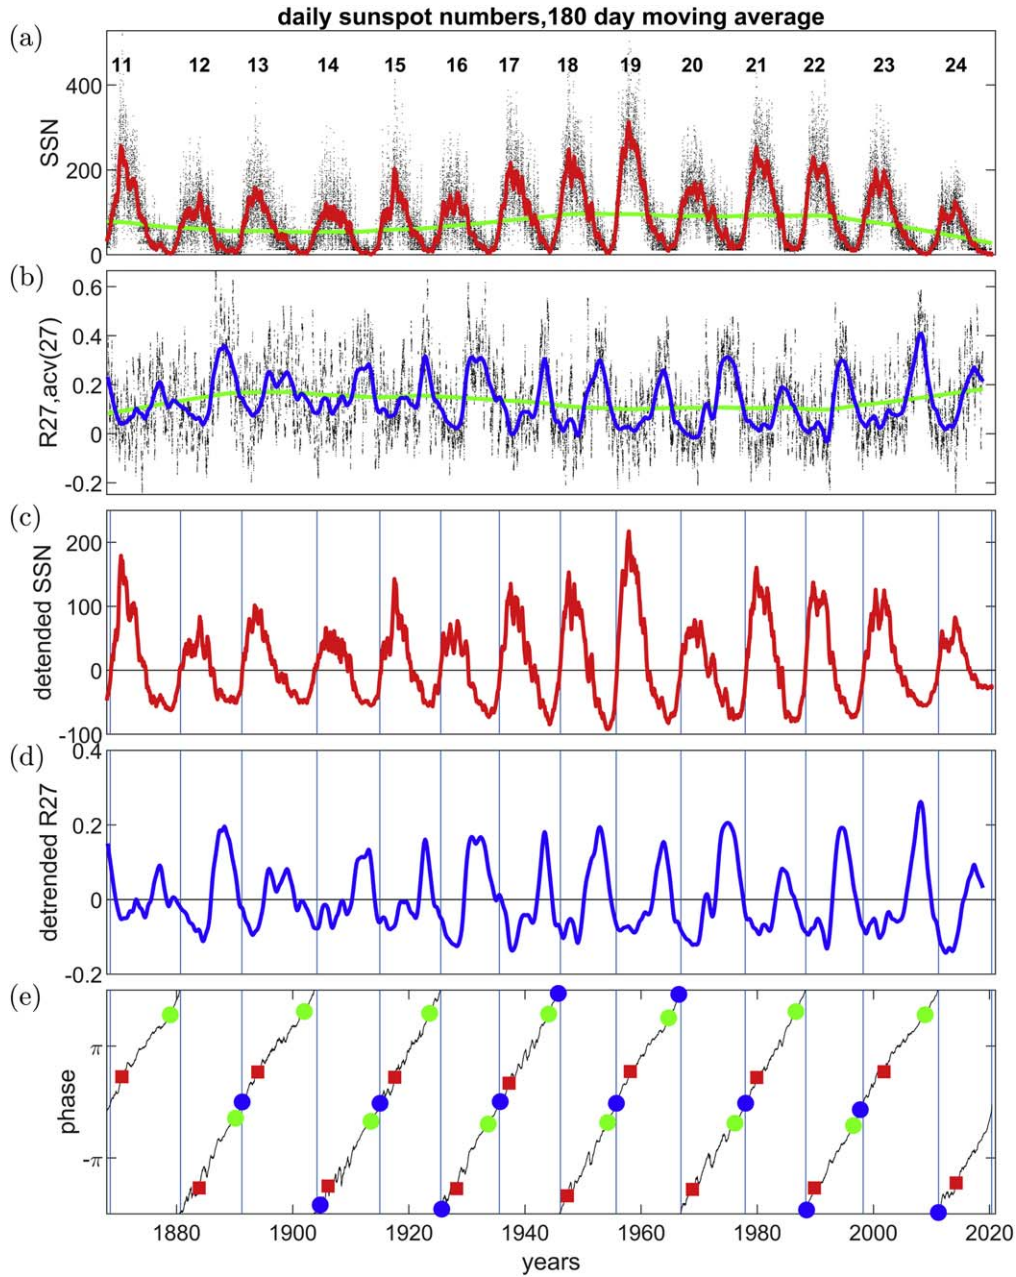


Figure 5. (a) Daily sunspot number (SSN) (black), 180 day moving average (red), and slow timescale trend obtained by local regression on a 40 yr window (green); (b) daily $acv(27)$ values (black), their 1000 day moving average $acv\text{-}R27$ (blue), and the slow timescale trend obtained by local regression on a 40 yr window (green); (c) smoothed SSN with slow timescale trend subtracted; (d) $acv\text{-}R27$ with slow timescale trend subtracted; (e) analytic phase of smoothed detrended SSN (black) plotted on a $[-2\pi, 2\pi]$ interval. The maxima and minima are indicated by red squares and green circles respectively and the blue circles indicate terminators from McIntosh et al. (2019). Zero phase is set to be at the average phase of the terminators. Solid blue vertical lines indicate terminator times of occurrence (see the text).

SSNs. A trial of this methodology on the model time series is shown in Figure 4. Each model time series is composed of 6000 “daily” values comprised of the same randomly generated noise as above, plus a random sequence of short runs of spikes with a 27 day recurrence that commences halfway across the time series. Figure 4(a) plots one realization of the model time series, and its autocovariance $acv(27)$ and $acv(10)$ are plotted in panel (b). The epoch average will be performed across non-overlapping month long bins and the result of averaging $acv(27)$ over these 30 day bins for this one realization is shown in panel (c). We have the aa index over seven Hale cycles and panel (d) plots the result of averaging $acv(27)$ over seven realizations of the model

time series. The rise time of the seven-epoch average, $\langle acv(27) \rangle$ then occurs on $\sim 2\text{--}3$ “monthly” bins, a timescale constrained by the $W=100$ “day” window over which $acv(27)$ is estimated.

3. Hilbert Transform—Building the Sun Clock

There have been several methods proposed to map successive Schwabe and Hale cycles onto a uniform epoch (see for example Thomas et al. 2014; Barnard et al. 2018). Chapman et al. (2020b) recently proposed a new method to order solar cycle variation on a timebase standardized to the solar cycle. The Hilbert transform of daily SSN is used to map

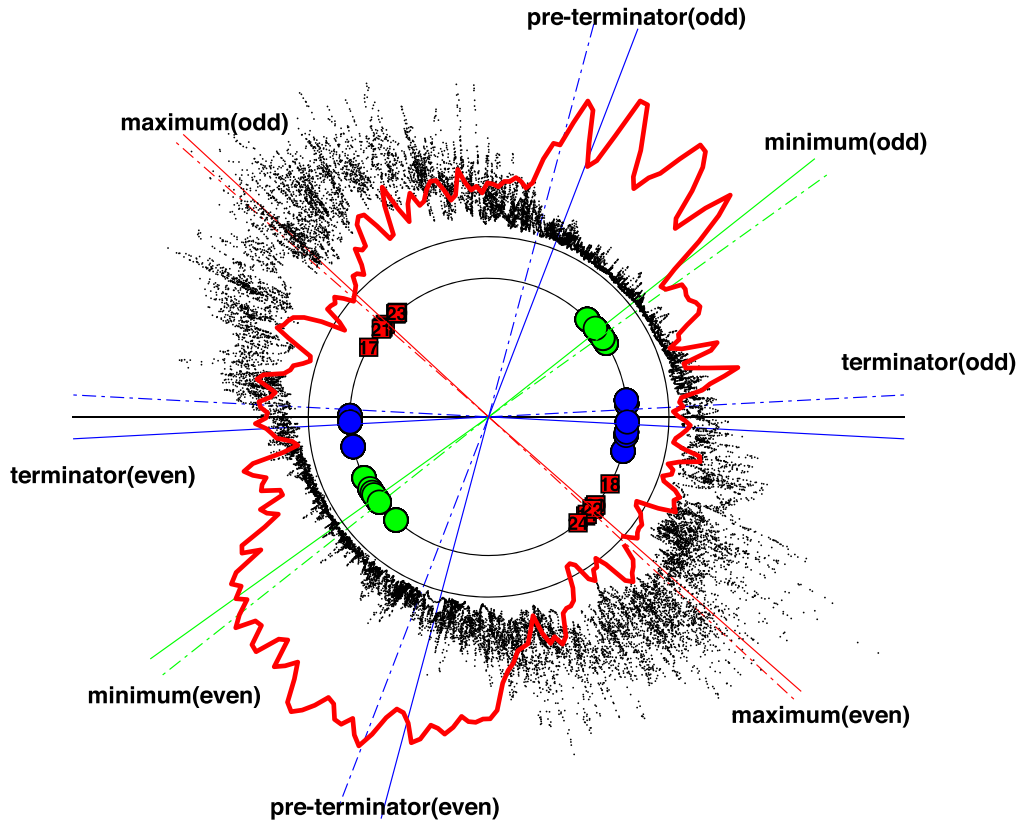


Figure 6. Hale cycle clock constructed such that increasing time (analytic phase) is read clockwise; $[0, 2\pi]$ in phase on the clock corresponds to the normalized 22 yr Hale cycle (phase $[0, 4\pi]$ from the Hilbert transform of daily SSN). The analytic phases of the maxima and minima of the last 18 solar cycles are indicated by red squares and green circles respectively and the blue circles indicate terminators for the last 12 solar cycles (McIntosh et al. 2019). Red, green, and blue solid lines indicate the average analytic phases for the maxima, minima, and terminators respectively; these averages are performed separately for odd and even Schwabe cycles. The pre-terminators (blue lines) are at phase $2\pi/5$ (4.4 normalized yr) in advance of the terminator. The solid lines indicating these average phases are continued across the clock with dashed-dotted lines; these indicate that the maxima occur regularly. The horizontal black line is located at the terminator occurrence averaged over all Schwabe cycles. These features form the basis of the Hale cycle clock. Black dots overplot daily F10.7. The red line is the $\langle acv(27) \rangle$ obtained by dividing the normalized 22 yr Hale cycle into 12×22 month long intervals and averaging the values of $acv(27)$ that fall within each of these intervals across multiple Hale cycles.

the irregular duration solar cycle in time onto a regular cycle in phase. We will use the same analysis as in Chapman et al. (2020b) to obtain a phase–time mapping from the daily SSN record onto a standardized Hale cycle. This mapping can then be used to (i) study the Hale cycle climatology of the acv -R27 index and (ii) perform an averaging across multiple standardized Hale cycles of $acv(27)$, giving a higher time resolution measure of the level of 27 day recurrent structuring of the solar wind, which we can then compare with the Hale climatology of other quantities available across multiple cycles.

We use the daily SSN record which provides an almost uninterrupted measure of solar coronal activity since 1818 January 1. It is plotted from the start of the aa index record (1868 January 1) in Figure 5(a). We can decompose this time series $S(t)$ in terms of a time-varying amplitude $A(t)$ and phase $\phi(t)$ by obtaining its analytic signal $A(t)\exp[i\phi(t)]$ (Gabor 1946; Boashash 1992) such that the real part of this signal is $S(t)$ and the imaginary part is obtained such that $A(t)\exp[i\phi(t)] = S(t) + iH(t)$ where $H(t)$ is the Hilbert transform of $S(t)$. This provides a mapping between time and signal phase, which converts the (variable) duration of each solar cycle into a corresponding uniform phase interval. The Hilbert transform takes the SSN time series as input and does not require particular markers for points in the cycle, such as the maxima and minima, to be identified; instead it provides the

phases of the maxima and minima as outputs. Chapman et al. (2020b) considered a mapping of each Schwabe cycle to phase in the range 0 to 2π . Here we will use this same mapping but will consider the Hale cycle that corresponds to the phase range 0 to 4π , that is, two successive Schwabe sunspot cycles.

While defined for an arbitrary time series, the analytic signal only gives a physically meaningful decomposition of the original time series if the instantaneous frequency $\omega(t) = d\phi(t)/dt$ remains positive (Boashash 1992). We therefore need to remove fast fluctuations and, for a positive-definite signal such as the daily SSN, a background trend. Following Chapman et al. (2020b), before performing the Hilbert transform we performed a $T_s = 180$ day moving average. We obtained a slow-timescale trend by performing a robust local linear regression which down-weights outliers (“rloless”) using a $T_B = 40$ yr window. As discussed in Chapman et al. (2020b) this analysis is robust across a reasonable range of T_s and T_B , so that the zero crossing times $\phi(t) = 0$ vary by less than ± 3 months when calculated over the range $T_s = [150360]$ days and $T_B = [2080]$ yr. We subtract the slow-timescale trend (green line in Figure 5(a)) to give a sunspot time series that is unambiguously zero-crossing (Figure 5(c)). We then obtain the Hilbert transform $H(t)$ for this smoothed and detrended signal which then gives the analytic signal. We obtain the slow-timescale trend for acv -R27 in exactly the same manner. These




Figure 7. Hale cycle clock constructed such that increasing time (analytic phase) is read clockwise; $[0 - 2\pi]$ in phase on the clock corresponds to the normalized 22 yr Hale cycle (phase $[0 - 4\pi]$ from the Hilbert transform of daily SSN). The format is as in the previous figure except that activity in the aa index is overlotted on the $(acv(27))$. Black dots plotted at successively increasing radii indicate days where the daily maximum aa index exceeded 100, 200, 300, 400, 500, 600 nT.

slow-timescale trends are also an aspect of solar climatology which we will consider in the context of the Gleissberg cycle in Section 6. The acv -R27 index is plotted in Figure 5(d) with this slow-timescale trend subtracted. Panel (e) plots the Hilbert analytic phase obtained from the daily SSN record; this is wrapped to a domain of a $[-2\pi, +2\pi]$ Hale cycle. Overplotted are the maxima (red squares), minima (green circles), and terminators (blue) of each Schwabe cycle. Zero ($\pm 2\pi$) phase is set to the average phase of the terminators and vertical blue lines are drawn at the times when the analytic phase crosses zero and $\pm 2\pi$ to demarcate one Schwabe cycle from the next. Panel (d) shows the same suggestion of a Hale cycle climatology in acv -R27 as that originally identified by Sargent (1985) in R27; alternate Schwabe cycles have longer/shorter peaks in acv -R27. Comparing panels (d) and (e) we also find a new result: that the downward sweeps in acv -R27 correspond to the observed terminator times identified in McIntosh et al. (2019). Termination of each Schwabe cycle is correlated with switch-on of disorder in the solar wind. However, as discussed above, acv -R27 and indeed the original R27 index are by necessity smoothed in time. We will overcome this limitation by mapping the observations over each cycle onto a uniform “clock.” While the cycle lengths are irregular, panel (e) shows how they can be mapped to a regular interval in phase. One can either construct a mapping based on the Schwabe cycle as in Chapman et al. (2020b) or, as we will do here, construct a mapping based on successive pairs of Schwabe cycles, that is, on Hale cycles, which corresponds to a mapping between irregular cycles in time, and a regular 4π interval in analytic phase.

4. Hale Cycles—The 22 yr Sun Clock

We now use the mapping between time and phase plotted in Figure 5(e) to construct plots of solar climatology on a regular, normalized Hale cycle. An example of the resulting Hale cycle clock is shown in Figure 6, which overplots successive Hale cycles on a $[0, 2\pi]$ interval corresponding to 22 normalized yr. This is mapped from $[0, 4\pi]$ in Hilbert analytic phase of the daily SSN record, that is, two successive Schwabe cycles.

The sunspot maxima and minima along with the terminator occurrences from the previous figure are plotted on the clock. The averages of their locations are indicated by solid lines; these averages are obtained independently over odd and even Schwabe cycles. The solid lines are extended as dashed lines across to the opposite half of the clock, so that the latter indicate where the average maxima, minima, and terminators would occur if they repeated exactly from one Schwabe cycle to the next. In interpreting these averages we stress that they are obtained in the regime of small sample number, $N = 9$ for the maxima and minima and $N = 6$ for the terminators. For a small sample number the standard error does not simply translate to an estimate of uncertainty. However the spread in observed values can be clearly seen on the plot. We then find that the even cycle maxima average is very close to π in phase in advance of the odd cycle maxima average; they differ by ~ 0.03 rad or, for a normalized 22 yr Hale cycle, ~ 0.6 normalized yr. The polarity reversal of the solar polar field is known to occur close to the solar cycle maxima (Thomas et al. 2014) so that this is consistent with the idea that the clock does indeed provide a robust epoch analysis tool for the Hale cycle. The daily F10.7 index since 1947 February 14 (Tapping 2013) is

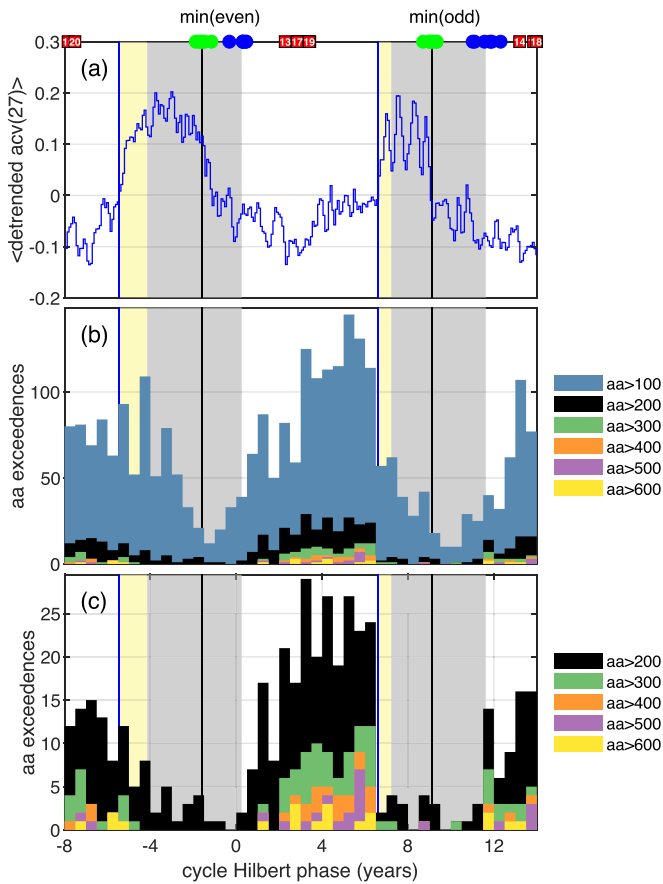


Figure 8. Hale cycle clock phase plotted linearly; the normalized 22 yr Hale cycle (phase $[0-4\pi]$ from the Hilbert transform of daily SSN). The analytic phases of the maxima and minima of the last 18 solar cycles are indicated by red squares and green circles respectively and the blue circles indicate terminators for the last 12 solar cycles (McIntosh et al. 2019). Black lines indicate the average analytic phase for the minima. The quiet interval around minimum, shaded gray, is of 4.4 normalized yr duration starting at the pre-terminator and ending at the average phase of the terminators. These averages are performed separately for odd and even cycles. The late-declining phase (see the text) is shaded yellow. Panel (a): $\langle acv(27) \rangle$ (blue trace) is obtained by dividing the normalized 22 yr Hale cycle into 12×22 month long intervals and averaging the values of $acv(27)$ that fall within each of these intervals over multiple Hale cycles. Blue vertical lines are at the zero upcrossings; panel (b): six (normalized) month binned counts of days where the maximum daily aa index value exceeded 100, 200, 300, 400, 500, 600 nT; panel (c): counts where the maximum daily aa index value exceeded 200, 300, 400, 500, 600 nT. Zero and 11 normalized yr on this plot correspond to the phase indicated by the black line in Figures 6 and 7.

overplotted (black dots) and we can see that, as identified in Leamon et al. (2020a) and Chapman et al. (2020b), the terminators coincide with an increase in activity seen in F10.7 as the activity of the next cycle commences (McIntosh & Leamon 2014; McIntosh et al. 2019). Each Schwabe cycle “switch-on” of activity at the terminators is preceded by a “switch-off” of activity, and Chapman et al. (2020b) identified this pre-terminator as being approximately 4.4 normalized yr in advance of the terminator which on this 22 yr Hale cycle clock is $2\pi/5$ in phase. We emphasize that this is only an estimate of where the pre-terminator should be located, based upon observed correspondence with changes in multiple measures of solar activity (Chapman et al. 2020b; Leamon et al. 2020b), such as low values of the F10.7 index as can be seen in Figure 6.

4.1. Hale Cycle Structure of the aa Index

Sargent (1985) noted that enhancements in the R27 index indicating recurrent behavior are longer (shorter) on alternate Schwabe cycles, which we can also see in acv -R27 in Figure 5. We can now use the Hale cycle clock to perform a direct average of $acv(27)$ which avoids smoothing in the time domain, tracking the variation in 27 day recurrences at higher time resolution. We divide the normalized 22 yr Hale cycle into 264 equal-width bins in phase (that is, 12×22 month long intervals) and average the values of $acv(27)$ that fall within each of these intervals over multiple Hale cycles. This is then detrended by subtracting the same (40 yr window, rowless) slow-timescale trend as for acv -R27. The resulting $\langle acv(27) \rangle$ will have a time resolution dictated by the time window over which the autocovariance is calculated, which is 100 days here. The $\langle acv(27) \rangle$ is plotted on Figure 6 (red line) and it reveals a clear Hale cycle dependence of 27 day recurrence in the aa index. The $\langle acv(27) \rangle$ peaks during quiet intervals of solar activity as seen in F10.7 but the interval over which $\langle acv(27) \rangle$ is enhanced is almost twice as long during even cycles compared to odd ones. The $\langle acv(27) \rangle$ sharply increases around the pre-terminator in odd cycles and before it in even cycles, then decreasing at minimum, decaying to its low value at the terminator. The timescales for these sharp changes determined here are at the time resolution of the autocovariance window, that is about three to four solar rotations. The physical timescale of these changes could be faster than this.

As established in our modeling in Section 2 (compare panels (a), (b), and (d) of Figure 1), both $\langle acv(27) \rangle$ and acv -R27 are constructed to be sensitive to the time structure, i.e., recurrences, in the signal while being independent of the signal amplitude. Thus any Hale cycle variation does not simply reflect the overall level of activity in the aa index. Furthermore, the aa index (units, nT) is discretized in amplitude (Bubenik & Fraser-Smith 1977; Chapman et al. 2020a) since the underlying K index (Bartels et al. 1939) is a quasi-logarithmic 0–9 integer scale that characterizes the maximum magnetic deviations that occur during each three hour period at a given observatory. Therefore we focus on the occurrences of aa exceeding a threshold value, rather than the absolute values of aa , or its time averages.

Figure 7 plots $\langle acv(27) \rangle$ on a Hale cycle clock as in Figure 6, and black dots are plotted at successively increasing radii to indicate days where the daily maximum aa index exceeded thresholds $T_{aa} = 100, 200, 300, 400, 500, 600$ nT. Extreme space weather events are seen as radial “spokes” on this plot. The most extreme events occur in a disordered sequence with relatively low values of $\langle acv(27) \rangle$. However, more moderate $aa \lesssim 200$ nT events do occur when $\langle acv(27) \rangle$ takes a high value. The occurrence rates are plotted in Figure 8, which shows histograms of counts of days within non-overlapping normalized six month bins in which the aa index exceeds T_{aa} . Panel (a) of Figure 8 plots $\langle acv(27) \rangle$ and vertical blue lines are drawn at the upcrossings of $\langle acv(27) \rangle = 0$. The vertical blue lines indicate the transition from an aa index signal that is disordered to one that contains 27 day recurrences, and $\langle acv(27) \rangle$ rises steeply at these transitions. It is well known that geomagnetic storms are preferentially triggered by irregularly occurring coronal mass ejections (CMEs) around solar maximum, whereas they are likely to be triggered by corotating interaction regions (CIRs) in the declining phase (Richardson 2018; Pulkkinen 2007); the CIR-driven events are more

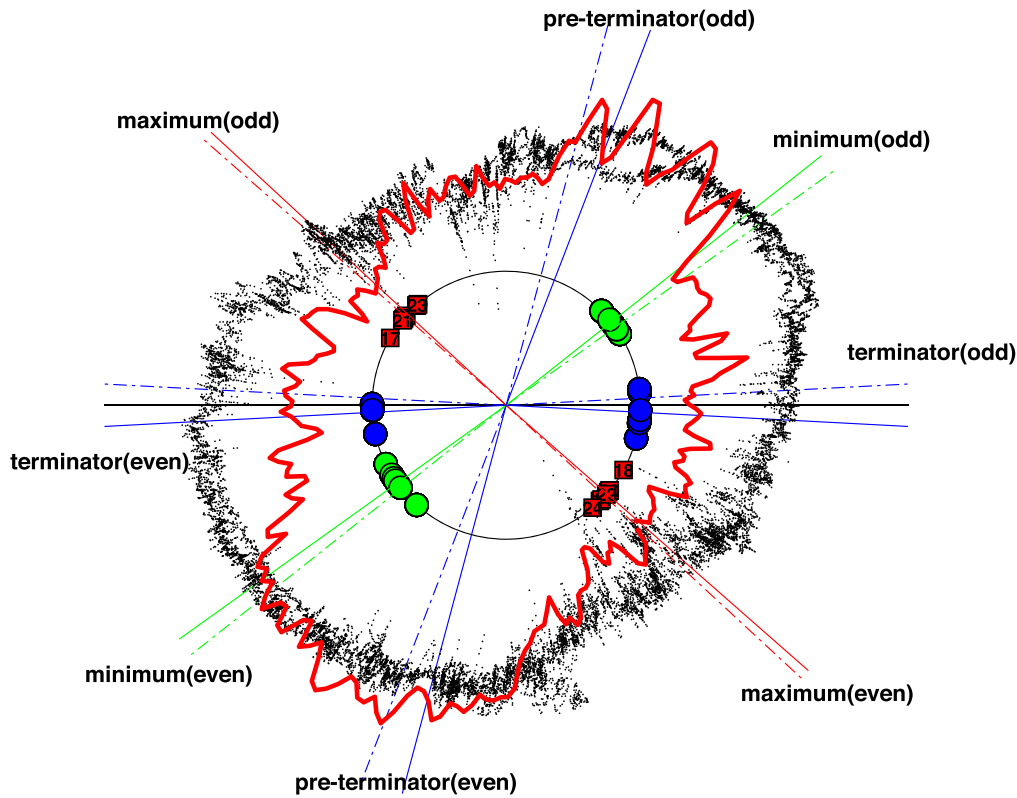


Figure 9. Hale cycle clock constructed such that increasing time (analytic phase) is read clockwise; $[0, 2\pi]$ in phase on the clock corresponds to the normalized 22 yr Hale cycle (phase $[0, 4\pi]$ from the Hilbert transform of daily SSN). The $\langle acv(27) \rangle$ (red) is overplotted on the daily detrended cosmic ray flux (black).

moderate than the CME-driven ones (Alves et al. 2006; Borovsky & Denton 2006). Our analysis is consistent with this picture and in addition reveals that the transition between these two behaviors is fast: from the time interval just following solar maximum where $\langle acv(27) \rangle$ is at its low value, to a late-declining phase (shaded yellow in Figure 8) where $\langle acv(27) \rangle$ rises steeply to its high value and geomagnetic storms decrease in their intensity. The frequency of occurrence of more moderate events does not decrease until around solar minimum. The duration of the late-declining phase shows a Hale cycle dependence; this is approximately twice as long for even cycles (~ 1.3 normalized yr) as for odd ones (~ 0.6 normalized yr).

4.2. Hale Cycle Structure in GCR Flux

GCR flux is also known to show a 22 yr pattern, with even cycles showing a longer, “flat top” enhancement compared to odd cycles. This has been attributed to different particle drift patterns when the northern solar pole has predominantly positive or negative polarities (Jokipii et al. 1977; Jokipii & Thomas 1981; Smith & Thomas 1986; Smith 1990; Ferreira & Potgeiter 2004). During positive polarities, cosmic ray protons reach Earth after approaching the poles of the Sun in the inner heliosphere and moving out along the heliospheric current sheet (HCS). During negative polarities, cosmic ray protons approach the Sun along the HCS plane and leave via the poles. Heliospheric modulation, both of the HCS and of the structure of the solar wind, has been found to play a significant role particularly in the declining phase of the solar cycle (Thomas et al. 2014 and references therein). Plots of annual mean SSN versus the annual mean GCR intensity are typically used to characterize the Hale cycle variation of GCR flux (see, e.g.,

Ross & Chaplin 2019 and references therein). Here, we can directly track GCR flux across the normalized Hale cycle. Since $\langle acv(27) \rangle$ is a measure of recurrence in the *aa* index, it directly depends on recurrent structure in the solar wind, so that we can infer that high values of $\langle acv(27) \rangle$ correspond to a solar wind dominated by 27 day recurrent high-speed streams, whereas low values of $\langle acv(27) \rangle$ correspond to a solar wind that is disordered.

Figures 9 and 10 plot the daily Oulu GCR flux (since 1964 April 30; see Usoskin et al. 2001) which given a low geomagnetic cutoff rigidity is sensitive to energetic particles that are significantly affected by plasma structures in the solar wind. It is plotted following subtraction of its slow-timescale (40 yr window rolloff) trend, and $\langle acv(27) \rangle$ on the Hale cycle clock. In Figure 8 we can see that the rises in GCR flux and $\langle acv(27) \rangle$ roughly track each other in the interval between the sunspot maximum and minimum, that is, through the declining phase. The GCR flux departs from the $\langle acv(27) \rangle$ at sunspot minimum, where $\langle acv(27) \rangle$ has passed its maximum value. The first half of the interval of enhanced GCR flux thus coincides with a solar wind dominated by recurrent streams, whereas the second half does not.

Figure 10 shows that that the increase in GCR flux is aligned with that in $\langle acv(27) \rangle$, particularly in even cycles. The increase in GCR flux during the declining phase thus coincides with a solar wind dominated by recurrent streams, and as suggested by Thomas et al. (2014) this may dominate over the polarity effect on proton drift paths during the declining phase. The comparison with $\langle acv(27) \rangle$ on a normalized 22 yr cycle shows how clearly demarcated these different intervals are. If we estimate the length of the interval of enhanced GCR flux as

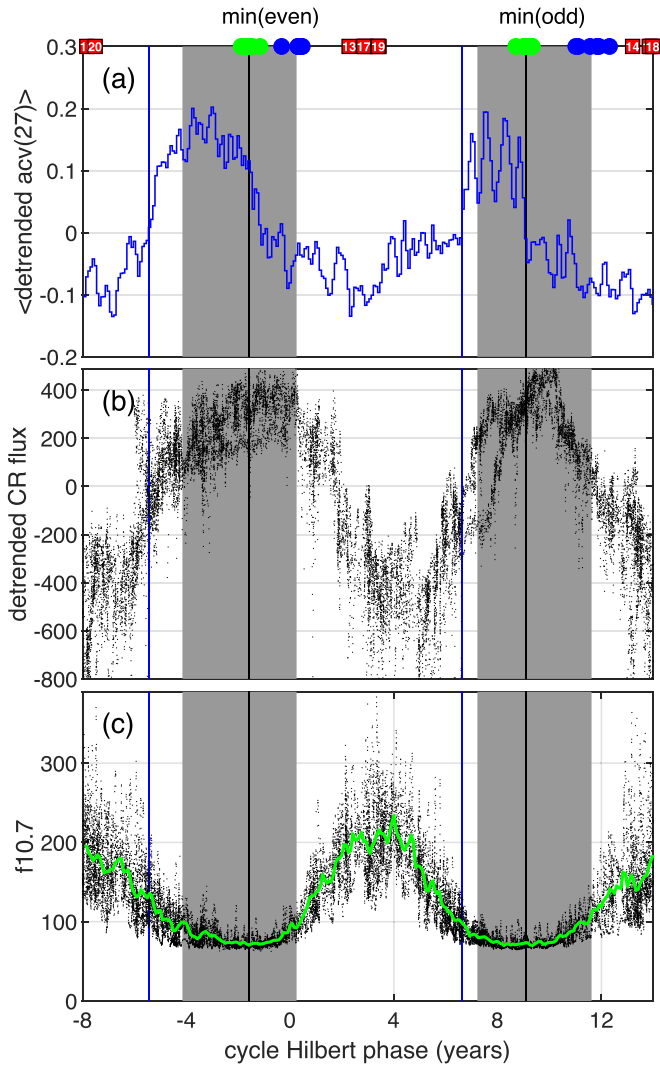


Figure 10. Hale cycle clock phase plotted linearly; the normalized 22 yr Hale cycle (phase $[0, 4\pi]$ from the Hilbert transform of daily SSN). The format follows that of Figure 7. Panel (a): $\langle acv(27) \rangle$ (blue trace) with blue vertical lines at the zero upcrossings; panel (b): detrended daily cosmic ray flux (black); panel (c): daily F10.7 (black) and its average (green) obtained by dividing the normalized 22 yr Hale cycle into 12×22 month long intervals and averaging the values of daily F10.7 that fall within each of these intervals over multiple Hale cycles.

times when the detrended GCR flux is above zero in Figure 10, we can see that in even cycles it has approximately twice the duration (~ 7.2 normalized yr) as in odd cycles (~ 4.5 normalized yr). The longer duration of the late-declining phase in even cycles identified in Figure 8 corresponds to ~ 1.3 normalized yr of enhanced GCR flux during even cycles. The GCR flux in even cycles only begins to decline after the terminator, crossing zero in Figure 10 at ~ 1.6 normalized yr after the terminator has occurred. Thus the extended duration “flat top” of GCR flux seen in even cycles compared to odd ones may be partially but not fully explained by the extended duration of the interval of enhanced recurrent structure dominating the solar wind as seen in high values of $\langle acv(27) \rangle$.

4.3. Hale Cycle and Sunspot Latitudes

The Greenwich record of sunspot latitudes and areas (1874 May 1–2016 September 30) provides a record of sunspot

activity extending over multiple Hale cycles which we directly compare with $\langle acv(27) \rangle$ in Figure 11. On this figure, panel (a) plots $\langle acv(27) \rangle$ as in previous figures. The modulus of the latitudes of sunspot centroids recorded each day are averaged, and the sequence of these daily averages are plotted in panel (b) (black dots). We obtained $\langle acv(27) \rangle$ by dividing the normalized 22 yr Hale cycle into 12×22 month long intervals and averaging the values of $acv(27)$ that fall within each of these intervals over multiple Hale cycles. We now perform the same operation on the daily averages of sunspot latitudes and this is plotted as the green trace on panel (b) of Figure 11 (we will denote this average sunspot latitude trace as $\langle SSL \rangle$). In Owens et al. (2011) the average modulus sunspot latitude was proposed as a useful parameterization of the Schwabe cycle as it shows a sudden increase at solar minimum due to the emergence of high-latitude sunspot pairs of new cycle polarity. The same overall sawtooth pattern identified by Owens et al. (2011) can be seen here in $\langle SSL \rangle$; it up-crosses latitude $\sim 15^\circ$ at the minima. However $\langle SSL \rangle$ also shows a Hale cycle dependence; the sawtooth pattern is more sharply defined for even cycles than for odd ones, and while both the pre-terminators correspond to $\langle SSL \rangle$ down-crossing latitude $\sim 10^\circ$, the terminators occur when $\langle SSL \rangle$ is at a higher latitude, $\sim 22^\circ$, for even cycles compared to odd ones, $\sim 20^\circ$. Panel (c) plots the sunspot latitudes as a classic butterfly diagram. Large sunspot areas are not seen in the quiet interval around minimum (gray), as we would expect. The largest switch off occurs at the start of the declining phase as identified by the sharp rises in $\langle acv(27) \rangle$ (blue vertical lines). This is consistent with a relatively sharp transition from a disordered solar wind, populated with CMEs which result in large $aa > 200$ geomagnetic storms, to a more ordered solar wind where storms, driven by recurrent high-speed streams, are more moderate.

5. Mayaud, Legrand, Simon, and the Hale/Extended Cycle

The analysis provided above allows us to briefly revisit, and contextualize, the pioneering work of Mayaud (1980), Legrand & Simon (1981, 1989, 1991), and Simon & Legrand (1989) (see also Feynman 1982). This series of papers culminated in the deduction that the aa index is indicating that the solar cycle has two distinct components which are out of phase relative to one another, one associated with active Sun and the other with quiescent (recurrent) high-speed solar wind streams. Their hypothesis was that the bimodal behavior of the aa index was consistent with the Sun exhibiting a “two-component cycle.” The earlier component in their view, the “dipole” component, occurs at mid-to-high latitudes and is exhibited some 5–6 yr before the second, the “bipolar” component, which originates in the sunspot-bearing latitudes before repeating.

Figure 12 overlays the trend-removed time series of the SSN and $acv-R27$ from Figure 5 and we can see that they are in antiphase. The $acv-R27$ can be seen to rise to its peak around or just after the pre-terminator regardless of its shape. This is the epoch of long-lived mid-latitude coronal holes (Krista et al. 2018; Hewins et al. 2020). Those coronal holes are the sources of Mayaud, Legrand, and Simon’s dipole component recurrent high-speed streams and belong to the magnetic band of the extended solar cycle that becomes host to the sunspots that follow some 5 or 6 yr later (see Figure 11(c), or McIntosh et al. 2021). Our analysis, when taken in concert, validates the insightful work of Mayaud, Legrand, and Simon.

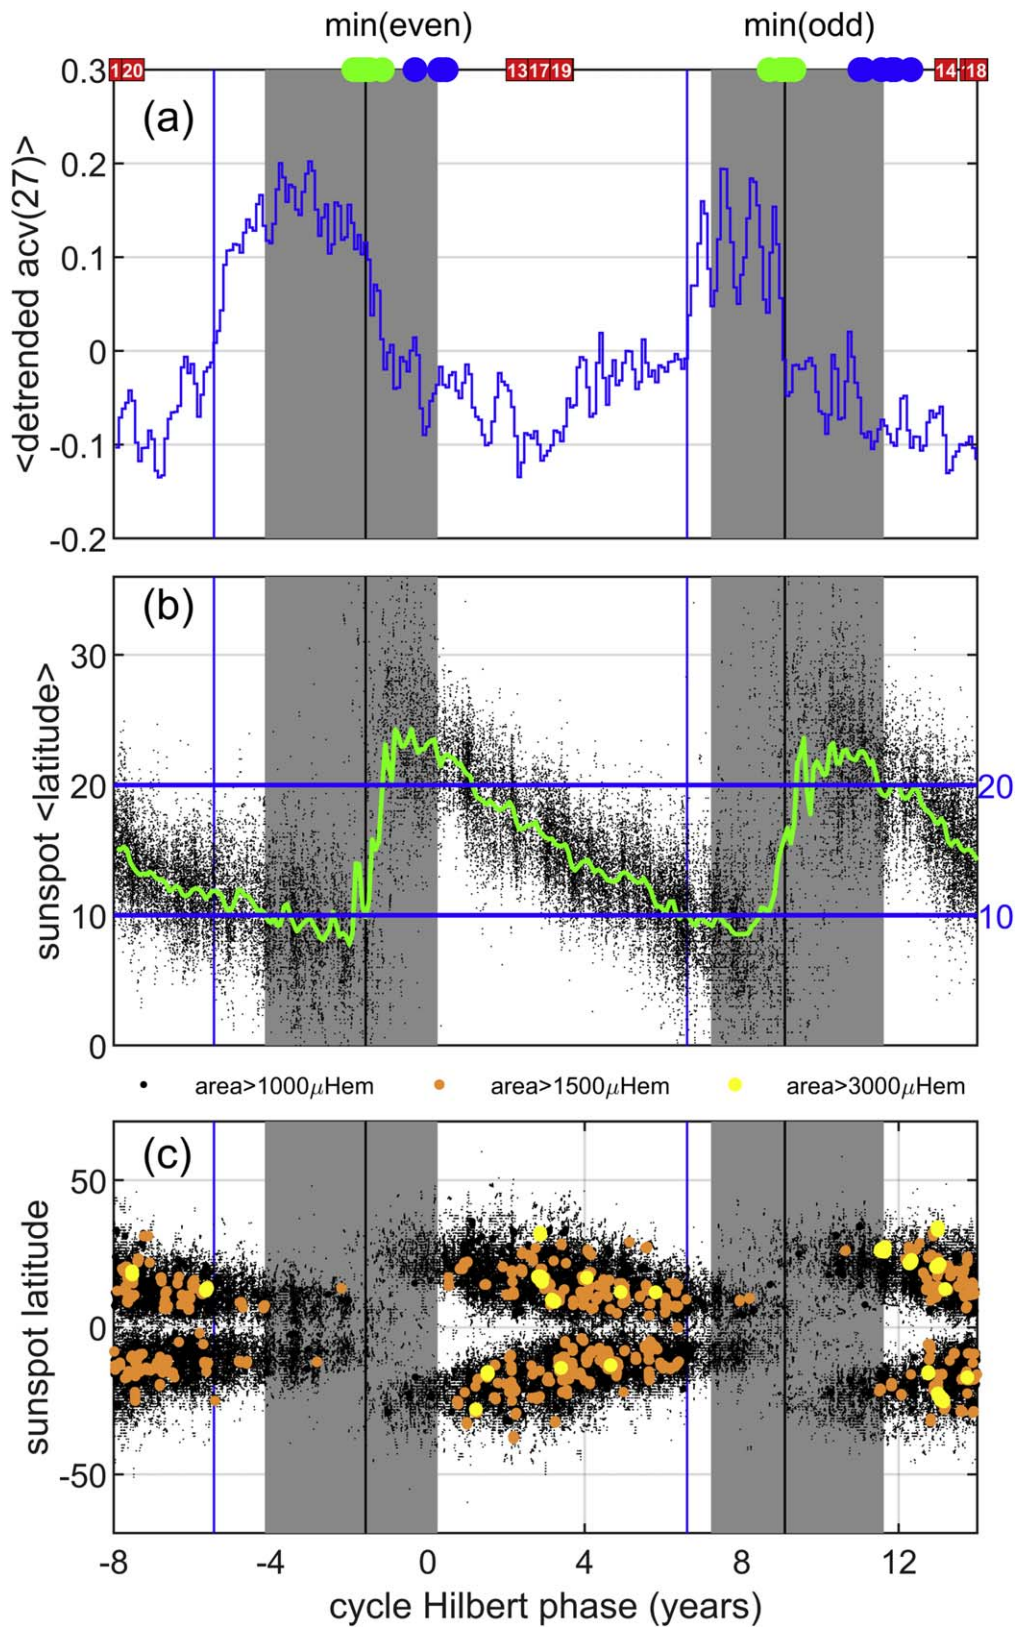


Figure 11. Hale cycle clock phase plotted linearly; the normalized 22 yr Hale cycle (phase $[0, 4\pi]$ from the Hilbert transform of daily SSN). The format follows that of the previous figure. Panel (a): $\langle \text{acv}(27) \rangle$ (blue trace) with blue vertical lines at the zero upcrossings; panel (b): daily averaged modulus of sunspot area centroid latitudes (black) and, overplotted, its multi-Hale cycle variation (green) obtained by dividing the normalized 22 yr Hale cycle into 12×22 month long intervals and averaging the (modulus) sunspot area latitudes that fall within each of these intervals over multiple Hale cycles. Panel (c): daily latitude positions of the sunspot area centroids; symbol colors and sizes differentiate areas $>1000, 1500, 3000 \mu\text{Hem}$.

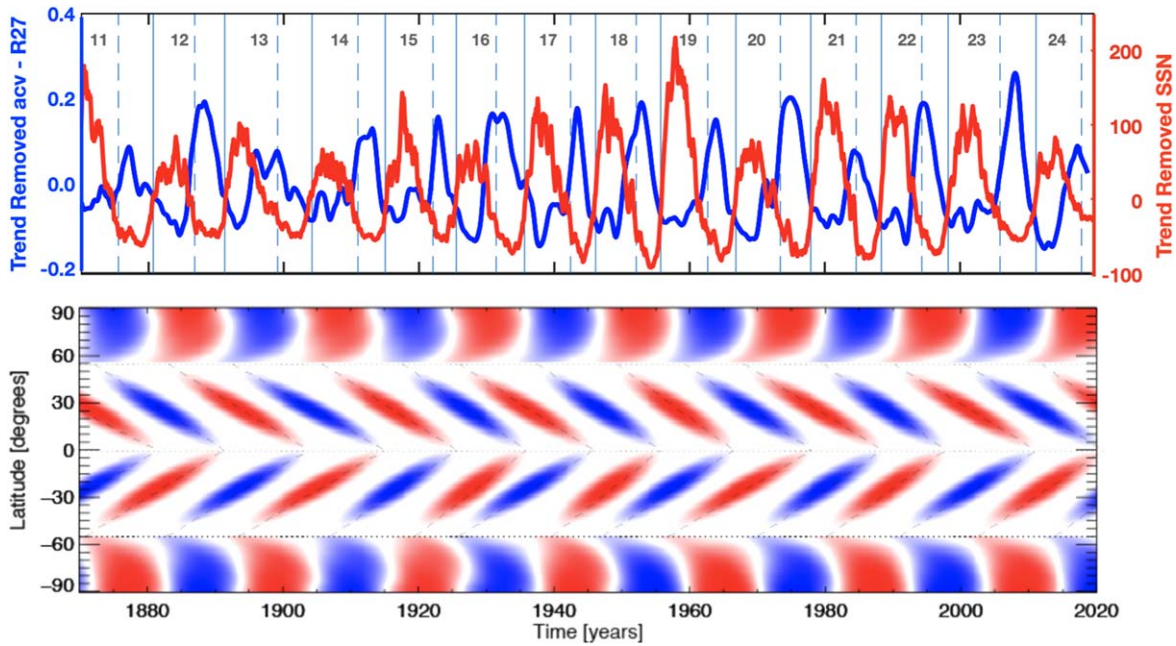


Figure 12. Placing the (trend-subtracted) SSN in context with the (trend-subtracted) acv-R27 and the “extended solar cycle” as a representation of the Hale cycle. The upper panel overlays the smoothed SSN with slow-timescale trend subtracted (red) and acv-R27 with slow-timescale trend subtracted (blue). The lower panel plots the data-motivated Hale cycle “band-o-gram” of McIntosh & Leamon (2014) and McIntosh et al. (2019, 2021). The pre-terminators and terminators are indicated by vertical dashed and solid blue lines respectively.

Going a step further, the temporal phasing of the dipolar and bipolar component signatures and the resulting correspondence with the magnetic systems of the Hale cycle provide insight into why the aa index around sunspot minimum is a stronger than average precursor of the *upcoming* sunspot cycle strength (see, e.g., Feynman 1982). The pre-terminator recurrent high-speed solar wind streams originate from the *same* Hale cycle magnetic bands as the sunspots that follow only a few years later, tying the “dark” and “light” sides of the Hale cycle together and illustrating the *why* of the strong precursor relationship, although the *how* requires further work.

6. Slow-timescale Trends and the Gleissberg Cycle

In the above analysis we subtracted a slow-timescale trend obtained by a robust local linear regression which down-weights outliers (“rlowess”) using a $T_B = 40$ yr window from the SSN, acv-R27 (Figure 5) and GCR flux (Figures 9 and 10). As part of its construction, $\langle acv(27) \rangle$ has the same trend subtracted as acv-R27. We did not subtract a slow-timescale trend from the average sunspot latitude (Figure 11) as it is only $\sim 2^\circ$ in amplitude; however, we will consider it here. These slow-timescale trends are compared to the Gleissberg cycle in Figure 13. The Gleissberg cycle $G^{(N)}$ is defined (Gleissberg 1967) in terms of successive solar maxima as

$$G^{(N)} = \frac{1}{8}(R_{\max}^{(N-2)} + 2R_{\max}^{(N-1)} + 2R_{\max}^{(N)} + 2R_{\max}^{(N+1)} + R_{\max}^{(N+2)}) \quad (3)$$

where $R_{\max}^{(N)}$ is the SSN at the maximum of the N th solar cycle. We used values for $R_{\max}^{(N)}$ determined by SILSO. The $G^{(N)}$ are plotted on panels (a)–(c) of Figure 13 (blue lines, symbols). Panel (a) overplots the SSN trend which can be seen to track the Gleissberg cycle so that the slow-timescale trend that we have removed in order to obtain the Hilbert transform of SSN is

just that of the Gleissberg cycle. Panel (b) shows that there is a slow-timescale trend in acv-R27 that is approximately in antiphase to the Gleissberg cycle. The average sunspot latitude (panel (c)) trend is in phase with the Gleissberg cycle.

This may relate to the slow-timescale trend in reconstructed interplanetary magnetic field and open solar flux which is relatively low during 1890–1910 and relatively high during 1960–1990 (Lockwood & Owens 2011; Owens & Lockwood 2012; see also Svalgaard & Cliver 2010). Those reconstructions (Lockwood & Owens 2011) rely in part on the magnitude of the aa index. Both acv-R27 and $\langle acv(27) \rangle$, as we have shown with the modeling in Section 2, are not sensitive to variation in the overall amplitude of aa ; instead they capture its time structure, from which Figure 13 would then suggest that enhanced 27 day recurrence coincides with reduced open solar flux. The aa index is known to contain a systematic drift in the offsets used in its construction, as discussed in detail in Lockwood et al. (2018b). Since acv-R27 and $\langle acv(27) \rangle$ are not sensitive to variation in the overall amplitude of aa , we would not expect this systematic drift to appear in the slow-timescale trend shown in Figure 13. We have verified that this is indeed the case by repeating our analysis with homogenized aa that has been corrected for these systematic effects by Lockwood et al. (2018c); our results are unchanged.

7. Conclusions

We have used the Hilbert transform of daily SSN to construct a Hale cycle clock which affords the comparison of quantities observed over multiple solar cycles. This clock normalizes successive Hale cycles onto a regular (normalized) 22 yr timebase. The autocovariance of the aa index at a lag of 27 days has been calculated directly and then averaged across these multiple normalized Hale cycles. This provides an index ($\langle acv(27) \rangle$) of solar rotation recurrences with time resolution limited by the autocovariance window length (here, 100 days) rather than being smoothed in time.

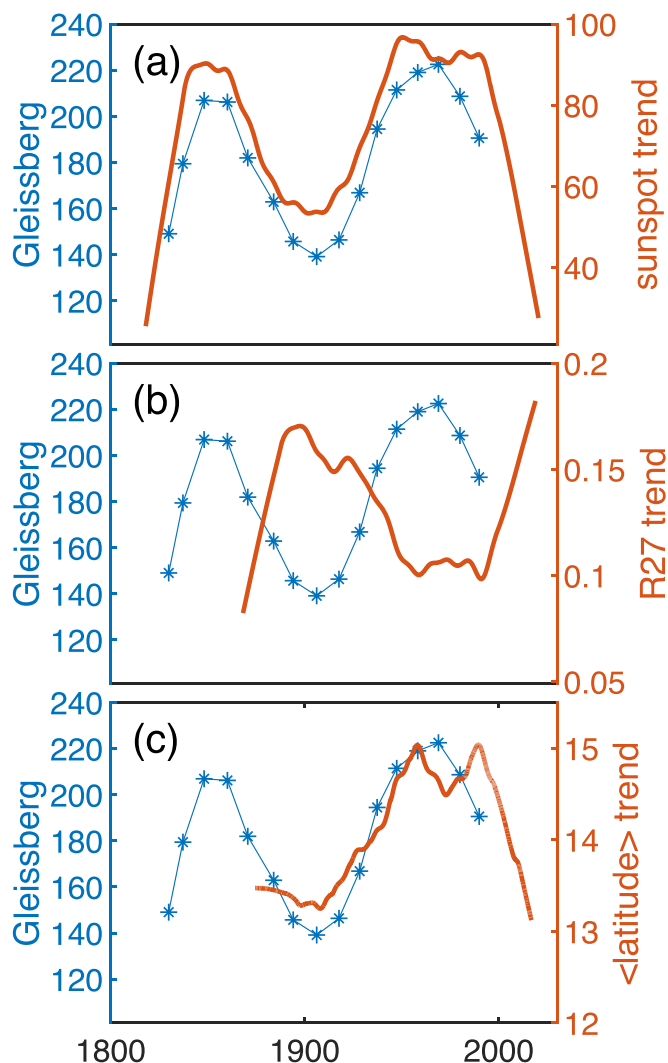


Figure 13. Gleissberg cycle overlaid on slow-timescale (40 yr rlowess local linear regression) trends of (a) SSN, (b) $\langle acv\text{-}R27 \rangle$, and (c) average modulus sunspot latitude.

We obtained the following results.

1. $\langle acv(27) \rangle$ rises sharply at the declining phase on a timescale with an upper bound of a few solar rotations. The declining phase is extended on even cycles—twice as long as odd.
2. $\langle acv(27) \rangle$ drops sharply, on a timescale of a few solar rotations, at solar minimum then slowly decays to its “low” value at the terminators.
3. The occurrence of the solar maxima shows almost no Hale cycle dependence so that the Hilbert transform-derived clock is indeed locked to polarity reversals which occur close to the maxima.
4. The GCR flux rises in step with $\langle acv(27) \rangle$ that is, at the onset of a more ordered solar wind, but then stays high.
5. Average sunspot latitude shows a Hale cycle dependence.
6. The fact that there are Hale cycle dependences on the parameters that we have explored would imply that there are previously unnoticed dependences on the Sun’s magnetic configuration. Pointing to the lower panel of Figure 11 for assistance, for odd numbered cycles (for example cycle 15) the toroidal band configuration is $-/+/-/+$ north to south while it is $+/-/+/-$ north to

south for even numbered cycles. Somehow these configurations produce different manifestations at Earth; this has been explored in the context of cosmic rays, but not for other proxies that we are aware of.

7. Slow-timescale trends, obtained here using a 40 yr (rflowess) local linear regression, give a trend in SSN which tracks the Gleissberg cycle. The slow trend in $\langle acv(27) \rangle$, and indeed in our estimate of the original R27 index, is in antiphase with the SSN trend. These (independent) measures of activity would suggest that the Gleissberg cycle is a very real phenomenon, one rooted in the dynamics controlling the production of the global scale magnetic field, as it is present in both the “dark” and “light” faces of solar activity. We also note that the average sunspot latitude shows a slow trend which tracks the Gleissberg cycle.

Our analysis codifies the Hale cycle-rooted relationship of recurrent solar wind activity and sunspot production. This provides new insight into why the aa index at solar minimum is a reasonable precursor of the upcoming sunspot cycle amplitude.

We thank H. Sargent for illuminating discussions. S.C.C. acknowledges ST/T000252/1 and FA9550-17-1-0054. S.M.C. is supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977. S.C.C., N.W.W., and R.J.L. appreciate the support of the HAO Visitor Program. R.J.L. acknowledges support from NASA’s Living With a Star Program.

All data used in this study are freely available from the following sources: SILSO Royal Observatory of Belgium, Brussels daily total sunspot number version 2.0 from 1818: <http://www.sidc.be/silso/home>. The dates of solar cycle maxima and minima are as determined from the smoothed sunspot number record by SILSO: <http://www.sidc.be/silso/cyclesmm>. Solar radio flux at 10.7 cm (the F10.7 index) from 1947: <https://www.spaceweather.gc.ca/solarflux/sx-en.php>. International Service of Geomagnetic Indices aa index data set from 1868: <http://isgi.unistra.fr/>. University of Oulu/Sodankyla Geophysical Observatory GCR flux from 1964: <http://cosmicrays.oulu.fi/>. Royal Observatory, Greenwich—USAF/NOAA sunspot Data sunspot areas from 1874-2016: <https://solarscience.msfc.nasa.gov/greenwch.shtml>.

ORCID iDs

S. C. Chapman <https://orcid.org/0000-0003-0053-1584>
 S. W. McIntosh <https://orcid.org/0000-0002-7369-1776>
 R. J. Leamon <https://orcid.org/0000-0002-6811-5862>
 N. W. Watkins <https://orcid.org/0000-0003-4484-6588>

References

- Alves, M. V., Echer, E., & Gonzalez, W. D. 2006, *JGRA*, **111**, A07S05
 Barnard, L., McCracken, K. G., Owens, M. J., & Lockwood, M. 2018, *JWSWC*, **8**, A23
 Bartels, J., Heck, N. H., & Johnston, H. F. 1939, *TeMAE*, **44**, 411
 Basu, S., Broomhall, A.-M., Chaplin, W. J., & Elsworth, Y. 2012, *ApJ*, **758**, 6
 Bendat, J. S., & Piersol, A. G. 2010, *Random Data: Analysis and Measurement Procedures* (4th ed.; New York: Wiley)
 Boashash, B. 1992, *Proc. IEEE*, **80**, 520
 Borovsky, J. E., & Denton, M. H. 2006, *JGRA*, **111**, A07S08
 Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. 2016, *Time Series Analysis Forecasting and Control* (5th ed.; New York: Wiley)

- Bubenik, D. M., & Fraser-Smith, A. C. 1977, *JGR*, **82**, 2875
- Chapman, S. C., Horne, R. B., & Watkins, N. W. 2020a, *GeoRL*, **47**, e2019GL086524
- Chapman, S. C., McIntosh, S. W., Leamon, R. J., & Watkins, N. W. 2020b, *GeoRL*, **47**, e2020GL087795
- Chree, C. 1913, *RSPTA*, **212**, 75
- Cliver, E. W. 2014, *SSRv*, **186**, 169
- Cliver, E. W., Boriakoff, V., & Bounar, K. H. 1996, *JGR*, **101**, 27091
- Ferreira, S. E. S., & Potgeiter, M. S. 2004, *ApJ*, **603**, 744
- Feynman, J. 1982, *JGR*, **87**, 6153
- Gabor, D. 1946, *J. Inst. Electr. Eng., Part 3*, 93, 429
- Glæssberg, W. 1967, *SoPh*, **2**, 231
- Hale, G. E., Ellerman, F., Nicholson, S. B., & Joy, A. H. 1919, *ApJ*, **49**, 153
- Hathaway, D. H. 2015, *LRSP*, **12**, 4
- Hewins, I. M., Gibson, S. E., Webb, D. F., et al. 2020, *SoPh*, **295**, 161
- Jokipii, J. R., Levy, E. H., & Hubbard, W. B. 1977, *ApJ*, **213**, 861
- Jokipii, J. R., & Thomas, B. 1981, *ApJ*, **243**, 1115
- Krista, L. D., McIntosh, S. W., & Leamon, R. J. 2018, *AJ*, **155**, 153
- Leamon, R. J., McIntosh, S. W., Chapman, S. C., & Watkins, N. W. 2020a, *SoPh*, **295**, 36
- Leamon, R. J., McIntosh, S. W., Chapman, S. C., et al. 2020b, *SoPh*, submitted (arXiv:2012.15186)
- Legrand, J. P., & Simon, P. A. 1981, *SoPh*, **70**, 173
- Legrand, J. P., & Simon, P. A. 1989, *AnGeo*, **7**, 565
- Legrand, J. P., & Simon, P. A. 1991, *SoPh*, **131**, 187
- Lockwood, M., Chambodut, A., Barnard, L. A., et al. 2018b, *JSWSC*, **8**, A53
- Lockwood, M., Finch, I. D., Chambodut, A., et al. 2018c, *JSWSC*, **8**, A58
- Lockwood, M., & Owens, M. J. 2011, *JGRA*, **117**, A04102
- Maunder, E. W. 1904, *MNRAS*, **64**, 747
- Mayaud, P.-N. 1972, *JGR*, **77**, 6870
- Mayaud, P. N. 1980, *GMS*, **22**, 607
- McIntosh, S. W., & Leamon, R. J. 2014, *ApJL*, **796**, L19
- McIntosh, S. W., Leamon, R. J., Egeland, R., et al. 2019, *SoPh*, **294**, 88
- McIntosh, S. W., Leamon, R. J., Egeland, R., et al. 2021, *SoPh*, submitted (arXiv:2010.06048)
- Owens, M. J., & Lockwood, M. 2012, *JGRA*, **117**, A04102
- Owens, M. J., Lockwood, M., Barnard, L., & Davis, C. J. 2011, *GeoRL*, **38**, L19106
- Pulkkinen, T. 2007, *LRSP*, **4**, 1
- Richardson, I. G. 2018, *LRSP*, **15**, 1
- Ross, E., & Chaplin, W. J. 2019, *SoPh*, **294**, 8
- Sargent, H. H. 1985, *JGR*, **90**, 1425
- Sargent, H. H. 2021, arXiv:2101.02155
- Schwabe, M. 1844, *AN*, **21**, 233
- Simon, P. A., & Legrand, J. P. 1989, *AnGeo*, **7**, 579
- Slutsky, E. 1937, *Econometrica*, **5**, 105
- Smith, E. J. 1990, *JGR*, **95**, 18731
- Smith, E. J., & Thomas, E. J. 1986, *JGR*, **91**, 2933
- Svalgaard, L., & Cliver, E. W. 2010, *JGRA*, **115**, A09111
- Tapping, K. F. 2013, *SpWea*, **11**, 394
- Thomas, S. R., Owens, M. J., & Lockwood, M. 2014, *SoPh*, **289**, 407
- Usoskin, I. G., Mursula, K., & Kangas, J. 2001, *Proc. ICRC (Hamburg)*, **27**, 3842
- Wilson, P. R., Altrrock, R. C., Harvey, K. L., Martin, S. F., & Snodgrass, H. B. 1988, *Natur*, **333**, 748
- Yule, G. U. 1926, *J. R. Stat. Soc.*, **89**, 1