

SMALL-SCALE AND GLOBAL DYNAMOS AND THE AREA AND FLUX DISTRIBUTIONS OF ACTIVE REGIONS, SUNSPOT GROUPS, AND SUNSPOTS: A MULTI-DATABASE STUDY

ANDRÉS MUÑOZ-JARAMILLO^{1,2,3}, RYAN R. SENKPEIL⁴, JOHN C. WINDMUELLER¹, ERNEST C. AMOUZOU¹, DANA W. LONGCOPE¹,
ANDREY G. TLATOV⁵, YURY A. NAGOVITSYN⁶, ALEXEI A. PEVTSOV⁷, GARY A. CHAPMAN⁸, ANGELA M. COOKSON⁸,
ANTHONY R. YEATES⁹, FRASER T. WATSON¹⁰, LAURA A. BALMACEA^{11,12}, EDWARD E. DELUCA¹³, AND PETRUS C. H. MARTENS¹⁴

¹ Department of Physics, Montana State University, Bozeman, MT 59717, USA; munoz@solar.physics.montana.edu

² Space Sciences Laboratory, University of California, Berkeley, CA 94720, USA

³ W.W. Hansen Experimental Physics Laboratory, Stanford University, Stanford, CA 94305, USA

⁴ Department of Physics, Purdue University, West Lafayette, IN 47907, USA

⁵ Kislovodsk Mountain Astronomical Station of the Pulkovo Observatory, Kislovodsk 357700, Russia

⁶ Pulkovo Astronomical Observatory, Russian Academy of Sciences, St. Petersburg 196140, Russia

⁷ National Solar Observatory, Sunspot, NM 88349, USA

⁸ San Fernando Observatory, Department of Physics and Astronomy, California State University Northridge, Northridge, CA 91330, USA

⁹ Department of Mathematical Sciences, Durham University, South Road, Durham DH1 3LE, UK

¹⁰ National Solar Observatory, Tucson, AZ 85719, USA

¹¹ Institute for Astronomical, Terrestrial and Space Sciences (ICATE-CONICET), San Juan, Argentina

¹² National Institute for Space Research (INPE), São José dos Campos, Brazil

¹³ Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, USA

¹⁴ Department of Physics and Astronomy, Georgia State University, Atlanta, GA 30303, USA

Received 2014 September 16; accepted 2014 October 22; published 2015 February 9

ABSTRACT

In this work, we take advantage of 11 different sunspot group, sunspot, and active region databases to characterize the area and flux distributions of photospheric magnetic structures. We find that, when taken separately, different databases are better fitted by different distributions (as has been reported previously in the literature). However, we find that all our databases can be reconciled by the simple application of a proportionality constant, and that, in reality, different databases are sampling different parts of a composite distribution. This composite distribution is made up by linear combination of Weibull and log-normal distributions—where a pure Weibull (log-normal) characterizes the distribution of structures with fluxes below (above) 10^{21} Mx (10^{22} Mx). Additionally, we demonstrate that the Weibull distribution shows the expected linear behavior of a power-law distribution (when extended to smaller fluxes), making our results compatible with the results of Parnell et al. We propose that this is evidence of two separate mechanisms giving rise to visible structures on the photosphere: one directly connected to the global component of the dynamo (and the generation of bipolar active regions), and the other with the small-scale component of the dynamo (and the fragmentation of magnetic structures due to their interaction with turbulent convection).

Key words: Sun: activity – Sun: magnetic fields – Sun: photosphere – sunspots

1. INTRODUCTION

In spite of the great advances in observations and techniques during the last 50 yr, direct observation of the magnetic fields inside the solar convection zone is still out of our reach. This leaves observations of the surface magnetic fields, along with detailed simulations of solar convection, as our only tools for probing what goes on beneath the photosphere. This is no easy task due to the staggering range of length scales and timescales involved, and the fact that the solar magnetic field is daunting in its complexity. Nevertheless, although many structures and events appear to be unique, studying them as part of a larger ensemble allows us to find clues to the underlying mechanisms behind their formation.

A classical example of such behavior is the arrangement of the photospheric magnetic field into patches of magnetic flux spanning many orders of magnitude in lifetime and size, whose presence is a major determinant of the structure and evolution of the solar corona. Furthermore, since the photosphere is the backdrop against which we observe the main signatures of the solar cycle (in the form of the emergence and decay of bipolar magnetic regions, BMRs), understanding how the magnetic field arranges itself in the photosphere also provides clues as to how the solar cycle operates.

Although there are many properties that can be measured in photospheric magnetic patches, one of the most important properties is the amount of flux they contain (which, as will be shown later, is directly related to physical size). This has led to a copious amount of work characterizing the size-distribution of magnetic structures observed on the surface of the Sun, with different studies fitting different analytical distributions to different databases (distributions that will be described in detail in Section 3.1). Tang et al. (1984; analyzing bipolar magnetic regions identified in Mount Wilson Observatory data) and Schrijver et al. (1997; focusing exclusively on the quiet network as measured by the *Solar and Heliospheric Observatory (SOHO)*/Michelson Doppler Imager (MDI) fitted exponential distributions to their data, albeit with different characteristics sizes. Bogdan et al. (1988; analyzing sunspot umbral areas), Baumann & Solanki (2005; analyzing sunspot group data from the Royal Greenwich Observatory, RGO), Zhang et al. (2010; analyzing bipolar magnetic regions detected in *SOHO*/MDI), and Schad & Penn (2010, analyzing sunspot data detected automatically using the NASA/NSO spectromagnetograph) have used log-normal distributions to fit their populations. Harvey & Zwaan (1993; analyzing bipolar magnetic regions identified in Kitt Peak Vacuum Telescope data, KPVT) fitted a third order polynomial to the logarithms of frequency and size of their observations. Parnell

(2002; analyzing ephemeral regions detected automatically on *SOHO*/MDI) found that a Weibull distribution fit the data better than a power law. Zharkov et al. (2005; analyzing sunspots identified automatically using *SOHO*/MDI), Meunier (2003; analyzing automatically detected features on *SOHO*/MDI), and Parnell et al. (2009 data; using the automatic detection of magnetic features on *SOHO*/MDI and *Hinode*/Solar Optical Telescope (SOT)) fitted a power law to their data. Jiang et al. (2011; analyzing sunspot group data from the RGO) fitted a power law to the small sunspot group end of the distribution and a log-normal distribution to the larger end. Finally, Kuklin (1980; analyzing sunspot group data from the RGO), and Nagovitsyn et al. (2012; analyzing sunspot area data taken by the Kislovodsk Mountain Astronomical Station) have used two separate normal distributions to characterize the logarithm of sunspot areas.

One characteristic of studies of the size-distribution of magnetic structures is the ad hoc selection of models to fit the data. Generally, the studies mentioned above include no analysis of the goodness of fit of the chosen distribution (with the exception of the work of Parnell 2002; Parnell et al. 2009), and only 1 work out of 11 (Parnell 2002) uses an objective quantitative criterion to discriminate between two different distributions (the rest include no explanation as to why a particular model was chosen). Furthermore, to the extent of our knowledge, no consistent effort has been made to understand why different studies reach different conclusions (considering that all of them are studying related databases).

In this work, we perform a long-overdue, quantitative, and comparative study of the area and flux distribution of magnetic structures using 11 different sunspot group, sunspot, and BMR databases (described in detail in Section 2). Our first objective is to identify which of the different distributions, used as potential candidates in the literature, is the most adequate to characterize the data. These distributions, as well as the methods used to fit them to the data, and the methods used for model discrimination are described in Section 3. Fitting our distribution candidates to each database (see Section 4), we find that different databases are better fitted by different databases. For this reason, in Section 5, we probe the relationship between flux and area to evaluate if different data types are associated with different distributions. Instead, in Section 6, we find evidence suggesting that different databases are sampling different sections of a universal distribution, and that all can be reconciled by a single proportionality constant. In Section 7, we demonstrate that our data are better fitted by a composite distribution. In Section 8, we discuss the implications of our results, and finish with a summary in Section 9.

2. DATA SELECTION

2.1. Sunspot Group Databases

Our first sunspot group database was compiled and published as the Greenwich Photo-heliographic Results by the RGO. The measurements include the heliographic positions and areas of sunspot groups observed from 1874 to 1976 by a small network of observatories: Cape of Good Hope, Kodaikanal, and Mauritius. In 1976, the program of daily solar observations was transferred to the Debrecen Observatory of the Hungarian Academy of Sciences. The RGO data, covering nine solar cycles, provide the longest and most complete record of sunspot group areas. We extract from this database a single area and position for each sunspot group. We assign to the group the

single largest reported area in all days of observation. The result is a set of 30,026 groups.

Our second sunspot group database has been compiled by the US Air Force, beginning after the RGO program ended operation in 1976, from a global network of ground-based solar observatories known as SOON (the Solar Observing Optical Network) with the aim of providing real time data in order to continuously monitor the Sun for any kind of activity that may affect defense systems. At present, SOON telescopes are providing data at the Holloman Air Force Base (New Mexico), Learmonth (Australia), and San Vito (Italy), but earlier data sets also include data from Sagamore Hill (Massachusetts), Palehua (Hawaii), and Ramey Air Force Base (Puerto Rico). Measurements carried out at the Mt. Wilson and Boulder observatories are also included in the files available up to 2013 on the NOAA website. As with the RGO set, we extract from this database a single area and position for each sunspot group. We assign to the group the single largest reported area in all days of observation. The result is a set of 6764 groups. Although a correction factor of about 1.4 is often applied to SOON areas in order to combine the RGO and SOON data sets (see a review by Hathaway, 2010, and references therein), in this work we leave the SOON data as it is.

Our third sunspot group set comes from the Pulkovo's catalog of solar activity (PCSA), which was compiled by Mstislav N. Gnevyshev and Boris M. Rubashev (between 1932–1937), and Raisa S. Gnevysheva (between 1938–1991), based on observations taken in a wide array of observatories in the framework of the Sun Service program of the USSR. This database contains 115,925 sunspot group observations taken from 1932 August 1 to 1991 December 31 (covering 8.5 solar cycles, from cycle 15 to cycle 22). Once again, we extract from this database a single area and position for each sunspot group. We assign to the group the single largest reported area in all days of observation. The result is a set of 19,038 groups. PCSA data is available at <http://www.gao.spb.ru/database/csa/>, and is described in detail by Nagovitsyn et al. (2008). Data are shown in Figure 1(b).

Our fourth sunspot group database comes from observations taken by the Kislovodsk Mountain Astronomical Station (KMAS) of the Central Astronomical Observatory at Pulkovo. The KMAS has been in continuous operation since 1948, making it one of the very few institutions performing a wide array of solar surveys through the entirety of the space age. This makes it quite valuable as a connecting set between modern missions and previous surveys. This database contains 108,364 sunspot group observations taken from 1954 February 9 to the present (covering 6.5 solar cycles, from cycle 18 to cycle 24). We extract from this database a single area and position for each sunspot group. We assign to the group the single largest reported area in all days of observation. The result is a set of 19,221 groups. KMAS data is available at http://158.250.29.123:8000/web/Soln_Dann/. Data are shown in Figure 1(c).

Our fifth sunspot group database comes from the semi-automatic detection of sunspots on data taken by the Helioseismic and Magnetic Imager (HMI) on the *Solar Dynamics Observatory* (SDO) (see Schou et al. 2012 for details about SDO/HMI) sunspots performed at the KMAS. The data include the heliographic coordinates of each group, its total area, the area of the largest sunspot, and the total number of sunspots and pores in a group. Prior to 2010, the measurements were made manually. Beginning in 2010, a semi-automatic procedure was implemented, when all measurements are made

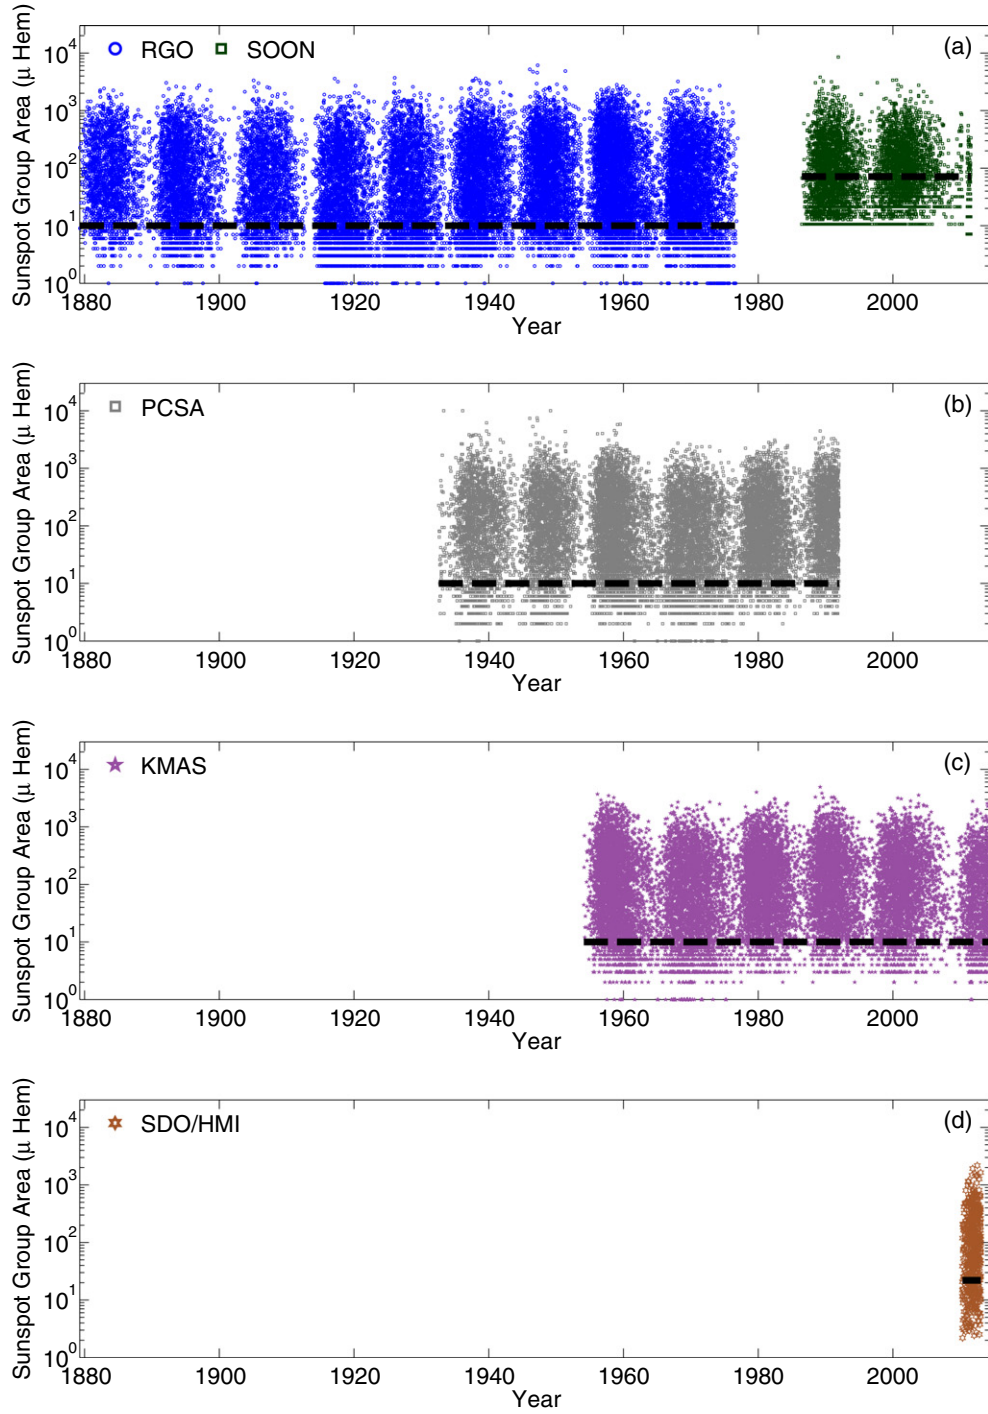


Figure 1. Logarithmic plot of sunspot group area as a function of time. Dashed black horizontal lines indicate the threshold above which data is fitted to the test distributions. This threshold is set an order of magnitude above the smallest structure of each set. (a) RGO/SOON, (b) PCSA, (c) KMAS, and (d) *SDO*/HMI sunspot group area database. Note the marked difference in span between the *SDO*/HMI sunspot group set and the rest. Remarkably, the small interval covered by *SDO* seems to be enough to sample most of the size distribution.

automatically, but the observer is given opportunity to verify the parameters and, if needed, make additional corrections. The detection algorithm identifies outer (quiet-Sun penumbra) and inner (penumbra-umbra) penumbral boundaries using two different methods: intensity threshold (e.g., Watson et al. 2009) and the border (gradient) method (e.g., Zharkova et al. 2005). The algorithm is applied to daily observations from *SDO*/HMI using 1 image per day, resulting in a set containing 18,341 sunspots. To minimize projection issues when measuring magnetic properties, we only use spots within 60 heliographic degrees of

disk center. These sunspots are then collected into groups using NOAA catalog index numbers. We extract from this database a single area and position for each sunspot group. We assign to the group the single largest reported area in all days of observation. The result is a set of 565 groups going between 2010 May 3 and 2014 January 14. More details of the detection algorithm can be found in Tlatov et al. (2014), and more details of its application to HMI data in Tlatov & Pevtsov (2014). This database is available at <http://158.250.29.123:8000/web/sdo/>. Data are shown in Figure 1(d).




Figure 2. Logarithmic plot of sunspot group area and magnetic flux as a function of time. Dashed black horizontal lines indicate the threshold above which data is fitted to the test distributions. This threshold is set an order of magnitude above the smallest structure of each set. (a) Sunspot area measured by SFO. (b) *SOHO*/MDI (light green diamonds) and *SDO*/HMI (dark red squares) sunspot areas detected using the STARA algorithm. (c) KPVT (magenta triangles) and MDI (dark green diamonds) unsigned BMR flux. (d) KPVT/SOLIS synoptic map unsigned BMR flux.

2.2. Sunspot Area Databases

Our first sunspot area database was compiled by A. M. Cookson, G. A. Chapman, & G. de Toma (see de Toma et al. 2013). Spots are detected by applying an automatic detection algorithm to 672.3 nm full-disk 512×512 images (Chapman et al. 1992) taken by the San Fernando Observatory (SFO) of the California State University-Northridge. The resulting database contains 34,697 entries, going from 1986 May 26 to 2013 December 31. One of the best features of these data is the detection of spots based on their photometric contrast. Since this is a physical property of sunspots related to the magnetic

field strength (Norton & Gilman 2004; Schad & Penn 2010), SFO areas are more accurate than areas derived from images that are not calibrated, giving it a high level of precision. More data processing details can be found in Walton et al. (1998). Data are shown in Figure 2(a).

Our second and third sunspot area databases have been compiled by Watson et al. (2011) by applying the sunspot tracking and recognition algorithm (STARA; see Watson et al. 2009) to *SOHO*/MDI (see Scherrer et al. 1995 for details about *SOHO*/MDI) and *SDO*/HMI data. These databases are of particular interest because they involve data from two different instruments, reduced using the exact same algorithm.

The resulting sets go from 1996 July 9 to 2010 October 26 for MDI, and from 2010 May 1 to 2013 July 12 for HMI. They include 16,141 entries for MDI and 9536 for HMI. It is important to note that these sets measure only the umbral area, whereas the SFO set combines umbral and penumbral areas. Data are shown in Figure 2(b).

2.3. Bipolar Magnetic Region Databases

Our first BMR database was assembled by Sheeley et al. (1985), and Wang & Sheeley (1989), using photographic prints of daily full disk magnetograms taken by the 512 channel magnetograph (Livingston et al. 1976) at the KPVT between 1976 August 16 and 1986 March 5 (covering solar cycle 21). Data reduction was performed manually using different techniques to estimate flux (for more details, see Sheeley et al. 1985, and Wang & Sheeley 1989). Special care was taken to count each BMR only once (even across a solar rotation) and measure its properties at the moment of full development. The resulting database contains 3046 BMRs. Data are shown in Figure 2(c).

Our second BMR database was assembled manually using a semi-automatic detection algorithm applied to *SOHO*/MDI magnetograms between 1996 November 12 and 2011 April 11 (covering solar cycle 23 and part of 24). One MDI full-disk line-of-sight magnetogram per day was inspected visually in search of new BMRs. When a new emergence was found, the region was followed until it was deemed to have fully developed, and then its two polarities were enclosed by a single hand-drawn (mouse-drawn) curve. The MDI pixels within each enclosing curve were used to compute the net flux and flux-weighted centroid of each polarity. The line-of-sight field strength was assumed to arise from a purely radial field and was therefore divided by the cosine of the angle from the disk center. The pixel areas were divided by the same factor to account for foreshortening. Pixels with field strength below 75 G were not included in these totals. Active regions that emerged on the backside of the Sun were characterized once they crossed the east limb (at a longitude of about 75°E). The result is a database containing 977 BMRs. Data are shown in Figure 2(c).

Our third BMR database was assembled using a semi-automatic detection algorithm applied to synoptic magnetogram data assembled using the KPVT and SOLIS from 1996 June 28 and 2014 January 15. Regions are defined as continuous pixel groups with radial field $|B_r|$ greater than a threshold of 50 G. A comparison is made with previous synoptic maps in order to ensure each region is counted only once. If a region is found to be too complex, unipolar, or in direct violation of Hale's law, it is flagged for human supervision. The result is a database containing 2412 BMRs. Data are shown in Figure 2(d). More details on the detection algorithm can be found in Yeates et al. (2007).

2.4. Truncation and Separation of Data

One of the observations often made when studying size distributions is the fact that the number of structures near the lower detection threshold is always undercounted. This is unavoidable when the cadence of the detection is similar in duration to the lifetime of small structures, even under perfect observational and detection conditions. Another problem affecting the detection of small structures arises from an unavoidable rounding error to which instruments are subject, resulting in an artificial binning of small objects into a small set of values. Figures 1 and 2, showing our data in a logarithmic scale, are quite illustrative of these problems.

In the case of human observers, the undercount of small structures is aggravated by changes in the quality of the observing conditions (for ground-based observations) and excessive complexity in the observed phenomenon (particularly evident in magnetograms taken during the active phases of the cycle). The time-dependent sensitivity of the MDI detection, where the observer is able to detect a larger number of small features during solar minimum than during solar maximum (see Figure 2(c)), is a clear example of this problem.

Another example of observational bias can be seen on the KPVT BMR set (see Figure 2(c)), where a slight declining trend is visible in terms of the flux of the smallest structures, which is caused by a combination of factors: first, early observations (1975–1977) had a larger number of noisy pixels (J. Harvey 2014, private communication), which would make the detection of smaller objects more difficult. Second, there was a selection effect since this BMR database was tailored for studying the large-scale magnetic field of the Sun (which is determined mainly by larger objects), which made small objects of secondary importance (N. Sheeley 2014, private communication). Finally, there is an unavoidable learning curve that allows the observer to be more effective at detecting smaller objects (N. Sheeley 2014, private communication). Altogether, they lead to an uneven detection of small structures across the different reduction campaigns.

In the case of automatic detection, other issues become evident. The first one is the difference in detection thresholds that can be used in *SOHO*/MDI and *SDO*/HMI (Figure 2(b)). This results in databases spanning different orders of magnitude which cannot be combined, and thus need to be analyzed separately. Another visible issue is the six-month modulation of areas in the smallest pores/sunspots of the SFO database (Figure 2(a)), caused by the yearly change in distance between the Sun and the Earth (compounded with the relatively large pixel size of the instrument). Furthermore, there seems to be a modulation in the discretization of the smallest values, with measurements more prone to collapsing into discrete values during certain parts of the year.

Our intention in highlighting these problems is not to reduce the legitimacy of our data sets for solar cycle studies, but rather to underline a fact that is very often overlooked: if one considers that small structures are also the most numerous, then it follows that these issues can skew the process of model distribution fitting quite significantly. Following a suggestion by C. DeForest (2014, private communication), we impose a truncation limit for all databases located one order of magnitude above the minimum size of detection, and only use data above this limit in our distribution fits and analysis. The location of these thresholds, shown in Figures 1 and 2 as dark horizontal lines, successfully isolates problematic data from the rest of each set.

3. MATHEMATICAL METHODS

Considering that different size distributions arise due to different physical processes, identifying which distribution fits the data best can be used to probe the mechanisms behind the creation of magnetic structures observed in the Sun. However, as mentioned above, we want to do this using an objective quantitative criterion and not the ad hoc model selection that is customary. In this section, we describe in detail the model distributions we will fit to the data, our method for fitting a given distribution to a data set, and the quantitative criteria that we use for model selection.

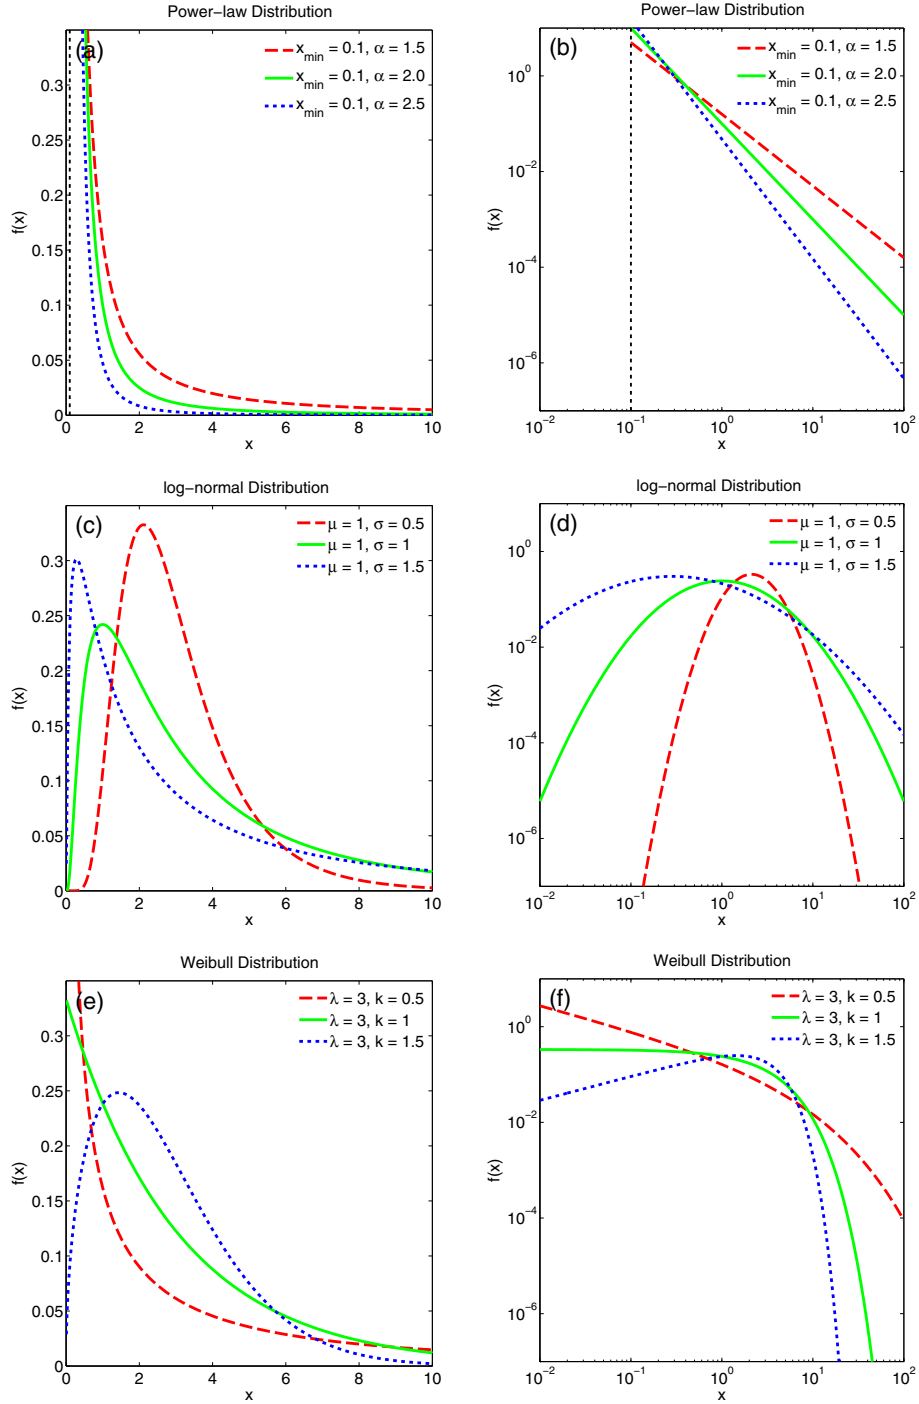


Figure 3. Power-law (top row; Equation (1)), log-normal (middle row; Equation (2)), and Weibull (bottom row; Equation (3)) distributions. All are plotted using both linear (left column) and logarithmic scales (right column). In all cases, three different parameter sets are shown. In the case of the power-law distribution, the minimum structure size is illustrated with a vertical dashed line.

3.1. Power-law, Log-normal, Exponential, and Weibull Distributions

The probability distributions that we fit to the data are the power-law distribution (see Figures 3(a) and (b)):

$$f(x; \alpha) = \frac{\alpha - 1}{x_{\min}} \left(\frac{x}{x_{\min}} \right)^{\alpha}, \quad (1)$$

where α is the power-law index and x_{\min} is the lower limit covered by the distribution; the log-normal distribution (see

Figures 3(c) and (d)):

$$f(x; \mu, \sigma) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x - \mu)^2}{2\sigma^2}}, \quad (2)$$

where μ and σ are the mean and standard deviation of the variable's natural logarithm; the Weibull distribution (see Figures 3(e) and (f)):

$$f(x; k, \lambda) = \frac{k}{\lambda} \left(\frac{x}{\lambda} \right)^{k-1} e^{-(x/\lambda)^k}, \quad (3)$$

where $k > 0$ and $\lambda > 0$ are its shape and scale parameters; and the exponential distribution:

$$f(x; \lambda) = \frac{1}{\lambda} e^{-(x/\lambda)}, \quad (4)$$

which can be seen as a Weibull distribution with a shape parameter $k = 1$ (included in Figures 3(e) and (f)).

Although a detailed explanation of the generative processes that lead to these distributions is beyond the scope of this paper, characterizing the size and flux distribution of magnetic structures gives us insight into the internal processes that give shape and structure to the solar magnetic field. All these distributions have been used to characterize a wide variety of processes, ranging from city growth to failure rate in communications, passing through income distribution and the sizes of living organisms (to name a few examples). Considering that the evolution of the solar magnetic field is primarily driven by its interaction with turbulent convection, in our brief review we focus on generative processes that lead to growth or fragmentation (in this case of the magnetic field).

In the case of the power-law and log-normal distributions, one of the possible generative processes is the fragmentation and aggregation of magnetic structures due to multiplicative iterations. In this kind of process, growth or shrinkage is governed by a random proportionality variable. In other words, the size of a structure in a subsequent step is always proportional to its size, and the proportionality constant is randomly distributed. What actually makes this process lead to either power-law or log-normal distributions is the fact that power-law distributions have a minimum size x_{\min} , beyond which structures cannot shrink (illustrated in Figures 3(a) and (b) as a vertical dotted line); whereas structures governed by a log-normal can become arbitrarily small (with the additional restriction that the proportionality constant is normally distributed). For more information on the generative processes behind power laws, log-normals, and their relationship, we recommend a very interesting review by Mitzenmacher (2003).

In terms of the Weibull and exponential distributions, one of the possible generative processes is sequential fragmentation, where a large structure is broken into smaller and smaller pieces through the application of mechanical forces. In fact, the Weibull distribution was first used to characterize the size-distribution of particles generated by grinding, milling, and crushing operations (Rosin & Rammler 1933), and the fracture of materials under repetitive stress (Weibull 1939). In the solar case, one can speculate that the repetitive fragmentation occurs on the magnetic field, and the mechanical agent is turbulent convection. In this case, as demonstrated by Brown & Wohletz (1995), the shape parameter k can be interpreted as a measure of the fractal dimension of the fragmentation process. It is important to mention that exponential and Weibull distributions have also been demonstrated to arise from generative processes involving emergence, coalescence, fragmentation, and cancellation of flux, depending on the assumptions made on the rates governing these different physical mechanisms. Please refer to the work of Schrijver et al. (1997) and the work of Parnell (2002) for the derivation of generative processes leading to exponential and Weibull distributions, respectively.

Figure 3 is quite illustrative of the intrinsic differences between these distributions. For example, processes leading to a log-normal distribution are characterized by very small and very large structures that are significantly less probable than mid-sized structures (arising from the fact that both growth and

fragmentation are involved). This is not the case for the power law, for which the hard limit imposed on fragmentation leads to an imbalance that inflates small structures compared with the larger ones. In contrast, in the case of Weibull distributions with shape parameter $0 < k \leq 1$ (which contains the exponential distribution as well), structures can become arbitrarily small and their relative abundance increases significantly with a decrease in size; however, large structures are less frequent when compared to the power-law distribution. This is related to the fact that one of the main generative processes of the Weibull distribution involves repetitive fragmentation.

Although a first principle derivation of each of these distributions is beyond the scope of this paper, it is that a detailed characterization of the size and flux distribution of magnetic structures can provide invaluable insight into the processes governing the evolution of the solar magnetic field.

3.2. Distribution Fitting

In order to fit distributions to the data, we use maximum likelihood estimates (MLE). This method is far superior to fitting functional forms to histograms because it is not sensitive to the details of data binning. The idea is to find the set of parameters that maximizes the likelihood of a statistical model M given the observed data $D = \{D_1, D_2, \dots, D_n\}$ by maximizing the likelihood (L) function:

$$L(M) \propto \text{pr}(D|M) = \prod_{i=1}^n \text{pr}(D_i|M). \quad (5)$$

This process of maximization is typically performed by first taking the logarithm of both sides of Equation (5), and maximizing the resulting log-likelihood (lk) function:

$$\text{lk}(M) = \sum_{i=1}^n \log(\text{pr}(D_i|M)). \quad (6)$$

More information about MLE can be found in most modern statistic books (for example, in Hoel 1984).

Since we are working with truncated sets, we use truncated distributions on our fits—building them from each probability distribution function (PDF) and cumulative distribution function (CDF) in the following manner:

$$\text{PDF}_{\text{trunc}}(x) = \frac{\text{PDF}(x)}{1 - \text{CDF}(x_{\text{trunc}})}, \quad (7)$$

and

$$\text{CDF}_{\text{trunc}}(x) = \frac{\text{CDF}(x) - \text{CDF}(x_{\text{trunc}})}{1 - \text{CDF}(x_{\text{trunc}})}, \quad (8)$$

where x_{trunc} denotes the limit value below which data is not used in the fit (see Section 2.4).

3.3. Model Selection

Ultimately, we want to compare the relative performance of different models to fit the data. To quantify relative performance, we use two separate criteria, the first one is the Kolmogorov–Smirnov (K-S) statistic, which corresponds to the biggest difference between the observed and model CDFs:

$$\text{KS} = \max |\text{CDF}(x) - \text{CDF}_{\text{emp}}(x)| \quad (9)$$

for $x_{\text{trunc}} \leq x \leq \infty$.

The second one is Akaike's information criterion (AIC; Akaike 1983). The AIC is a powerful tool for discriminating between different non-nested models by making an estimate of the expected, relative distance between the fitted model and the unknown true mechanism that generated the observed data. The AIC for a model M_j is defined as:

$$\text{AIC}_j = -2 \text{lk}(M_j) - 2n_j, \quad (10)$$

where $\text{lk}(M_j)$ is the log-likelihood of model M_j (as defined above) and n_j the number of parameters of model j . The model with the minimum AIC is chosen as the best. In a sense, by minimizing AIC, one is looking for the model with the largest log-likelihood. However, log-likelihood alone is not sufficient to discriminate between models because it is biased as an estimation of the model selection target. This bias was found by Akaike (1983) to be approximately equal to each model's number of parameters (n), and thus the presence of the second term in Equation (10). Together, log-likelihood and n are used to strike a balance between bias and variance (or the trade-off between underfitting and overfitting). It is very important to highlight that the significance of AIC is strongly dependent on an appropriate choice of models. Applying AIC to a set of very poor models will always select one estimated to be the best (even though that model may still be poor in an absolute sense).

The relative nature of the AIC is better represented by calculating the relative AIC differences:

$$\Delta_j^{\text{AIC}} = \text{AIC}_j - \min(\text{AIC}). \quad (11)$$

This in turn can be used to estimate the likelihood of a model given the data:

$$\mathcal{L}(M_j|D) \propto \exp\left(-\frac{\Delta_j^{\text{AIC}}}{2}\right), \quad (12)$$

and use it to calculate the Akaike weights:

$$Aw_j = \frac{\exp\left(-\frac{\Delta_j^{\text{AIC}}}{2}\right)}{\sum_{k=1}^K \exp\left(-\frac{\Delta_k^{\text{AIC}}}{2}\right)}, \quad (13)$$

which are a measure of the probability that the model M_j is the best model given the data. For more information about AIC, we recommend the excellent book by Burnham & Anderson (2002).

4. SINGLE DISTRIBUTION FIT RESULTS

The results of fitting log-normal, power law, exponential, and Weibull distributions to our data are tabulated in Table 1 and shown in Figure 4 for sunspot group area, Figures 5(a), (c), and (e) for sunspot area, Figures 5(b), (d), and (f) for BMR unsigned flux. Due to the large amount of data in almost every set, AIC (see columns 6 and 7 in every section of Table 1) unambiguously selects one of the models with likelihoods above 0.99 when compared with the other models (with relative AIC differences of the order of thousands). In every case, the smallest K-S statistic also coincides with the most likely model defined by AIC.

In agreement with previous results, no single distribution fits all data sets. However, even though in every case there is a clear indication of what distribution yields the best fit, very few of the fits pass the K-S test (in which the null hypothesis assumes

that the observed data is drawn by the fitted distribution). This is illustrated in column 5 of Table 1, which, for each set and distribution, shows the estimated probability that the observed data (or a more extreme set) was drawn randomly from each given distribution. The only fits yielding significant probabilities (4/11) are the Weibull distribution fit to HMI sunspot group area ($P = 0.75$), the log-normal distribution fit to SFO sunspot area ($P = 0.56$), the log-normal distribution fit to manual MDI BMR flux data ($P = 0.78$), and the log-normal distribution fit to KPVT/SOLIS BMR flux data ($P = 0.83$). This suggests that, even though in each case we can find a best fit, neither of these models is capturing the real distribution giving rise to these populations.

We find that no database is better fitted by either power-law or exponential distributions. Instead, databases are better fitted by either Weibull or log-normal distributions. Interestingly, there seems to be a preferred distribution fit depending on the kind of data used. On the one hand, for all sunspot group area sets (RGO, SOON, PCSA, KMAS, and HMI), as well as the two STARA umbral area sets (MDI and HMI), the best fit is the Weibull distribution. On the other hand, the SFO sunspot area set, as well as the BMR flux sets (KPVT, MDI, and KPVT/SOLIS), are better fitted by log-normal distributions. In the next sections, we explore why our databases are either fitted by Weibull, or log-normal distributions, as well as the possible implications.

5. RELATIONSHIP BETWEEN FLUX AND AREA

As mentioned above, one of the intriguing results of fitting our databases to the different distributions is the separation of our databases into those better fitted by a Weibull distribution and those better fitted by a log-normal distribution—a separation that does not appear to occur randomly, but which clearly differentiates between data types (i.e., sunspot group area data are better fitted by a Weibull distribution, whereas BMR flux data are better fitted by a log-normal distribution with sunspot area data falling between). An obvious question arises: Can flux be compared with area? Or, in other words, is the fact that flux and area data are better fitted by different distributions evidence that they cannot be compared?

In a recent paper, Tlatov & Pevtsov (2014) reported an approximately linear relationship between sunspot area and sunspot magnetic flux. Figure 6 shows a reproduction of this relationship for sunspot groups automatically detected on SDO/HMI (Figure 6(a)), as well as the relationships we obtain using sunspot umbras detected using the STARA algorithm on MDI (Figure 6(b)) and HMI (Figure 6(c)). Fitting this relationship using the least squares method to a power law of the form:

$$f(x; a, b) = ax^b, \quad (14)$$

we find $a = (1.95 \pm 0.14)10^{19}$ and $b = 0.98 \pm 0.01$, with a coefficient of determination of $R^2 = 0.98$ for HMI groups, $a = (6.21 \pm 0.11)10^{19}$ and $b = 0.97 \pm 0.01$, with a coefficient of determination of $R^2 = 0.94$ for MDI umbras, and $a = (5.20 \pm 0.03)10^{19}$ and $b = 1.08 \pm 0.01$, with a coefficient of determination of $R^2 = 0.99$ for HMI umbras. It is to be expected that the proportionality constant between flux and area for sunspot groups (that include penumbrae) is less than if one considers only umbrae. We find significantly more scatter for MDI than we do for HMI data. Several factors may be playing a role, and one is the difference in spatial resolution: MDI pixels are 16 times larger than those of HMI, which would make the areas measured with MDI appear larger

Table 1
Fitting Parameters and Model Selection Quantities for the Sunspot Group Area Distributions

Sunspot Group Area RGO						
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	3.94	1.67	0.049	<0.001	806.4	<0.001
Power Law	α	X_{\min}^*	0.132	<0.001	7,862	<0.001
	1.47	1.00				
Exponential		λ^*	0.211	<0.001	9,742	<0.001
		187.61				
Weibull	k	λ^*	0.045	<0.001	0	>0.999
	0.49	68.30				
Sunspot Group Area SOON						
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	4.60	1.18	0.027	0.065	10.66	0.005
Power Law	α	X_{\min}^*	0.084	<0.001	215.32	<0.001
	2.09	10.00				
Exponential		λ^*	0.119	<0.001	326.34	<0.001
		252.57				
Weibull	k	λ^*	0.024	0.131	0	0.995
	0.48	43.56				
Sunspot Group Area KMAS						
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	4.40	1.55	0.050	<0.001	687	<0.001
Power Law	α	X_{\min}^*	0.164	<0.001	7,763	<0.001
	1.42	1.00				
Exponential		λ^*	0.179	<0.001	4,948	<0.001
		230.6				
Weibull	k	λ^*	0.031	<0.001	0	>0.999
	0.56	115.89				
Sunspot Group Area PCSA						
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	4.29	1.61	0.048	<0.001	554	<0.001
Power Law	α	X_{\min}^*	0.153	<0.001	6,712	<0.001
	1.43	1.00				
Exponential		λ^*	0.202	<0.001	6,369	<0.001
		234.95				
Weibull	k	λ^*	0.035	<0.001	0	>0.999
	0.52	99.60				
Sunspot Group Area SDO/HMI						
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	4.61	1.25	0.047	0.284	13	0.001
Power Law	α	X_{\min}^*	0.181	<0.001	175	<0.001
	1.56	2.20				
Exponential		λ^*	0.112	<0.001	53	<0.001
		204.06				
Weibull	k	λ^*	0.032	0.754	0	0.999
	0.66	123.17				
Sunspot Umbral Area MDI						
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	2.59	1.03	0.016	0.030	95	<0.001
Power Law	α	X_{\min}^*	0.117	<0.001	1,673	<0.001
	1.89	0.68				
Exponential		λ^*	0.082	<0.001	477	<0.001
		21.86				
Weibull	k	λ^*	0.012	0.197	0	>0.999
	0.66	11.55				
Log-Normal	μ	σ	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
	1.02	1.40	0.034	<0.001	143	<0.001

Table 1
(Continued)

Sunspot Umbral Area HMI						
Power Law	α 1.60	X_{\min}^* 0.09	0.126	<0.001	1542	<0.001
Exponential		λ^* 7.40	0.157	<0.001	1252	<0.001
Weibull	k 0.54	λ^* 2.88	0.022	0.004	0	>0.999
Sunspot Area SFO						
Log-Normal	μ 4.41	σ 1.08	K-S St. 0.006	K-S Pr. 0.559	Δ_j^{AIC} 0	Aw >0.999
Power Law	α 1.89	X_{\min}^* 4.40	0.126	<0.001	3,260	<0.001
Exponential		λ^* 149.81	0.102	<0.001	2,407	<0.001
Weibull	k 0.56	λ^* 51.94	0.020	<0.001	103	<0.001
BMR Flux KPVT						
Log-Normal	μ 49.93	σ 0.99	K-S St. 0.105	K-S Pr. <0.001	Δ_j^{AIC} 0	Aw >0.999
Power Law	α 1.96	X_{\min}^\dagger 0.20	0.209	<0.001	411	<0.001
Exponential		λ^\dagger 7.14	0.127	<0.001	88	<0.001
Weibull	k 0.88	λ^\dagger 6.67	0.131	<0.001	23	<0.001
BMR Flux MDI						
Log-Normal	μ 51.20	σ 0.77	K-S St. 0.024	K-S Pr. 0.785	Δ_j^{AIC} 0	Aw 0.983
Power Law	α 2.07	X_{\min}^\dagger 0.88	0.139	<0.001	185	<0.001
Exponential		λ^\dagger 21.00	0.074	<0.001	9	0.011
Weibull	k 1.12	λ^\dagger 22.46	0.073	0.001	10	0.005
BMR Flux KPVT/SOLIS						
Log-Normal	μ 50.05	σ 0.75	K-S St. 0.014	K-S Pr. 0.834	Δ_j^{AIC} 0	Aw >0.999
Power Law	α 1.95	X_{\min}^\dagger 0.22	0.168	<0.001	666.31	<0.001
Exponential		λ^\dagger 6.45	0.065	<0.001	25.68	<0.001
Weibull	k 1.13	λ^\dagger 6.91	0.061	<0.001	24.34	<0.001

Notes. Fitting parameters and model selection quantities for the sunspot group area, sunspot area, and BMR unsigned flux distributions. Quantities accompanied by a * are in units of μ Hem, quantities accompanied by a \dagger are in units of 10^{21} Mx, and other quantities are dimensionless. K-S St. denotes the K-S distance described in Equation (9). K-S Pr. is the probability of observing each database (or a more extreme set) given a fitted distribution function. Δ_j^{AIC} is the relative AIC difference described by Equation (11). Aw is the Akaike weight described by Equation (13). Best fit is highlighted in bold letters.

than they really are due to partial filling factors. Another difference may be the fact that MDI measures magnetic field averaged throughout the pixel, blending positive and negative flux together. Finally, MDI magnetic fields are corrected line-of-sight field measurements, whereas the HMI fields come from Milne-Eddington inversions. Nevertheless, in all cases, results are consistent with a proportional relationship between area and flux, suggesting that they can be considered in a joint analysis

and that the underlying reasons leading to different distributional fits go beyond the nature of the measured quantity.

6. RECONCILIATION OF DATA SETS AND EVIDENCE IN FAVOR OF A COMPOSITE DISTRIBUTION

Once we move beyond the different quantities that have been measured, there is another striking difference between




Figure 4. Distribution fits to sunspot group area: (a) RGO, (b) SOON, (c) KMAS, (d) PCSA, and (e) *SDO/HMI*. Figures show a logarithmic histogram and fits to the distributions described in Section 3.1. Histograms include all data in each set, but only data shown in a dark shade are included in the fits.

the databases that are better fitted by Weibull and log-normal distributions: the range covered by each set. As a general rule, those databases that cover the greatest number of decades (sunspot group areas) are better fitted by a Weibull distribution, whereas those that cover the smallest number of decades (BMR flux) are better fitted by a log-normal (see Figures 4 through 6). That in and of itself would not be remarkable, were it not for the different nature of structures that make it into each of those sets. On the one hand, BMR flux databases are extremely selective, focusing on the largest objects that appear in the photosphere and further limiting the selection of magnetic structures to those that are bipolar and in close flux balance. On the other hand, sunspot group databases include both the structures that are part of the BMR databases, as well as their fragmentation into individual sunspots and pores. This is significant because while BMR sets are only sampling the larger end of the true solar distribution, fits to sunspot group databases are being driven by smaller

structures which are significantly more numerous (effectively oversampling the smaller end of the true solar distribution).

In order to quantify and visualize this trend, we take advantage of AIC as an estimate of the expected relative distance between the fitted model and the unknown true mechanism that generated the observed data (see Section 3.3). For each database, we calculate a normalized AIC relative difference between the Weibull and log-normal AICs:

$$\Delta_{W-LN}^{AIC} = -\frac{(AIC_{Wb} - AIC_{LN})}{N}, \quad (15)$$

where AIC_{Wb} and AIC_{LN} are calculated using Equation (10), and N is the number of points in the data set. This quantity is positive (negative) when the distribution is better fitted by a Weibull (log-normal) distribution, and its magnitude is indicative of how much better the fit is. The $1/N$ factor is a rough normalization factor used to standardize all databases (whose

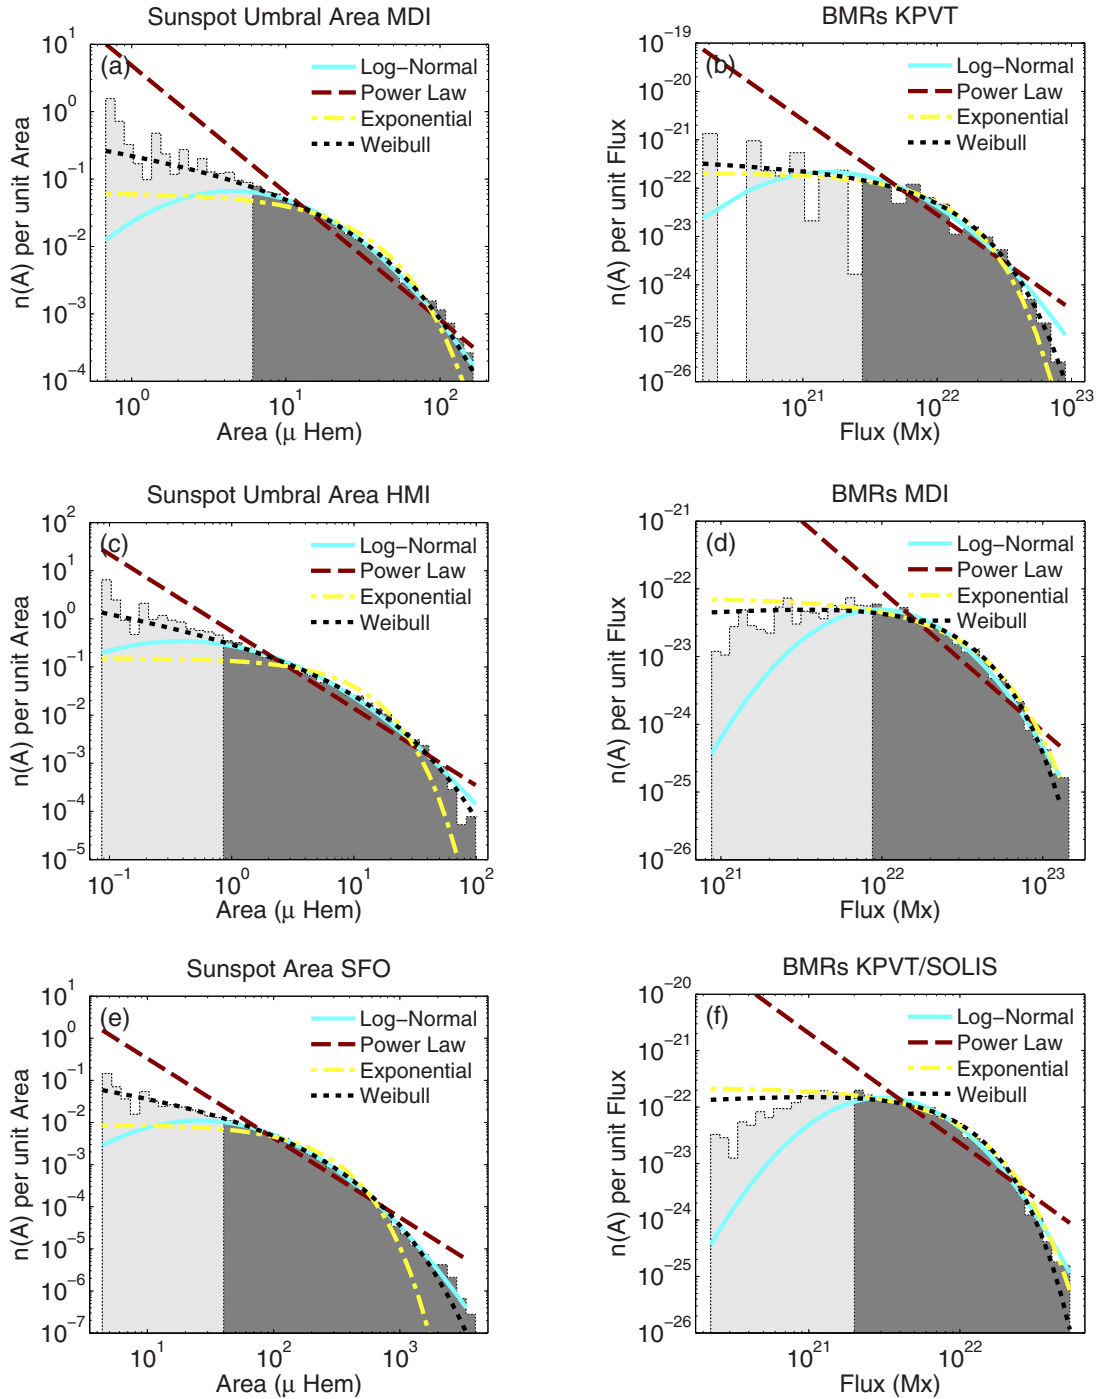


Figure 5. Distribution fits to sunspot area: (a) SOHO/MDI, (c) SDO/HMI, and (e) SFO; and distribution fits to BMR flux: (b) KPVT, (d) SOHO/MDI, and (f) KPVT/SOLIS. Figures show a logarithmic histogram and fits to the distributions described in Section 3.1. Histograms include all data in each set, but only data shown in a dark shade are included in the fits.

size differ significantly), so that they can be compared with each other. We find a very clear relationship between this quantity and the logarithmic data range (ratio between the smallest and largest object on a database; see Figure 7). We propose that different sets are actually sampling different sections of a universal composite distribution.

6.1. Database Cross-calibration

In order to look for evidence of a composite distribution, we use the empirical distribution of the RGO database as reference, and make comparisons between sections of this

reference distribution and the empirical distribution of the rest of our databases. This comparison can be performed all across our databases due to the proportional relationship existing between magnetic flux and area (shown in Figure 6).

Our procedure, which we perform separately for each of our databases, consists of the following steps.

1. Choose a proportionality constant out of a range of possible values.
2. Multiply all sizes (or fluxes) in the database by this proportionality constant (effectively shifting the empirical distribution left or right in logarithmic scale).




Figure 6. Log-log scatter plot of sunspot group area vs. sunspot group unsigned magnetic flux as measured by HMI (a), and umbral sunspot area vs. umbral unsigned magnetic flux as measured by MDI (b) and HMI (c). The dashed lines correspond to a power-law fits of the form ax^b . For HMI sunspot groups, we find a proportionality constant $a = (1.95 \pm 0.14)10^{19}$ and an exponent $b = 0.98 \pm 0.01$. For MDI umbrae, we find a proportionality constant $a = (6.21 \pm 0.11)10^{19}$ and an exponent $b = 0.97 \pm 0.01$. For HMI umbrae, we find a proportionality constant $a = (5.20 \pm 0.03)10^{19}$ and an exponent $b = 1.08 \pm 0.01$. The coefficients of determination of the fits are $R^2 = 0.98$, $R^2 = 0.94$, and $R^2 = 0.99$, respectively.




Figure 7. Logarithmic data range vs. normalized AIC relative difference for the Weibull and log-normal distributions. Logarithmic range is the ratio between the largest and smallest object in each database (not counting data below the accuracy threshold; see Section 2.4). The normalized AIC relative difference quantifies how much better a database is fitted by either the Weibull or log-normal distributions—a positive (negative) value indicates that the database is better fitted by the Weibull (log-normal) distribution and is denoted using solid (open) markers in the plot. Different marker shapes and colors are used to denote different types of data: Sunspot group area (blue circles), sunspot area (red squares), and BMR flux (magenta triangles).

3. Evaluate if the resulting empirical distribution overlaps with the reference RGO distribution.
4. Find the root mean square error (RMSE) between the overlaps.
5. After trying all possible proportionality values in a set, identify which one minimizes RMSE.

Besides the proportionality constant (which shifts the empirical distribution left or right), we also add a normalization constant that accounts for the fact that each set contains a different number of datapoints (which shifts the empirical distribution up or down).

The results of this experiment, shown in Figure 8, support our hypothesis that different sets are actually sampling different sections of a universal composite distribution, and demonstrate

Table 2
Calibration Constants between our Sunspot Group Area and MDI BMR Unsigned Flux Databases

Sunspot Group Area Databases		
	From RGO SG Area	To RGO SG Area
SOON	1.11	0.90
KMAS	1.07	0.93
PCSA	1.22	0.82
HMI	1.10	0.91
Sunspot Area Databases		
	From RGO SG Area	To RGO SG Area
MDI	0.06	15.43
HMI	0.03	30.57
SFO	0.71	1.41
BMR Flux Databases		
	From RGO SG Area (Mx/ μ Hem)	To RGO SG Area (μ Hem/Mx)
KPVT	2.05×10^{19}	4.88×10^{-20}
MDI	4.68×10^{19}	2.14×10^{-20}
KPVT/SOLIS	1.60×10^{19}	6.22×10^{-20}

Notes. Sunspot group area constants (top four rows) are in units of Mx/ μ Hem. The BMR unsigned flux constants (bottom row) is dimensionless.

that a simple proportionality constant is sufficient to connect them. Additionally, as can be observed in Figures 8(f), (g) and (h), the distribution of sunspot sizes is contained within the distribution of sunspot group sizes. This is consistent with a picture in which the generation process that leads to the formation of BMRs and sunspot groups is the same process that leads to the fragmentation of these structures to form individual sunspots and smaller magnetic elements.

Based on the excellent agreement between reference and test distributions found for every database, we argue that this method can be useful for cross-calibrating data sets (even if there is no time overlap between them). In fact, as can be seen in Figure 8(e), four years' worth of HMI sunspot groups (numbering only 565 in contrast with the 30,026 contained in the RGO database) seems to be enough to sample most of the distribution.

Although the focus of this work is not to perform calibrations (nor thoroughly reconcile different data sets), as an interesting exercise, in Table 2 we show the conversion factors needed to transform all our databases to and from RGO sunspot group area. It is reassuring to find that the calibration factors




Figure 8. Overplot of the empirical distribution of our databases against the reference empirical distribution of RGO sunspot group data (a). Each color indicates a different type of data. Blue shows the empirical distributions of sunspot group area: (b) SOON, (c) KMAS, (d) PCAS, and (e) HMI groups. Red shows the empirical distributions of sunspot areas: (f) MDI, (g) HMI, and (h) SFO. Green shows the empirical distributions of unsigned BMR flux: (i) KPVT, (j) MDI, and (k) KPVT/SOLIS. The location of each empirical distribution, within the reference distribution of RGO, is obtained by using the proportionality constants shown in Table 2. This converts all sets to units of sunspot group area (i.e., μHem). Histograms include all data in each set, but only the sections shown in a dark shade are included in the cross calibration.

obtained between sunspot group area and BMR flux databases (by fitting the empirical distributions) is similar to the one obtained using direct measurements of area and flux (obtained by fitting direct measurements using a power-law; see Figure 6 and Section 5). This supports the usefulness of this method for database calibration.

7. FIT TO A COMPOSITE DISTRIBUTION

Although there is an understandable hesitancy to increase the number of fitting parameters for fear of over-fitting the

data, our results strongly suggest that fitting a combination of distributions is the correct approach. This has been performed in the past by Kuklin (1980) and Nagovitsyn et al. (2012), who fitted two log-normal distributions to their data. In particular, Nagovitsyn et al. (2012) showed that a histogram of sunspot group area using logarithmic binning shows two distinct peaks, one at $17 \mu\text{Hem}$ and the other at $174 \mu\text{Hem}$ (and that bin count in such histogram can be fitted using normal distributions).

The top row of Figure 9 shows the RGO, KMAS, and PCAS data cast in a histogram using logarithmic binning showing a double-peaked structure. When translated into empirical




Figure 9. (Top row) Histogram using logarithmic binning of RGO (a), KMAS (b), and PCSA (c) sunspot group area. (Middle row) Empirical PDF of RGO (d), KMAS (e), and PCSA (f) sunspot group area. The arrows point at the change in the curvature of the PDF. (Bottom row) RGO (a), KMAS (b), and PCSA (c) empirical PDFs, overplotted with a fit using a linear combination of Weibull (dashed blue line) and log-normal distributions (dotted yellow line). The composite fit is shown as a solid dark red line. In all cases, the improvement in the fit goes beyond what is expected statistically from the increased number of parameters.

distributions (shown in the middle row of Figure 9), the presence of these peaks turns into a weak depression that deviates from a pure Weibull or log-normal distribution.

Due to the fact that the leftmost part of the peak around $17 \mu\text{Hem}$ is populated by data near the detection threshold (which, as demonstrated in Section 2.4, is troublesome and generally under-represented), it is possible to see the trend as increasing for smaller objects. Because of this, and based on the results of Section 4, we propose a change in the approach of Nagovitsyn et al. (2012), which is to substitute the log-normal distribution used to fit the peak around $17 \mu\text{Hem}$ for a Weibull distribution. The combination of a Weibull and log-normal distributions becomes:

$$f(x; k, \lambda, \mu, \sigma, c) = \frac{ck}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k} + \frac{(1-c)}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x - \mu)^2}{2\sigma^2}}, \quad (16)$$

where $k > 0$ and $\lambda > 0$ are the shape and scale parameters of the Weibull distribution, μ and σ are the mean and standard deviation characterizing the log-normal, and $0 \leq c \leq 1$ is the proportionality constant that blends these distributions together.

The results of this fit are shown (tabulated) in the bottom row of Figure 9 (Table 3) and represent a significant improvement over the single function fitting. This is not only visible qualitatively in terms of a tight fit of the distribution's ankle and knee, but also qualitatively in terms of a reduced K-S statistic (see Equation (9)) for the three databases, shown in column 6 and Table 3.

Perhaps more importantly is how, for all three sets, the recalculated relative AIC differences (see Section 3.3) find the composite function to be the most likely model out of all the models presented in this paper (with likelihoods above 0.99). This is very important because AIC factors a penalization for the addition of parameters. This means that, out of all the fitting

Table 3
Fitting Parameters of the Composite Distribution to RGO, KMAS,
and PCSA Sunspot Group Data

Composite Fit to RGO sunspot group data								
Weibull		Log-Normal		c	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
k	λ^*	μ	σ	0.57	0.024	<0.001	0	>0.999
0.57	16.21	5.62	0.85					
Composite Fit to KMAS sunspot group data								
Weibull		Log-Normal		c	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
k	λ^*	μ	σ	0.64	0.022	<0.001	0	>0.999
0.61	40.34	5.93	0.79					
Composite Fit to PCSA sunspot group data								
Weibull		Log-Normal		c	K-S St.	K-S Pr.	Δ_j^{AIC}	Aw
k	λ^*	μ	σ	0.67	0.020	<0.001	0	>0.999
0.55	34.03	5.96	0.82					

Notes. Quantities accompanied by a * are in units of μHem , and other quantities are dimensionless. K-S St. denotes the K-S distance described in Equation (9). K-S Pr. is the probability of observing each database (or a more extreme set) given a fitted distribution function. Δ_j^{AIC} is the relative AIC difference described by Equation (11). Aw is the Akaike weight described by Equation (13). Both Δ_j^{AIC} and Aw are re-calculated including all the models fitted to RGO, KMAS, and PCSA shown in Table 1.

models presented in this paper, the composite fit is the best and not just because it has more parameters. This should not come as a surprise if one considers that we are dealing with databases that have significantly more entries than fitting parameters.

Unfortunately, the composite distribution function still does not pass a K-S test and has a very low probability of surfacing as a random draw. This indicates that, despite being the best model presented in this paper, there are still subtleties in the data that need to be captured and understood. In a preliminary analysis, we have found that this is caused in part by changes in the statistical properties of magnetic structures with the progression of the cycle. A detailed exploration of this time dependence will be performed in a future article.

8. IMPLICATIONS OF A COMPOSITE FLUX-AREA DISTRIBUTION

Taking advantage of both the proportional relationship that we find between all our databases (see Section 6) and the fitting of RGO data to a composite distribution (see Section 7), we can return to the question as to why some of them are better fitted by Weibull or log-normal distributions. Figure 10 shows what happens if we overplot the fitted composite distribution, as well as its Weibull and log-normal components on the calibrated databases. It can be observed that there is very good agreement between the single distribution fits found to be the best for each database (see Section 4), and whether or not their range includes a significant portion of the Weibull component.

Focusing on the overplots of the distributions of BMR flux and the composite fit to RGO data (Figures 10(i), (j), and (k)), we find a remarkable coincidence between the location and shape of BMR data and the location and shape of the log-normal component of the composite distribution. Although this can only be treated as circumstantial evidence, it suggests that the log-normal component of the flux-area distribution is inherently related with the appearance of BMRs in the photosphere (i.e., clearly bipolar structures whose poles appear simultaneously), whereas the mechanisms giving rise to smaller magnetic structures are different (and characterized by a Weibull distribution). Invoking the generative processes associated with

the log-normal and Weibull distributions (see Section 3.1), our results suggest that large-scale flux-tubes are formed in a process that allows for both growth and fragmentation (i.e., there is a preferred set of scales that are more likely to occur than much larger or smaller objects), whereas only the repetitive fragmentation inherent to the Weibull distribution can explain the significant amount of smaller magnetic structures observed in the empirical distribution (coupled with a reduced frequency for large structures). We propose this as evidence in favor of the formation of BMR flux-tubes in the stable layer at the bottom of the convection zone, whereas the distribution of small-scale magnetic fields arises from the interaction of these structures, as well as their fragments, with convection throughout the convection zone (and at the photosphere).

By taking advantage of our characterization of the composite distribution, we can identify the length scales at which magnetic structures originate from either of the proposed generation mechanisms. Figure 11 shows the relative contribution of the Weibull and log-normal distribution to the composite (using both sunspot group area and BMR unsigned flux). We find that the transition from one regime to the other takes place during a full order of magnitude between roughly 10^{21} and 10^{22} Mx (30 and $300 \mu\text{Hem}$). It is to be expected that some of the objects in this flux range are either small emergent BMRs or the result of the initial fragmentation of mid to large BMRs (i.e., the largest sunspots). Although this may be coincidental, this transitional range is roughly the same as the transitional range found by Tlatov & Pevtsov (2014) that separates sunspots into two distinct populations (small and large) with different average properties. Perhaps part of the reason behind such separation resides in the fact that sunspots belonging to each of these categories arise from different generation mechanisms.

8.1. Consistency with the Results of Parnell et al. (2009)

The final issue that we address is the apparent discrepancy between our results (in which a power-law distribution is clearly the worst model that can be used to characterize any of our databases) and the results of Parnell et al. (2009) (in which, applying six different detection algorithms on MDI/HR, MDI/FD, and SOT/NFI magnetograms, they find a power-law distribution covering more than five orders of magnitude in flux).

Before addressing this issue, it is important to clarify that Parnell et al. (2009) are characterizing a slightly different quantity than the one we are characterizing in this work. The difference arises from the fact that Parnell et al. (2009) used features detected in instantaneous magnetic snapshots, whereas our databases encompass all features observed within a period of several (to more than a hundred) years. The difference is subtle but very important because both approaches fold in time-dependent information of the size distribution. On the one hand, the time span of all our databases is orders of magnitude above the longest lived structures inside them, which means that we are folding cycle dependencies into our fits. On the other hand, the time span of the databases of Parnell et al. is orders of magnitude below the longest lived structures inside them, which means that they are folding the comparative lifetime of different structures into their fit.

In spite of these differences, it is interesting to explore the behavior of our composite distribution as it extends into the length scales observed by Parnell et al. Looking at Figure 11(b), it is clear that a Weibull distribution shows the expected behavior of a power law for small scales, since it displays a nearly log-linear behavior for more than five orders of magnitude.

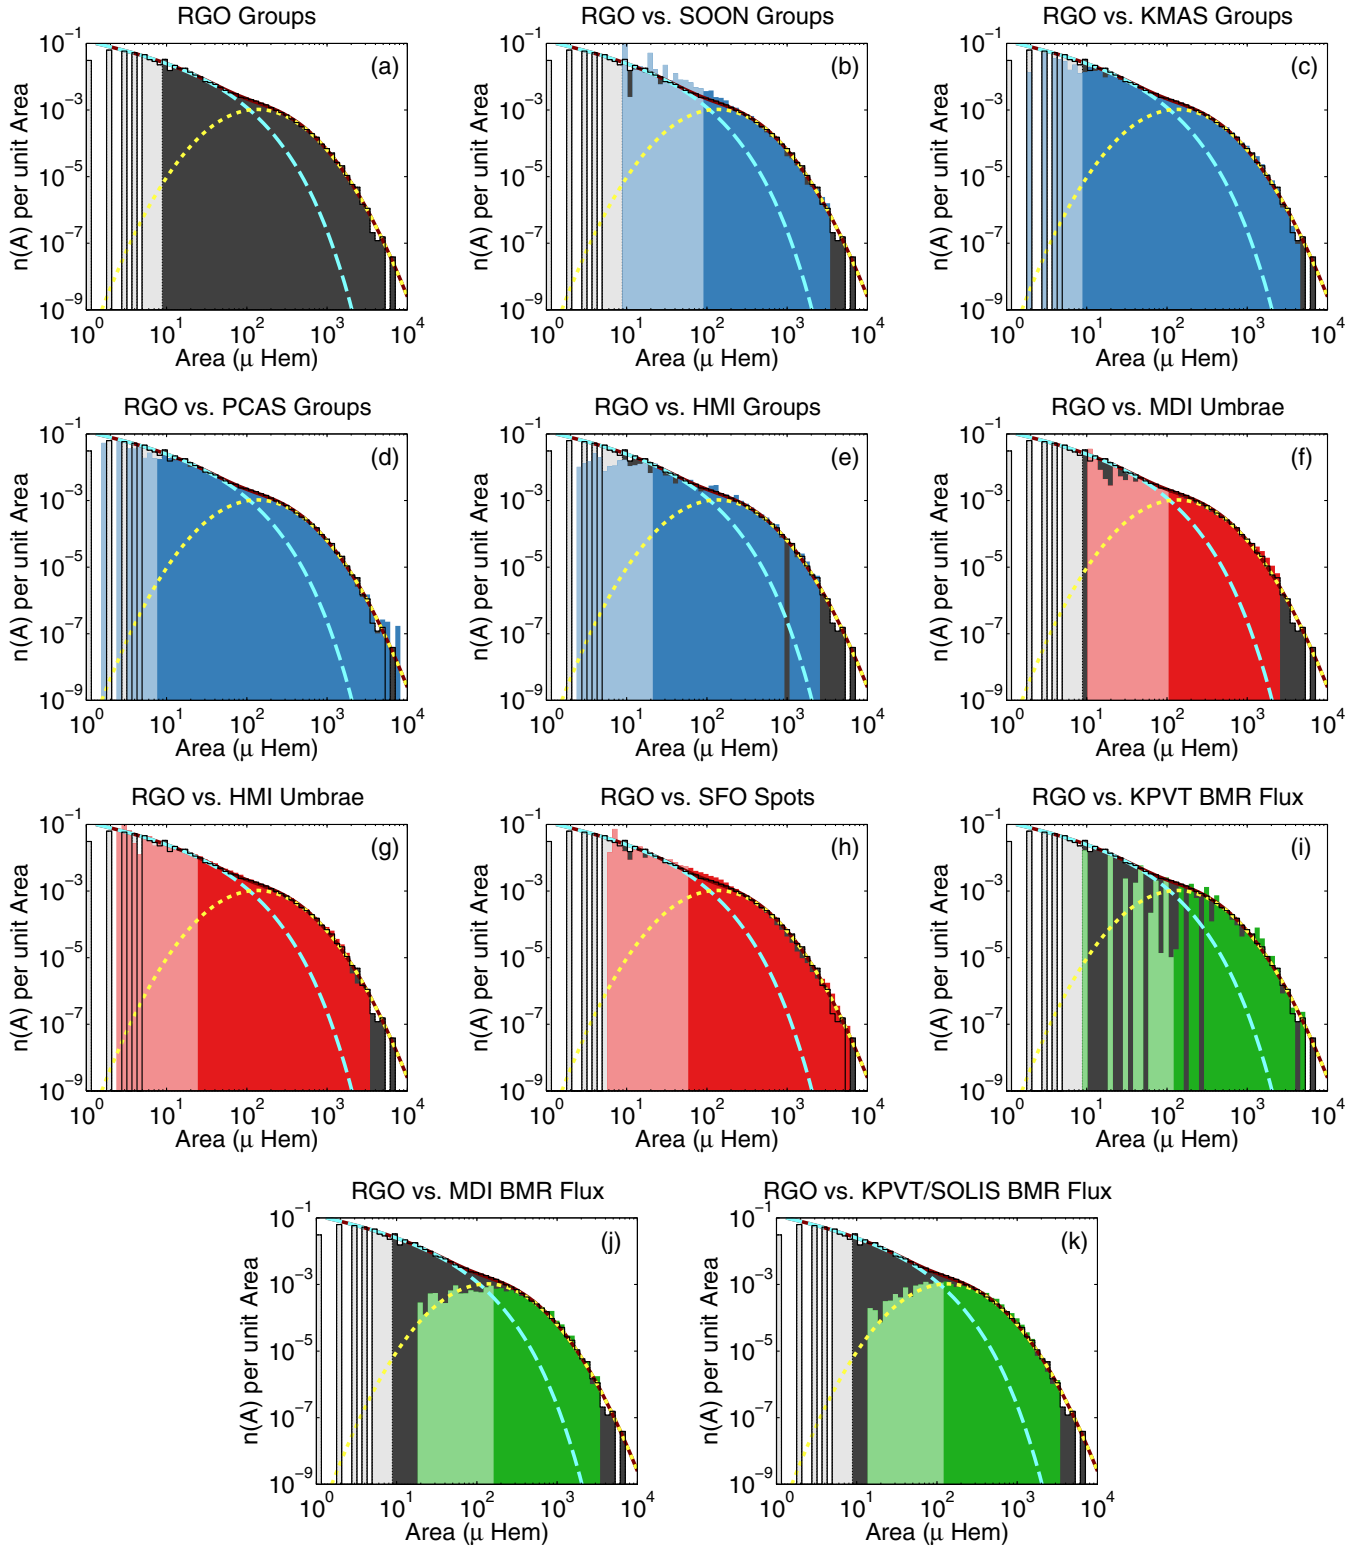


Figure 10. Overplot of the RGO fit to a composite distribution over all shifted databases. The composite fit is shown as a solid dark red line. The Weibull (dashed blue line) and log-normal distributions (dotted yellow line) that form part of the composite are shown as well. The same composite distribution is overplotted on all figures, and it is the composite fit to RGO data shown in Figure 9(g). For additional information on colors and background empirical distributions, see the caption of Figure 8.

We propose that perhaps what Parnell et al. (2009) are observing is indeed a Weibull distribution. This agrees with the results of Parnell (2002), who, analyzing ephemeral regions detected automatically on *SOHO*/MDI between 10^{18-20} Mx, performed a quantitative comparison between Weibull and power-law distributions and found the Weibull distribution to be superior to

the power law. It is clear that with this analysis we are pushing the limits of our databases, barely scratching at the distribution of small magnetic structures. However, the analysis of Parnell et al. (2009) also involves very limited timescales. Only the careful analysis of long-term magnetic data will be able to truly characterize these distributions.

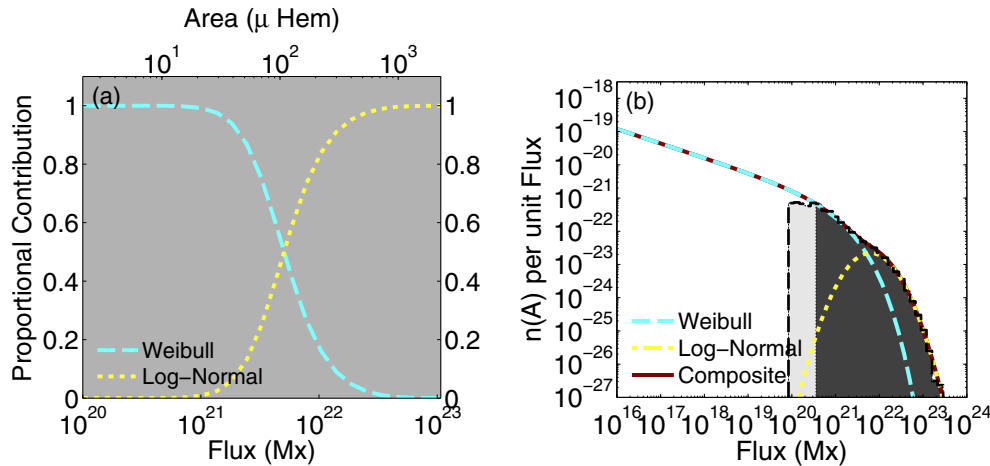


Figure 11. (a) Relative contribution of the Weibull and log-normal components to the composite distribution. (b) Extrapolation of the composite distribution toward smaller domains showing behavior consistent with the log-linearity of a power law.

9. SUMMARY AND CONCLUDING REMARKS

The focus of this work has been the characterization of the flux-area distribution of sunspot groups, sunspots, and bipolar magnetic regions. This is largely motivated by a wide array of different competing results in the literature, and a general lack of a quantitative comparison between candidate distributions. For this purpose, we use 11 different databases: 5 sunspot group area databases (Royal Greenwich Observatory, the USAF’s Solar Observing Optical Network, Pulkovo’s catalog of solar activity, Kislovodsk Mountain Astronomical Station, and *SDO/HMI*), 3 sunspot area databases (San Fernando Observatory, *SOHO/MDI*, and *SDO/HMI*), and 3 unsigned BMR flux databases (The 512 Channel magnetograph at the Kitt Peak Vacuum Telescope, *SOHO/MDI*, and synoptic maps assembled by the Kitt Peak Vacuum Telescope and SOLIS).

Using the Kolmogorov–Smirnov statistic and Akaike’s information criterion we test which—power-law, log-normal, exponential, or Weibull distributions—is the best distribution that fits each of our databases. We find that for six of our databases (RGO groups, SOON groups, KMAS groups, PCSA groups, HMI groups, MDI spots, and HMI spots) the best fit is the Weibull distribution, and for the remaining four (SFO spots, KPVT BMR flux, MDI BMR flux, and KPVT/SOLIS BMR flux) the best fit is a log-normal. In every single case, we find the power law to be the worst distribution for describing the data.

Motivated by the work of Kuklin (1980) and Nagovitsyn et al. (2012), we test the possibility that the flux-area distribution of magnetic structures is better described by a composite distribution combining Weibull and log-normal distributions. Furthermore, we test whether the reason why some databases are better fitted by Weibull or log-normals is that different databases sample different sections of this composite distribution. Our results demonstrate that all our databases can be made compatible by the simple application of a proportionality constant, and that all our databases are indeed sampling different parts of a composite flux-area distribution. We find that those better fitted by log-normals span only the largest structures, whereas those better fitted by Weibull distributions contain a significant amount of small structures. We find that the transition between the Weibull and log-normal components of the composite distribution occurs for fluxes (areas) between 10^{21} and 10^{22} Mx (30 and 300 μ Hem). For structures with fluxes (areas) below 10^{21} Mx (30 μ Hem) the composite distribution is essentially a Weibull and for structures

with fluxes (areas) above 10^{22} Mx (300 μ Hem) the composite distribution is essentially a log-normal.

We find a remarkable coincidence between the log-normal part of the composite distribution and the shape and location of the distributions of BMR unsigned flux. At the same time, only a Weibull distribution (arising from processes of repetitive fragmentation) can explain both the significant amount of small structures present in the data, and the relative decrease in large ones. We propose that this is evidence of two separate mechanisms giving rise to visible structures on the photosphere: one directly connected to the global component of the dynamo (and the generation of bipolar active regions), and the other with the small-scale component of the dynamo (and the fragmentation of magnetic structures due to their interaction with turbulent convection).

Although our results (in which the power law yields the worst fits) seem to be at odds with the results of Parnell et al. (2009), who reported a power-law distribution covering more than five orders of magnitude in flux, we demonstrate how a Weibull distribution shows the expected linear behavior of a power-law distribution for small scales. We propose that the flux-area distribution for small-scale structures is not a power law, but a Weibull distribution, as proposed originally by Parnell (2002). Ultimately, only a multi-scale analysis of the flux-area distribution involving all length scales of interest, as well as solar cycle timescales, can truly settle this issue.

Our discovery that a proportionality constant is sufficient to harmonize the size-flux distribution of different databases creates a useful framework within which multiple databases can be cross-calibrated. Furthermore, the existence of a proportional relationship between flux and area (see Tlatov & Pevtsov 2014) makes this method useful for cross-calibration between magnetic and optical contrast data. Additionally, the applicability of this method seems to be independent of the observational particularities of each database (automatic versus human, ground-based versus space-based, etc.), and valid irrespective of whether or not the databases overlap in time. We believe that this method will help promote a better consolidation of long-term databases spanning all our instruments and decades of observation, thereby enhancing the usefulness of historic data in a modern context.

Although our results are suggestive, and we have made an effort to interpret them from a physical point of view, a solid theoretic framework is still necessary to take maximum advantage of the characteristics of the observed flux-area distributions. Of

particular interest would be to perform studies of the size distribution of magnetic structures in MHD simulations of turbulent convection. Not only this will provide an additional constraint to those simulations, but, together, simulations and observations will help us further our understanding of flux-emergence and transport throughout the convection zone.

We thank our anonymous referee for a very detailed and conscientious report, which significantly improved the quality of this paper. Additionally, we thank Giuliana de Toma, Neil Sheeley, Jack Harvey, Craig DeForest, William Dean Pesnell, Steve Cranmer, and María Navas Moreno for useful discussions and suggestions. We are very grateful to Neil R. Sheeley Jr. for sharing his KPVT BMR database with us. This research was supported by the NASA Living With a Star Jack Eddy Postdoctoral Fellowship Program, administered by the UCAR Visiting Scientist Programs, contract SP02H1701R from Lockheed-Martin to the Smithsonian Astrophysical Observatory, and the CfA Solar Physics REU program, NSF grant number AGS-1263241. Andrés Muñoz-Jaramillo is very grateful to George Fisher and Stuart Bale for their support at the University of California - Berkeley, and Phil Scherrer for his support at Stanford University. The National Solar Observatory (NSO) is operated by the Association of Universities for Research in Astronomy, AURA Inc under cooperative agreement with the National Science Foundation (NSF).

REFERENCES

- Akaike, H. 1983, *Bull. Int. Stat. Inst.*, 50, 277
- Baumann, I., & Solanki, S. K. 2005, *A&A*, 443, 1061
- Bogdan, T. J., Gilman, P. A., Lerche, I., & Howard, R. 1988, *ApJ*, 327, 451
- Brown, W. K., & Wohletz, K. H. 1995, *JAP*, 78, 2758
- Burnham, K., & Anderson, D. 2002, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (New York: Springer)
- Chapman, G. A., Herzog, A. D., Lawrence, J. K., et al. 1992, *JGR*, 97, 8211
- de Toma, G., Chapman, G. A., Cookson, A. M., & Preminger, D. 2013, *ApJL*, 771, L22
- Harvey, K. L., & Zwaan, C. 1993, *SoPh*, 148, 85
- Hathaway, D. H. 2010, *LRSP*, 7, 1
- Hoel, P. G. 1984, *Introduction to Mathematical Statistics*, Vol. 1 (New York: Wiley)
- Jiang, J., Cameron, R. H., Schmitt, D., & Schüssler, M. 2011, *A&A*, 528, A82
- Kuklin, G. V. 1980, *BAICz*, 31, 224
- Livingston, W. C., Harvey, J., Pierce, A. K., et al. 1976, *ApOpt*, 15, 33
- Meunier, N. 2003, *A&A*, 405, 1107
- Mitzenmacher, M. 2003, *Internet Math.*, 1, 226
- Nagovitsyn, Y. A., Miletsky, E. V., Ivanov, V. G., & Guseva, S. A. 2008, *CosRe*, 46, 283
- Nagovitsyn, Y. A., Pevtsov, A. A., & Livingston, W. C. 2012, *ApJL*, 758, L20
- Norton, A. A., & Gilman, P. A. 2004, *ApJ*, 603, 348
- Parnell, C. E. 2002, *MNRAS*, 335, 389
- Parnell, C. E., DeForest, C. E., Hagenaar, H. J., et al. 2009, *ApJ*, 698, 75
- Rosin, P., & Rammler, E. 1933, *J. Inst. Fuel*, 7, 29
- Schad, T. A., & Penn, M. J. 2010, *SoPh*, 262, 19
- Scherrer, P. H., Bogart, R. S., Bush, R. I., et al. 1995, *SoPh*, 162, 129
- Schou, J., Scherrer, P. H., Bush, R. I., et al. 2012, *SoPh*, 275, 229
- Schrijver, C. J., Title, A. M., van Ballegoijen, A. A., Hagenaar, H. J., & Shine, R. A. 1997, *ApJ*, 487, 424
- Sheeley, N. R., Jr., DeVore, C. R., & Boris, J. P. 1985, *SoPh*, 98, 219
- Tang, F., Howard, R., & Adkins, J. M. 1984, *SoPh*, 91, 75
- Tlatov, A. G., & Pevtsov, A. A. 2014, *SoPh*, 289, 1143
- Tlatov, A. G., Vasil'eva, V. V., Makarova, V. V., & Otkidych, P. A. 2014, *SoPh*, 289, 1403
- Walton, S. R., Chapman, G. A., Cookson, A. M., Dobias, J. J., & Preminger, D. G. 1998, *SoPh*, 179, 31
- Wang, Y., & Sheeley, N. R., Jr. 1989, *SoPh*, 124, 81
- Watson, F., Fletcher, L., Dalla, S., & Marshall, S. 2009, *SoPh*, 260, 5
- Watson, F. T., Fletcher, L., & Marshall, S. 2011, *A&A*, 533, A14
- Weibull, W. 1939, *A Statistical Theory of the Strength of Materials, Ingeniörsvetenskapsakademiens handlingar (Generalstabens litografiska anstalts förlag)*
- Yeates, A. R., Mackay, D. H., & van Ballegoijen, A. A. 2007, *SoPh*, 245, 87
- Zhang, J., Wang, Y., & Liu, Y. 2010, *ApJ*, 723, 1006
- Zharkov, S., Zharkova, V. V., & Ipson, S. S. 2005, *SoPh*, 228, 377
- Zharkova, V. V., Aboudarham, J., Zharkov, S., et al. 2005, *SoPh*, 228, 361