Supporting Information

A recent method for estimating and removing noise from dMRI data relies on principal component analysis (PCA), where the highest principal components, accounting for the majority of the variation in the data, are assumed to be signal, and the lowest components are assumed to be noise (reference [18] in main article). This strategy was employed to estimate and compare the contributions of noise from the single band (SB) and multiband (MB) imaging sequences employed in this work. PCA was performed on all 3 sets (SB, MB with matched TR, and MB with full acceleration) of diffusion-weighted images (DWIs), and the least significant principal component was considered as noise. The local standard deviation of the noise in each image was estimated by calculating the standard deviation of the noise value within a 3x3x3 neighborhood, and assigning that value to the central voxel of that neighborhood. This process is diagramed in Supporting Information Figure S1. Additionally, the L_1 norm of the noise images was calculated from every voxel within the brain, and used to determine an estimate of the overall difference in noise levels between SB and MB images.

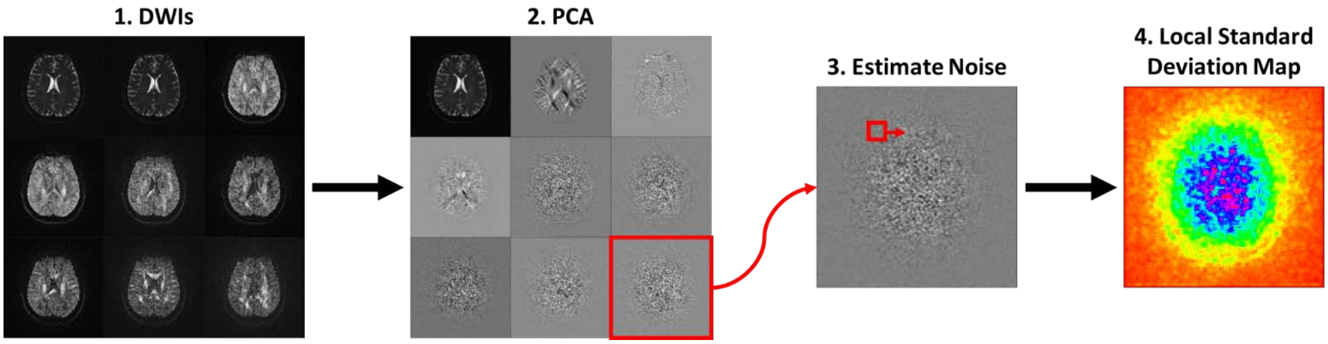

Supporting Information Figure S1: Pipeline for estimating the relative contribution of noise from each acquired sequence. Panel 1 shows representative (9 out of 130) DWIs acquired for the SB sequence. Panel 2 shows example components (9 out of 130) from the PCA decomposition of the DWIs. Panel 3 shows the resulting noise map (approximated as the least significant principle component), and local standard deviation estimation depicted as a small, 3x3x3 moving red box. Finally, panel 4 shows an example of the resulting standard deviation map.

Results: Supporting Information Figure S2 shows the generated noise maps (A-C), and their corresponding local standard deviation maps (D-F). As expected from any sequence employing parallel imaging (in-plane and in slice direction), the magnitude of noise is increased in the center of the image, where there is increased ambiguity in the coil sensitivity profiles. This is reflected in the noise maps and the local standard deviation maps. Further, both datasets collected using MB show increased noise in the center of the image when compared to the SB image. Lastly, the L_1 norm of the MB image with maximal acceleration was 1.1% larger than the L_1 norm computed across the SB image. The L_1 norm of the MB with matched TR was 2.3% larger than the L_1 norm of the SB image.

Discussion: As might be expected, the data generated using a MB imaging sequence has increased noise when compared to that generated using a traditional SB sequence. The effect is more prominent in the central portion of the FOV, likely due to the increased parallel imaging factor. This is also reflected in the analysis of the L_1 norm. It is possible that the noise maps derived from the PCA decomposition are not purely noise, but the above methodology provides a straightforward means of directly comparing the noise characteristics between the three sequences. This analysis, while somewhat qualitative, is consistent with the findings in dMRI-derived metrics presented in the main body of the work.

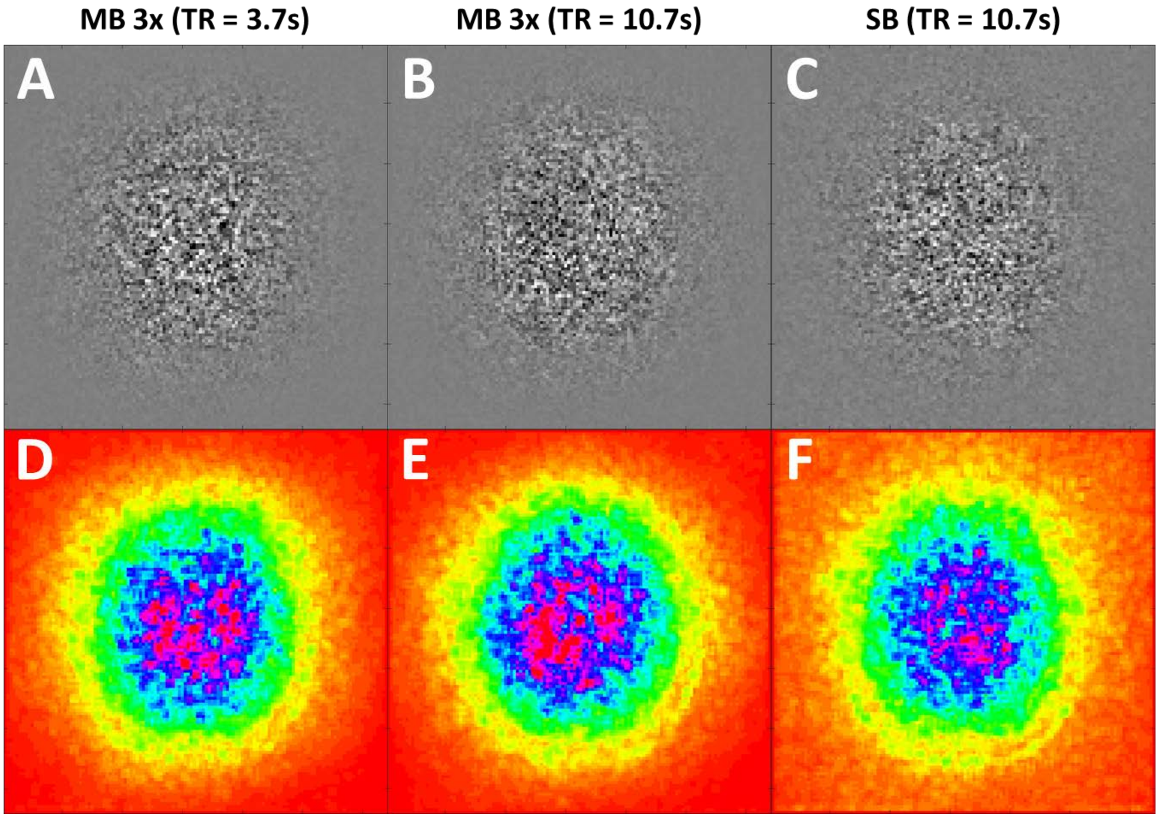

Supporting Information Figure S2: Noise maps (A-C) derived from a PCA decomposition of the data: A) fully accelerated MB sequence, B) TR-matched MB sequence, and C) SB sequence. Panels D-F show maps of the local standard deviation of the noise, estimated from a 3x3x3 moving window throughout the noise maps in A-C. D) fully accelerated MB sequence, E) TR-matched MB sequence, and F) SB sequence. Panels A-C have the same gray scale (arbitrary units) and D-F are shown on the same color scale (arbitrary units).