Supporting Online Material

Determining pulse shape parameters

The pairs of pulse profiles in Fig. 3 of the journal paper were inspected and fitted with two or three gaussian components, as required to provide satisfactory descriptions of the profiles. The same components were fitted to each observed profile, and a synthetic profile produced, from which the value of the chosen shape parameter used in Fig. 4 of the journal paper was determined. W_{10}, W_{50} and W_{75} are the widths measured at 10%, 50% and 75% of the pulse peak. W_{eq} is the equivalent width, being the area under the pulse divided by the peak amplitude, and A_{pc}/A_{mp} is the ratio of the peak amplitudes of the precursor and main pulse components. Note that this procedure has the virtue of applying a quasi-optimum filter to the data in order to minimise the effects of high-frequency noise on the values of the parameters.

Determining times-of-arrival

We note that the times-of-arrival used to derive the timing residuals of Fig. 1 in the journal paper were usually obtained by matching the observed pulse profiles with templates derived from average profiles obtained over large sections of the data. Hence, for the objects in Fig. 3 of the journal paper, the templates have intermediate shapes, and are usually not perfectly matched to the observed profile, giving rise to possible systematic offsets in the timing residuals. The maximum magnitudes of the offsets have been estimated by comparing the times-of-arrival obtained by fitting the two profiles in Fig. 3 with the corresponding template. For the six objects, the differences in the offsets were respectively 0.70, 0.05, 0.25, 0.17, 0.04, and 0.14 ms. Since these are much smaller than the variations seen in Fig. 1 of the journal paper, we conclude that the profile switching and use of a single template has little impact upon the structures seen in the timing residuals. In practice, for PSR B1828−11 the procedure described in (*S1*) was used and therefore the TOAs are not prone to such systematic errors.

Observational limitations

Within the relatively small number of the pulsars in our sample which have high signal-to-noise ratio profiles, pulse-shape variations are observed which are correlated with spin-down rate. The results show that multiple pulse-profile and associated spin-down states that switch on timescales of weeks to years is a common phenomenon seen in many pulsars. We believe that it is possible that all pulsars which display timing noise may show the same correlated pulse-shape behaviour, although it is so far unobserved in most pulsars, because of a combination of modest profile changes, poor signal-to-noise ratio and often relatively-short available data spans. The profile changes may be small, for instance, if changes in the pulsar emission beam happen to be small in that part of the beam which crosses the line-of-sight to the Earth.

Simulations of timing residuals (Fig. S1)

We have calculated some simulations of a two-state spin-down model for pulsar timing noise in which only two parameters determine the form of the simulations, namely the ratio R of time spent in high and low spin-down states, and the rms fractional dither D in the switching period. In detail, simulated values of $\dot{\nu}$ were determined for a regular time sampling. $\dot{\nu}$ switched between two modes $\dot{\nu}_1$ and $\dot{\nu}_2$, spending a time t_1 in the first mode and t_2 in the second mode (typically t_1 and t_2 are a few hundred days). The ratio of t_1 and t_2 equals R. For simulations involving dithering the timescales, t_1 and t_2 are slightly modified by adding a Gaussian random deviate with an rms of Dt_1 and Dt_2 respectively. The resulting $\dot{\nu}$ values are numerically integrated twice to produce pulse phase, followed by appropriate sampling and a quadratic polynomial removed to form the resulting simulated timing residuals. Note how dither in the switching period can give rise to low-frequency structure in the residuals (Fig. S1c).

References

[S1] I. H. Stairs, A. G. Lyne, S. Shemar, *Nature* **406**, 484 (2000).

[S2] G. Hobbs, A. G. Lyne, M. Kramer, *MNRAS* **402**, 1027 (2010).

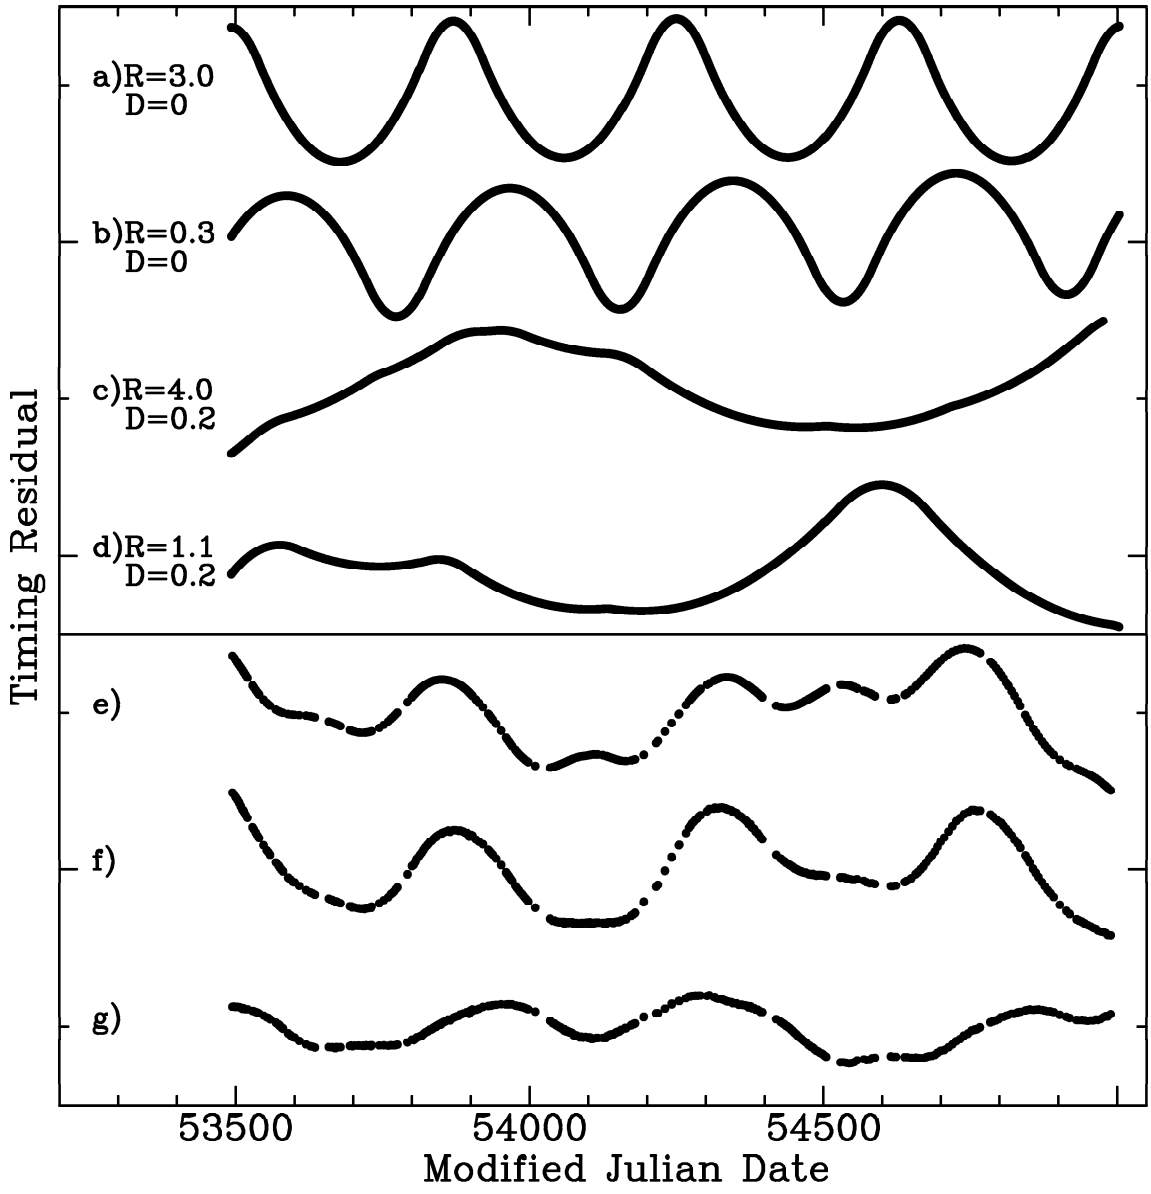

Figure S1: Simulations of timing residuals. a)-d) Simulations of a two-state spin-down model for pulsar timing noise in which only two parameters determine the form of the simulations, namely the ratio R of time spent in high and low spin-down states, and the rms fractional dither D in the switching period (See supporting text). Note how dither in the switching period can give rise to low-frequency structure in the residuals. e) simulated timing residuals for PSR B1828−11 in which the spin-down state is determined purely from the observed pulse shape parameter f) observed timing residuals for PSR B1828−11 from a simple spin-down model, which shows most of the features predicted by e), and g) the difference between the observed and simulated timing residuals. In spite of the severe undersampling of the shape parameter due to telescope availability ($< 1\%$ of the time), this demonstrates how it might be possible to "correct" the times-of-arrival for spin-state variations indicated by the pulse shape.

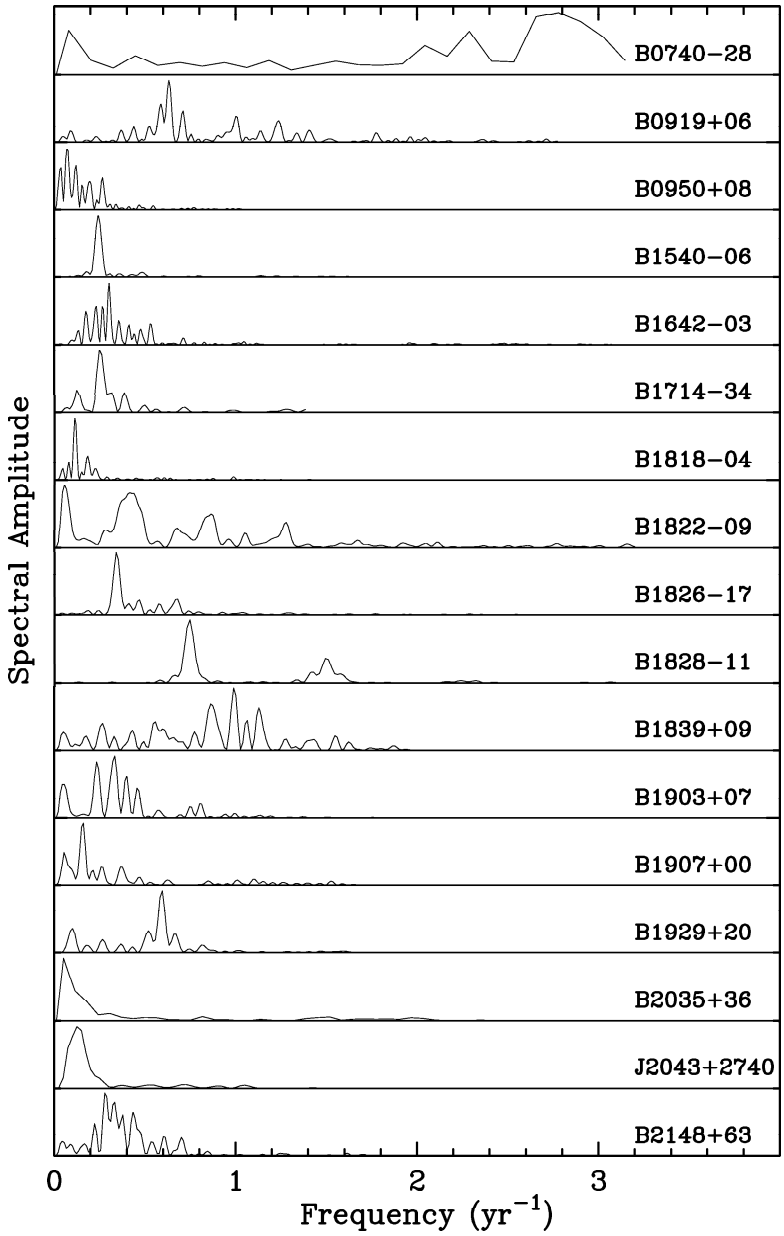

Figure S2: Lomb-Scargle spectra of the spin-down rates $\dot{\nu}$ presented in Fig. 2 of the journal paper for 17 pulsars. The peaks of all the spectra have been normalised to the same amplitude.

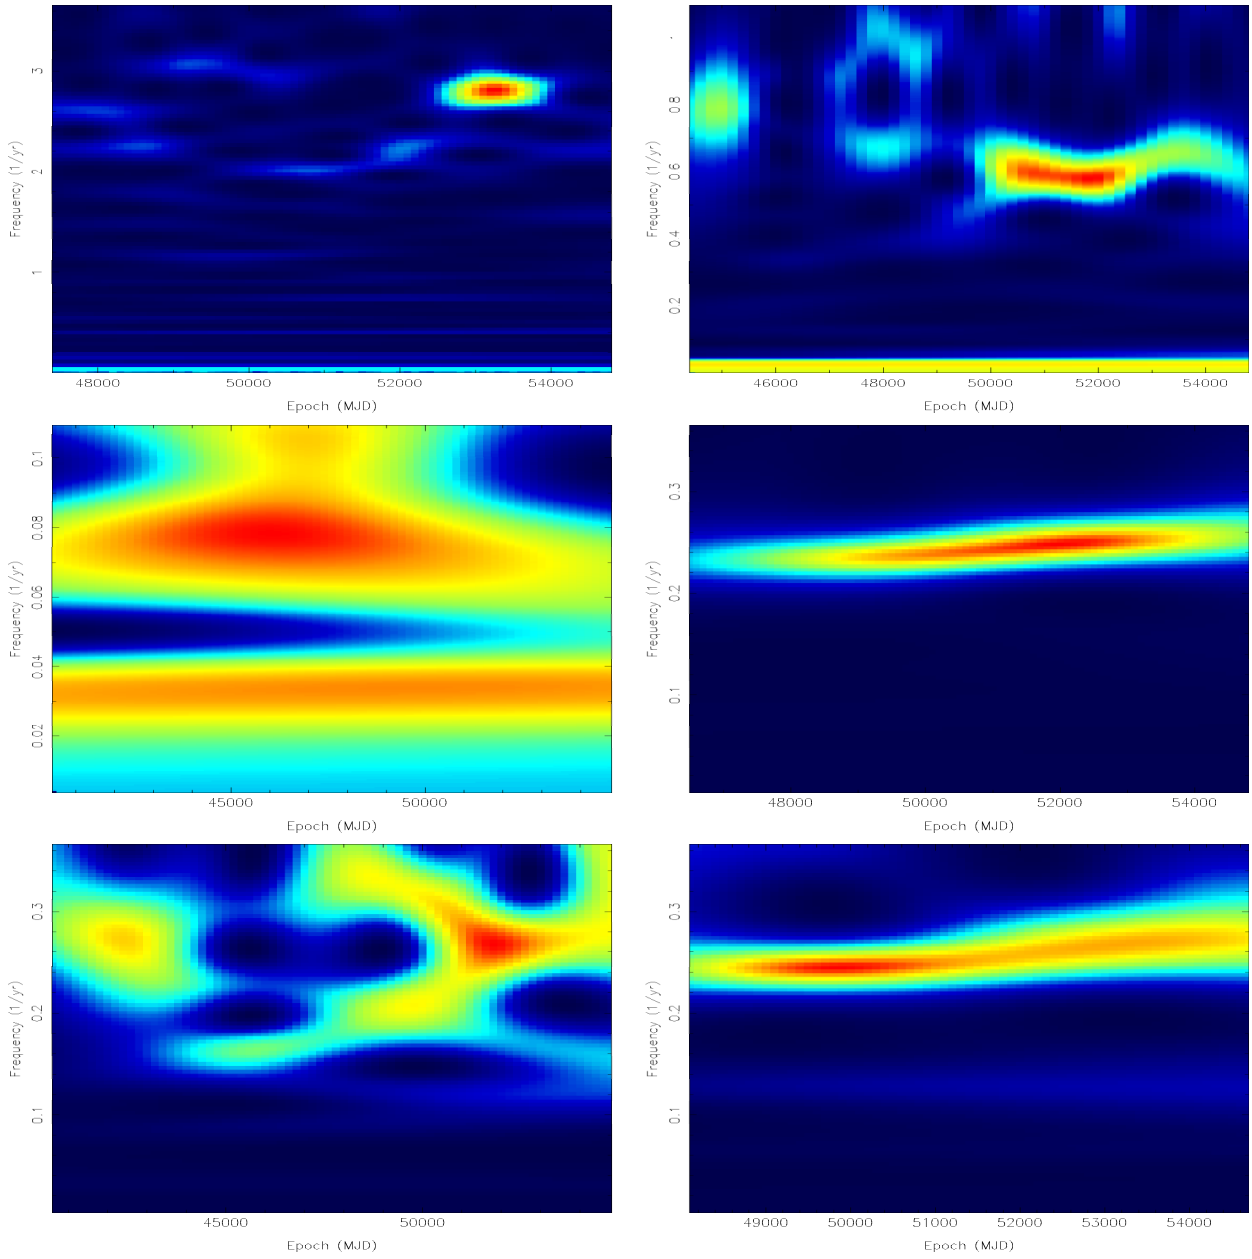

Figure S3: a) Wavelet spectra for pulsars B0740−28 and B0919+06 (top row), B0950+08 and B1540−06 (middle row) and B1642−03 and B1714−34 (bottom row). The frequency ranges shown cover the periodicities suggested in Fig. S2. The wavelet Z-statistic is computed as a function of both time and frequency. The wavelet "window" can be specified by a "decay constant", c, that defines the number of cycles of a given frequency f expected within the window. Values between 0.001 and 0.01 were chosen in an attempt to obtain the best compromise between frequency and time resolution, given the data in Fig. 2. The results agree well with the periodicities derived from the Lomb-Scargle analysis but demonstrate that for some sources the effective frequencies are varying or not always present.

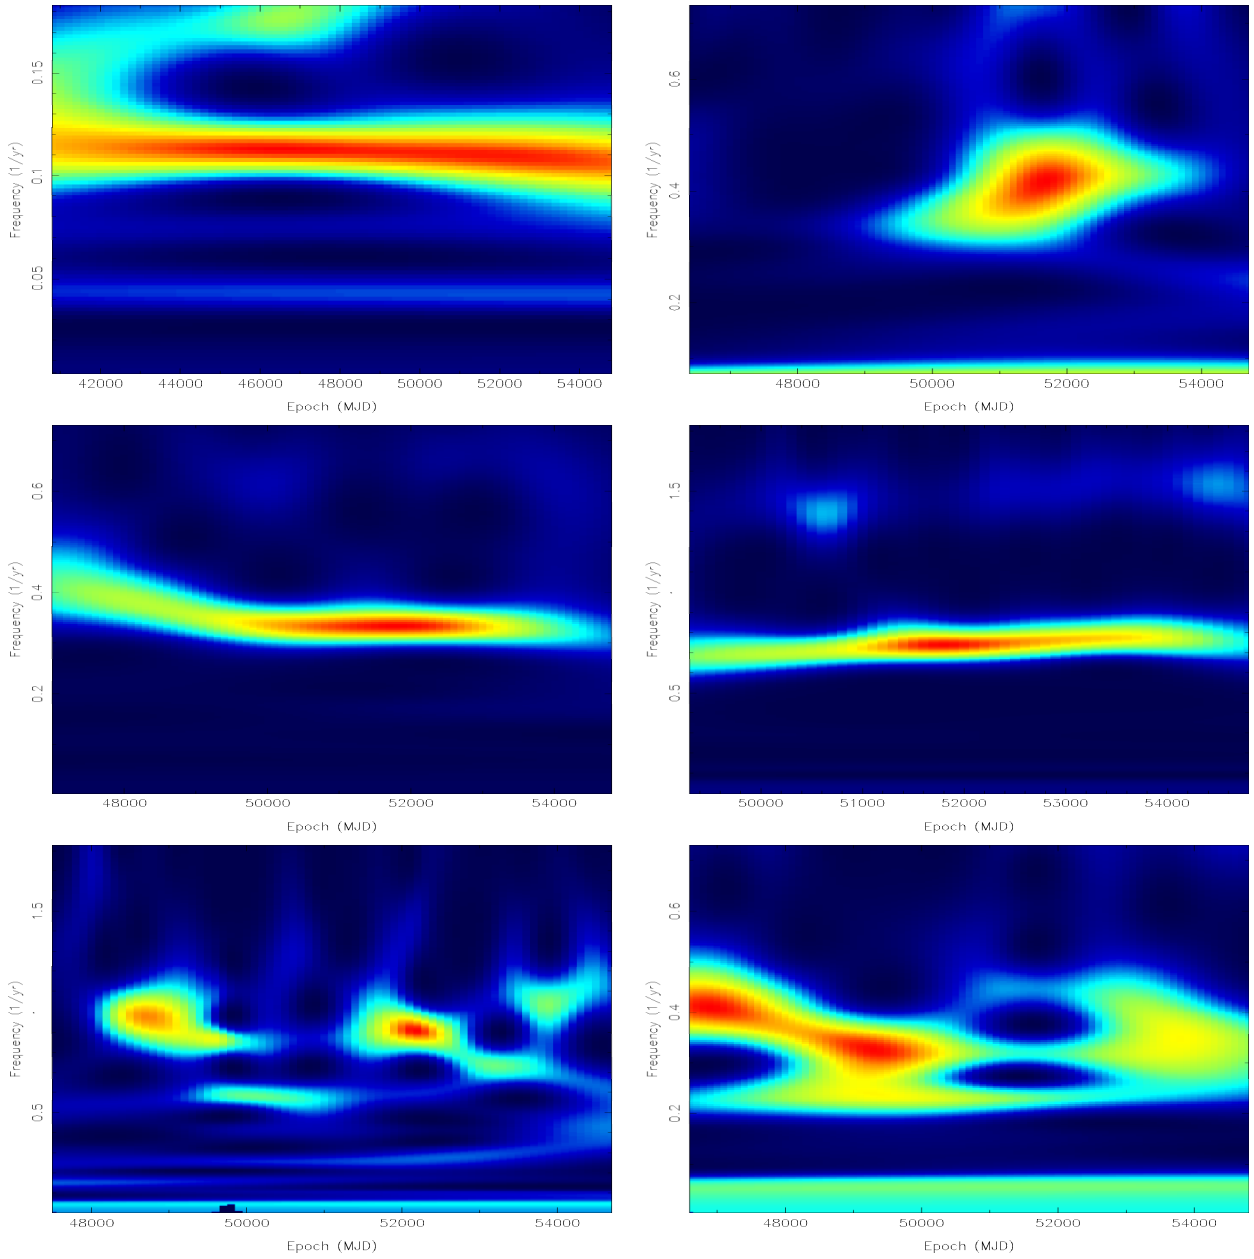

Figure S3: b) Wavelet spectra for pulsars B1818−04 and B1822−09 (top row), B1826−17 and B1828−11 (middle row) and B1839+09 and B1903+07 (bottom row).

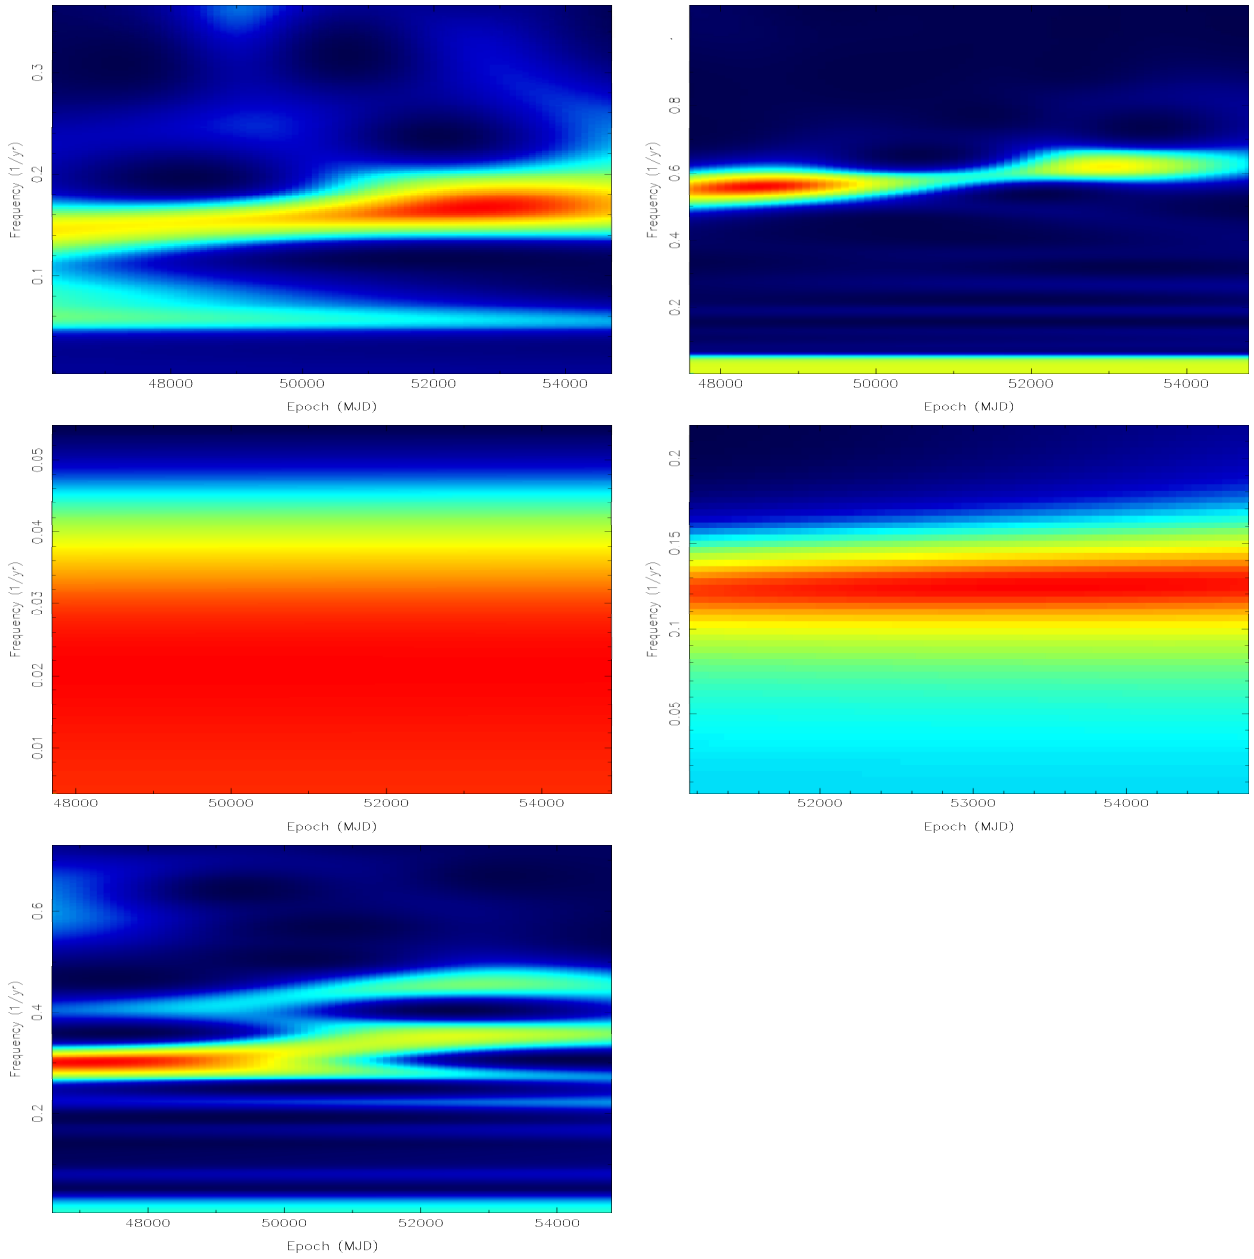

Figure S3: c) Wavelet spectra for pulsars B1907+00 and B1929+20 (top row), B2035+36 and J2043+2740 (middle row) and B2148+63 (bottom row).

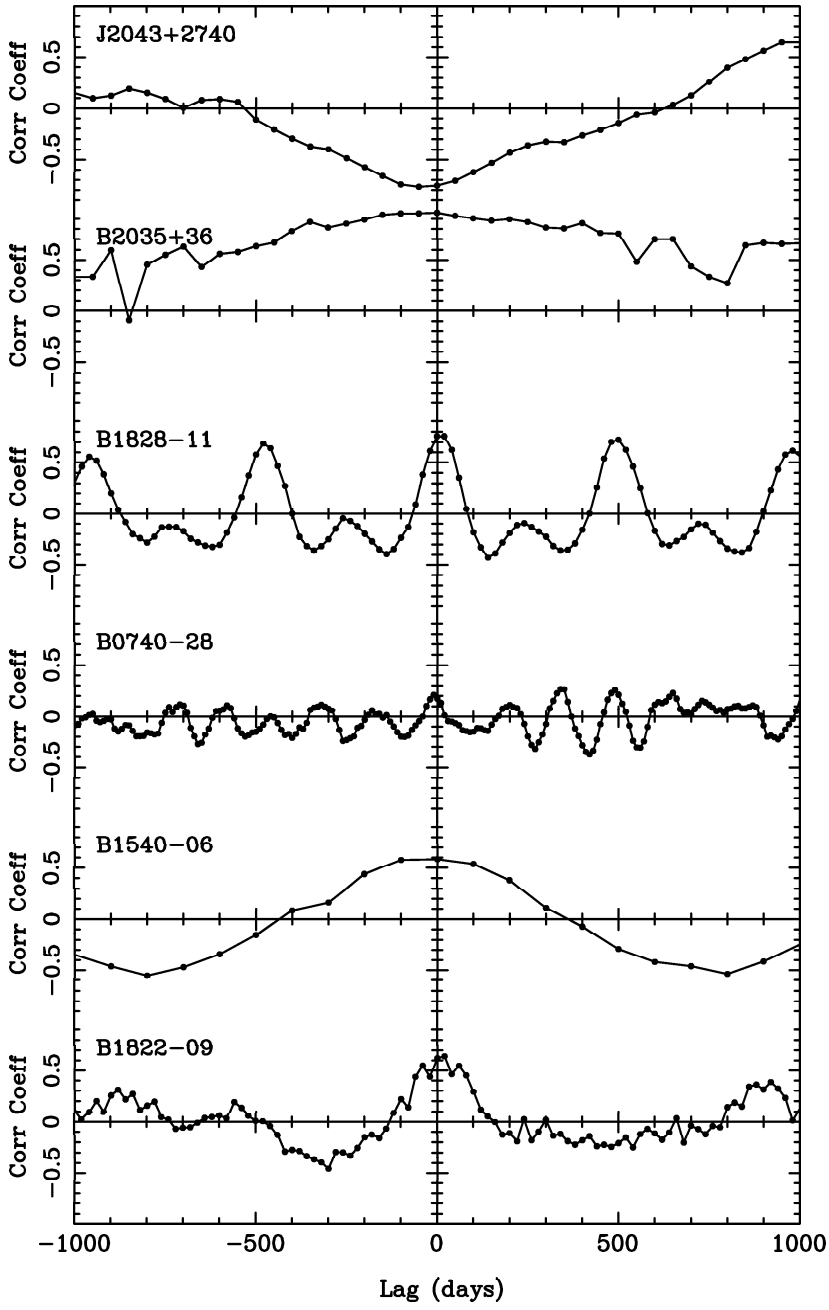

Figure S4: Cross-correlation functions between the average values of pulse shape parameters and the spin-down rates $\dot{\nu}$ shown in Fig. 4 of the journal paper for six pulsars. Note that in all cases, the magnitude of correlation coefficient always peaks close to zero lag.

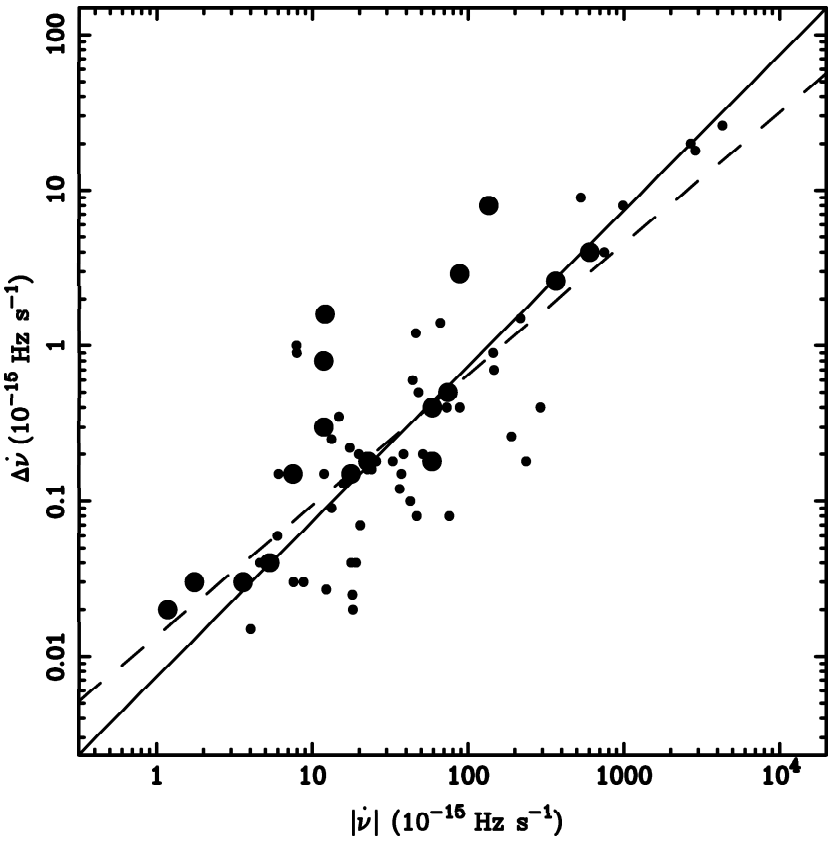

Figure S5: The dependence of the magnitude of the switching in $\dot{\nu}$ upon the magnitude of the average spin-down rate $|\dot{\nu}|$ for the 17 pulsars in Table 1 excepting PSR B1931+24 (large symbols) and another 51 pulsars which were studied in (*S2*) (small symbols). There is an excellent correlation between the logarithm of the variables with a correlation coefficient of 0.80. The solid line has a slope of unity and the dashed line is the best-fitting straight line, which has a slope of 0.84 ± 0.08.