SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

TABLE OF CONTENTS

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

SUPPORTING FIGURES

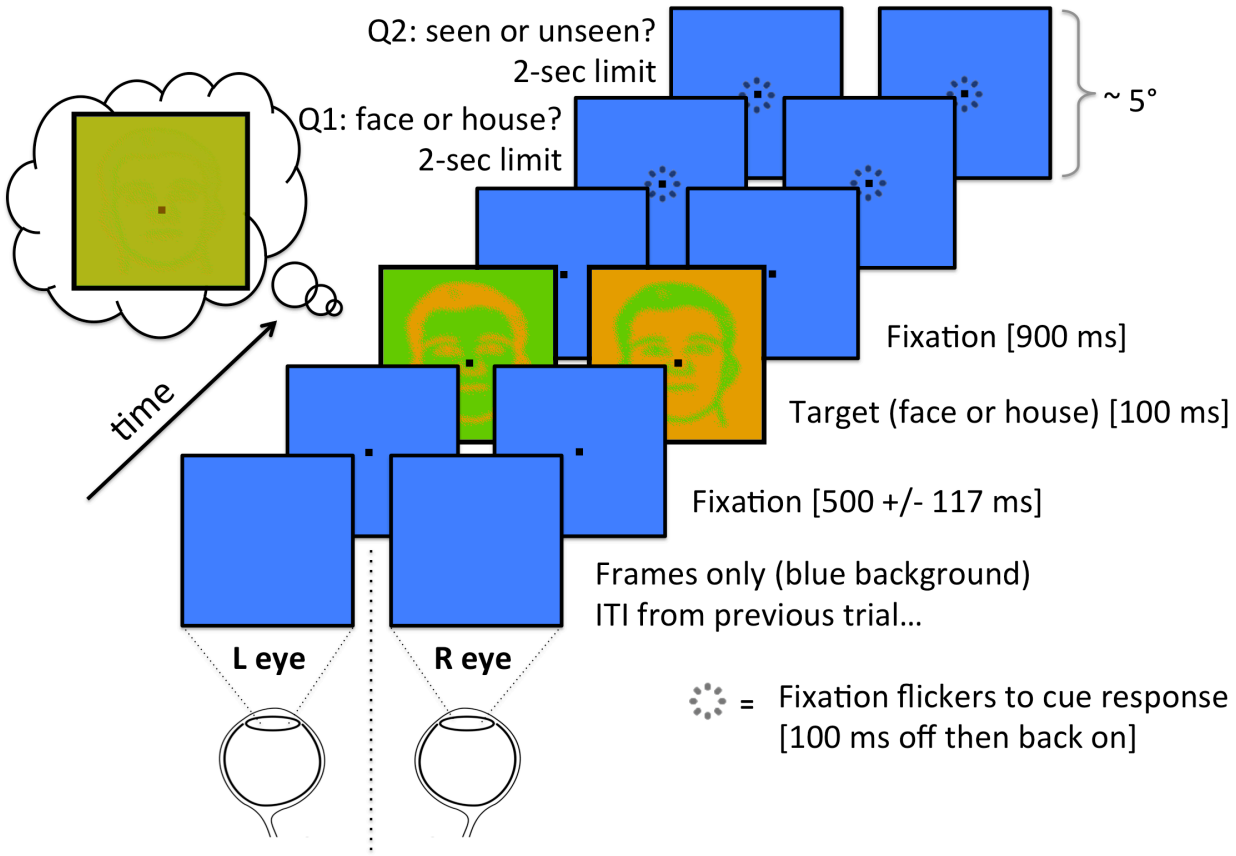

Figure S1: Trial sequence

Two-color isoluminant line drawings of faces (shown) and houses were used as stimuli, were presented separately to the two eyes using prism lenses and a cardboard divider (Schurger, 2009, J Neurosci Methods 177(1): 199-202). Stimuli were always presented with the opposite figure-ground color assignment (as shown), and the color contrast was varied to manipulate awareness of the object. At low color contrast, dichoptic fusion of the two images may result in the subjective impression of a uniform yellow square. The appearance of the fixation point ~ 500 ms before stimulus onset cued the subject to maintain steady fixation for the next two seconds. The stimulus duration was always 100 ms followed by 900 ms of fixation only. The fixation point then blinked to cue the subject to respond to the first question (2-second timeout) and then blinked a second time to cue the subject to respond to the second question (2-second timeout).

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

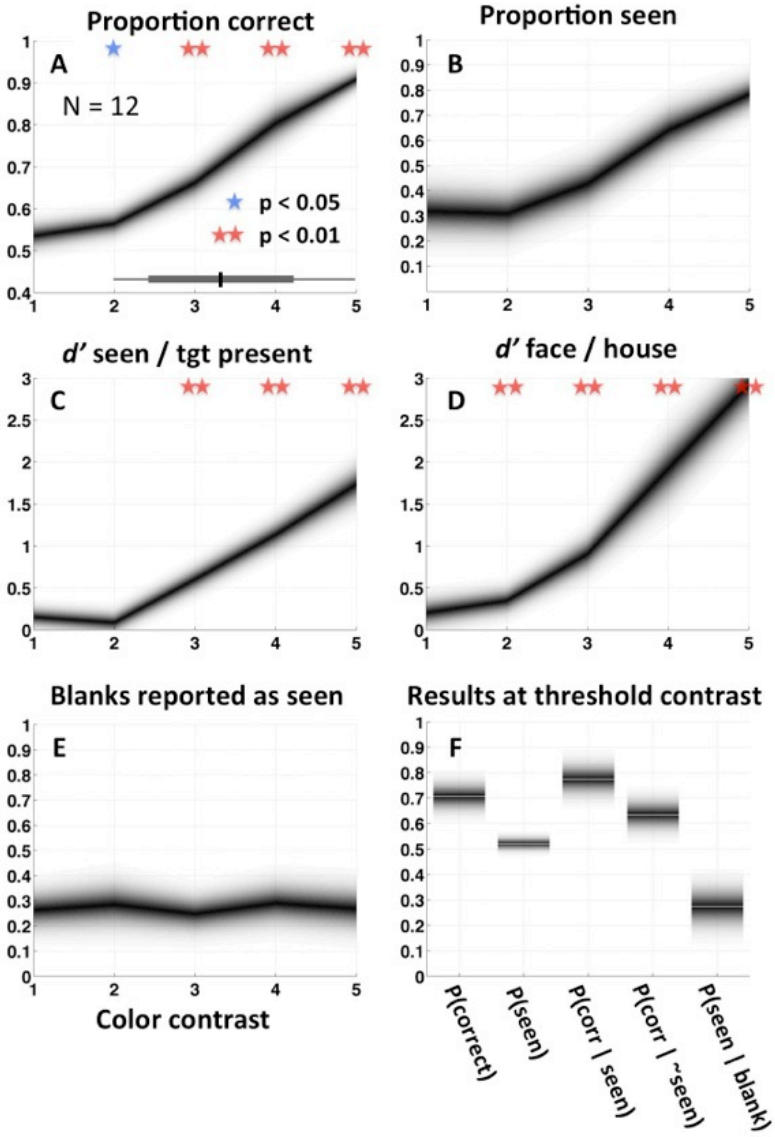

Figure S2: Behavioral results

Proportion correct (***A***), proportion seen (***B***), detection d-prime (***C***), face-house d-prime (***D***), and proportion of blank control stimuli reported as 'seen' (***E***) at each of the five levels of color contrast. Detection d-prime was computed by considering each 'seen' response on a target-present trial to be a "hit" and each 'unseen' response on a target-absent to be a "correct rejection" (see Methods). Face-house d-prime was computed by considering *face* trials to *be target-present* and *house* trials to be *target-absent*, with a "face" response on a *face* trial considered to be a "hit". At the top of each panel, a single-star means "significant at $p < 0.05$" and a double-star means "significant at $p < 0.01$" (signed rank test). Shades-of-gray error boundary extends out to 99% confidence on a t-distribution (to give a qualitative impression; statistics were computed using a non-parametric test). **F** shows the same five proportions as in **A** – **E**, but for the specific threshold-level color contrast chosen independently for each subject. Horizontal error bar at the bottom of A shows the span of the contrast levels used (min, max, mean +/- stddev).

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

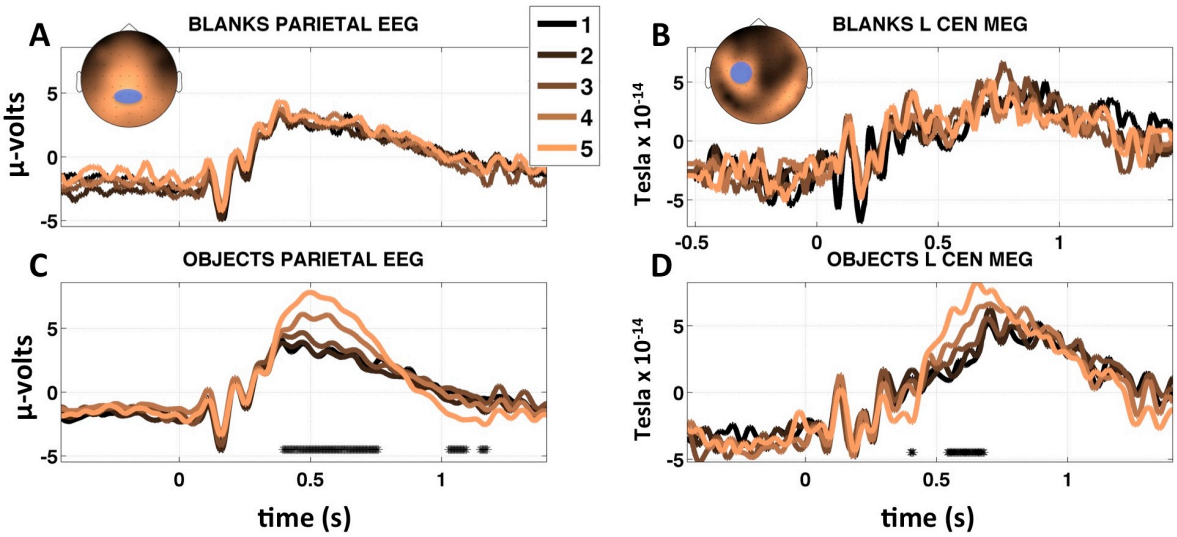

Figure S3: Evoked potentials and evoked fields, control and object stimuli

Left panels show the avereage over three parietal EEG sensors (P1, Pz, P2) for blank control stimuli (**A**) and object stimuli (**C**) at the five different levels of color contrast (color coded). Right panels show the average over two left-central magnetometers (MEG-0441, MEG-1811) for blank control stimuli (**B**) and object stimuli (**D**) at the same five levels of color contrast. Small black stars placed at the bottom of each axis mark time points where a one-way ANOVA across the five contrast levels was significant at $p < 0.01$ (corrected for temporal non-independence). No significant effect of color contrast was found at any sensor (EEG or MEG) for the blank control stimuli. N = 12.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

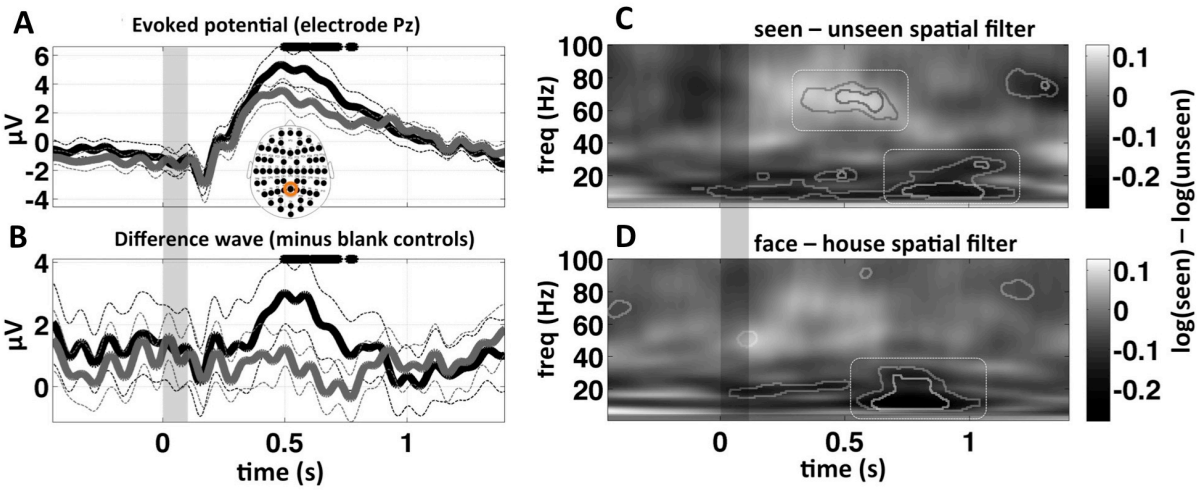

Figure S4: Evoked potential and spectral correlates of *seen* versus *unseen* objects

(**A**) average evoked potential at EEG electrode Pz for "seen object" (black) and "unseen object" (gray) trials at a fixed, threshold level of color contrast chosen individually for each subject. Error boundary (thin dashed line) is standard error of the mean. (**B**) same as (**A**), but showing the difference "seen" object – *control* (black) and "unseen" object – *control* (gray). Black stars at the top of A and B indicate timepoints where the difference seen – unseen was significant at $p < 0.01$ (corrected for temporal non-independence). (**C**) difference between the mean time-frequency (TF) decomposition of "seen" trials and the mean TF decomposition of "unseen" trials (computed using Morlet wavelets; see Methods), from MEG data. The TF decomposition was performed on the output of a spatial filter trained at each time point on the difference between "seen" and "unseen" trials ("S-U filter"), and applied to the gradiometer data using a leave-one-out cross validation. (**D**) same as (**C**) except that the spatial filter was trained on the difference between *face* and *house* trials ("F-H filter"). In both **C** and **D**, outer (lighter) contour lines mark regions in TF space where the difference was significant at $p < 0.05$, and inner (darker) contour lines mark regions were the difference was significant at $p < 0.01$ (point-wise signed rank test, N=12). Dashed rounded squares highlight regions, bounded by the outer contour line, where the difference was significant at $p < 0.01$ after a cluster-based permutation test (see Methods) to correct for multiple comparisons. N = 12.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

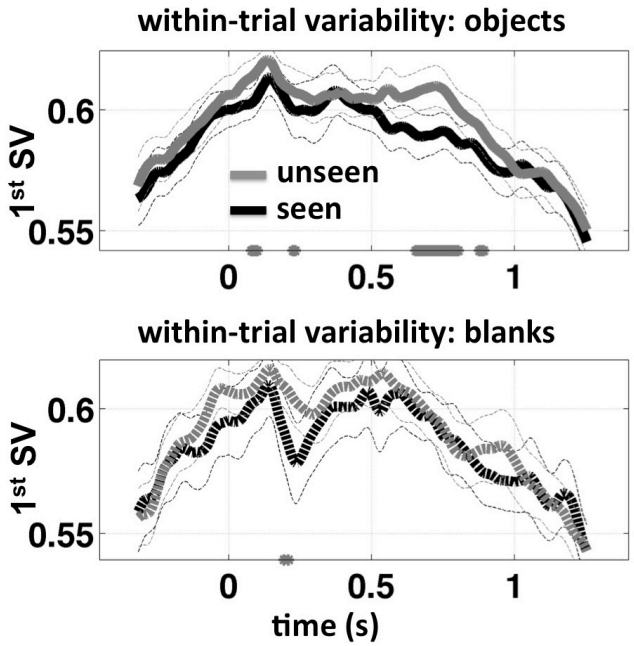

Figure S5: Using a singular-value decomposition (SVD) to measure stability

Another way to measure stability of a set of variables (in this case, MEG sensors) over a window of time is to apply a singular-value decomposition[1] (SVD) to the time-x-sensor matrix and look at the variance accounted for by the first singular vector (SV). . This was computed as the ratio of the square of the first singular value to the sum of all the singular values squared ($(S_1)^2/\sum_i(S_i)^2$). In the limit, if the exact same pattern were repeated across the entire time window (the upper limit of stability) then all of the variance would be accounted for by the first SV. At the other extreme, if the patterns were all orthogonal to one another (the lower limit of stability) then the variance would be shared evenly among the different factors, in which case much less than the total amount of variance would be accounted for by the first SV. Hence, the variance accounted for by the first SV also gives an estimate of stability, when applied to a matrix where the observations are measurements made at equally spaced time intervals. As an independent verification of the result obtained with directional variance (fig. 2C and 3C), we used the SVD in this way as a means of measuring stability for object (top) and blank (bottom) stimuli, and obtained qualitatively similar results. Here we present $1 - v$ instead of v (where v is the variance accounted for the by the first SV and $0 \leq v \leq 1$) in order to be consistent with figs. 3 and 4. Gray stars at the bottom of the plot show time points where the difference between "seen" and "unseen" is significant at $p < 0.05$ (corrected for temporal non-independence).

[1] We also repeated the same procedure using principal components analysis (PCA, which centers the data around zero) and obtained the same result.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

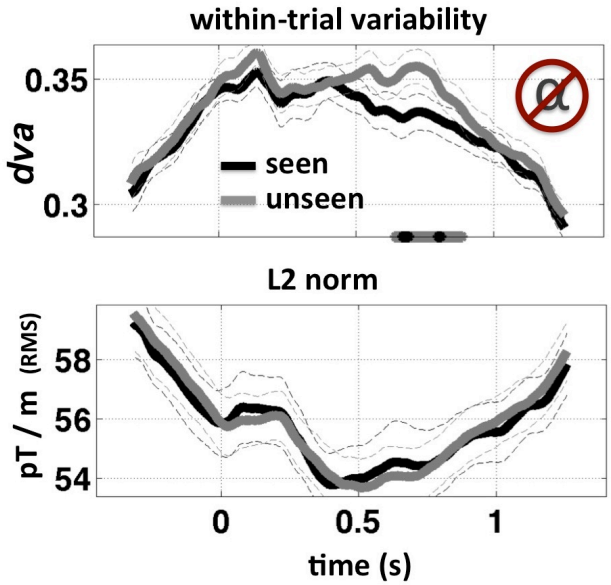

Figure S6: Control for alpha suppression

Same as fig. 2C and E except that a 7 – 13 Hz band-stop filter was applied to the data before running the analysis to compute *dva* and the spatial L2 norm. This was done to rule out a possible difference in stability due merely to the coincident suppression of alpha power on "seen" trials (see fig. S4D). Stars at the bottom of each plot indicate time points where the difference between seen and unseen trials was significant at $p < 0.05$ (gray) or $p < 0.01$ (black), corrected for temporal non-independence using a cluster-based permutation test.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

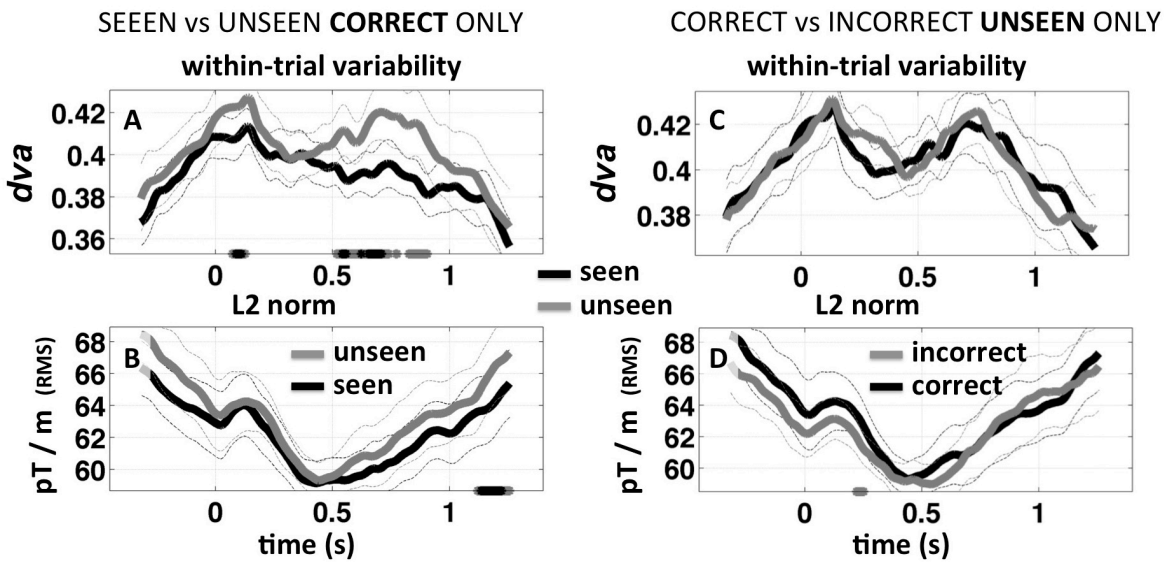

Figure S7: Control for task performance

Task performance tends to be higher on "seen" as compared to "unseen" trials, for a fixed level of stimulus contrast (fig. S2F). In order to rule out an effect of task performance on the difference between "seen" and "unseen" *dva* we performed the control analyses shown above. **A** and **B** (left) show the within-trial directional variance and L2 norm for seen (black) and unseen (gray) trials on which a correct response was given. The effect is still present even when the data are restricted to correct trials. A comparison of correct versus incorrect "unseen" trials (**C** and **D**) shows no effect of task performance alone (there were too few incorrect "seen" trials to perform the same analysis on "seen" trials). Stars at the bottom of each plot indicate time points at which the difference (seen – unseen in **A** and **B**, correct – incorrect in **C** and **D**) was significant at $p < 0.05$ (gray) or $p < 0.01$ (black), corrected for temporal non-independence.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

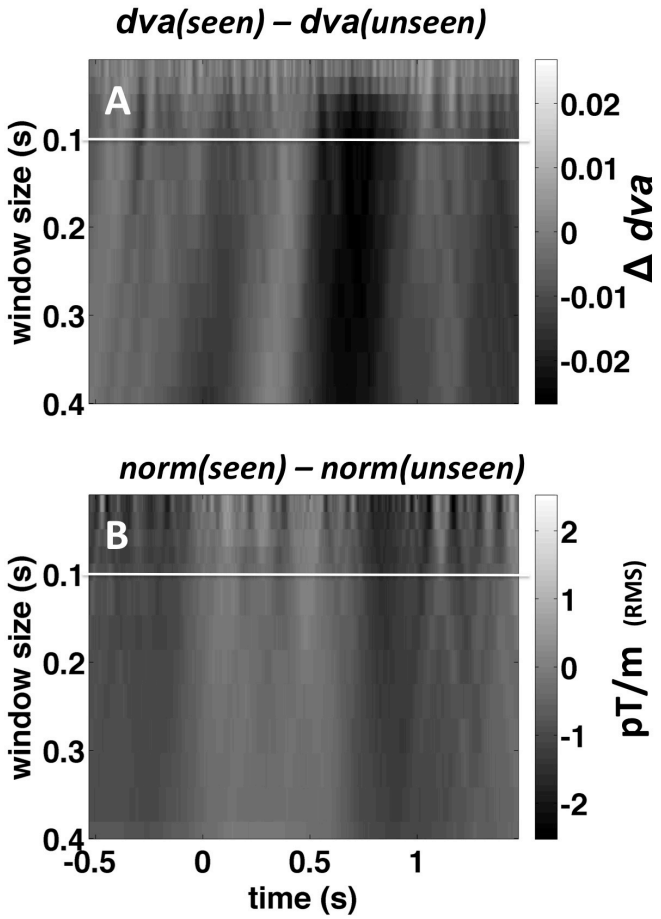

Figure S8: Within-trial directional variance and norm for a range of different sliding window widths

Horizontal white line marks the window width used for figures 2 and 3. The period of relative stability appears as a vertical black trough in panel **A**, and its duration is roughly constant at ~ 300 ms regardless of the width of the sliding window.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

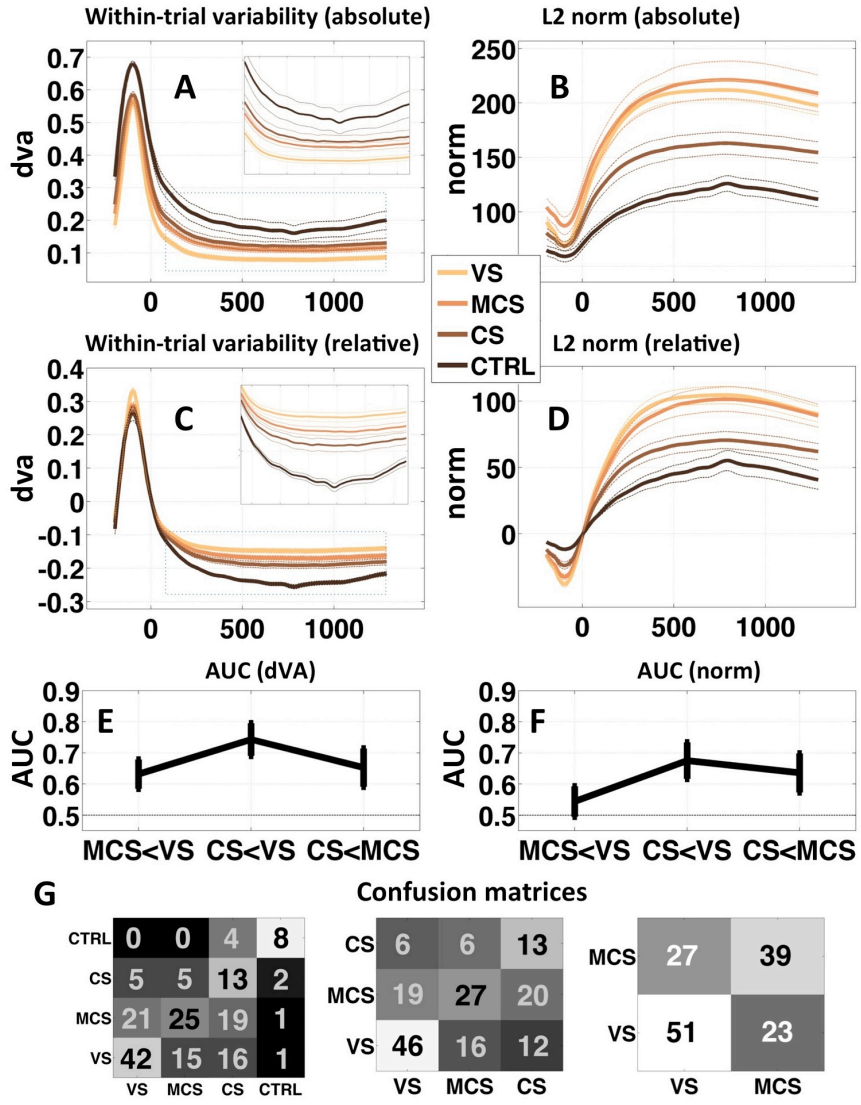

Figure S9: Patient data

A and **B** (top row) show the mean absolute time course of within-trial dva (**A**) and norm (**B**) for each patient category and the healthy control subjects (color coded). Thin dashed lines are standard error of the mean. **C** and **D** (second row) show the same time courses relative to the value at t_0 (onset of the first tone). **E** and **F** show the AUC for discrimination between each pair of patient categories based on dva (**E**) and norm (**F**). Error bars are standard error. **G** (bottom) shows confusion matrices resulting from the use of linear-discriminant analysis (LDA) to classify patients based on their combined mean dva and norm. The three confusion matrices, from left to right, were from separate analyses run either with all patients and healthy controls (left), only the patients (middle), or only the VS and MCS patients (right). The classifier performed significantly above chance ($p < 0.01$) in all three cases (49.6% correct, 50% correct, and 65% correct, respectively, with bootstrap-estimated chance performance being 26.5%, 33%, and 50%, respectively).

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

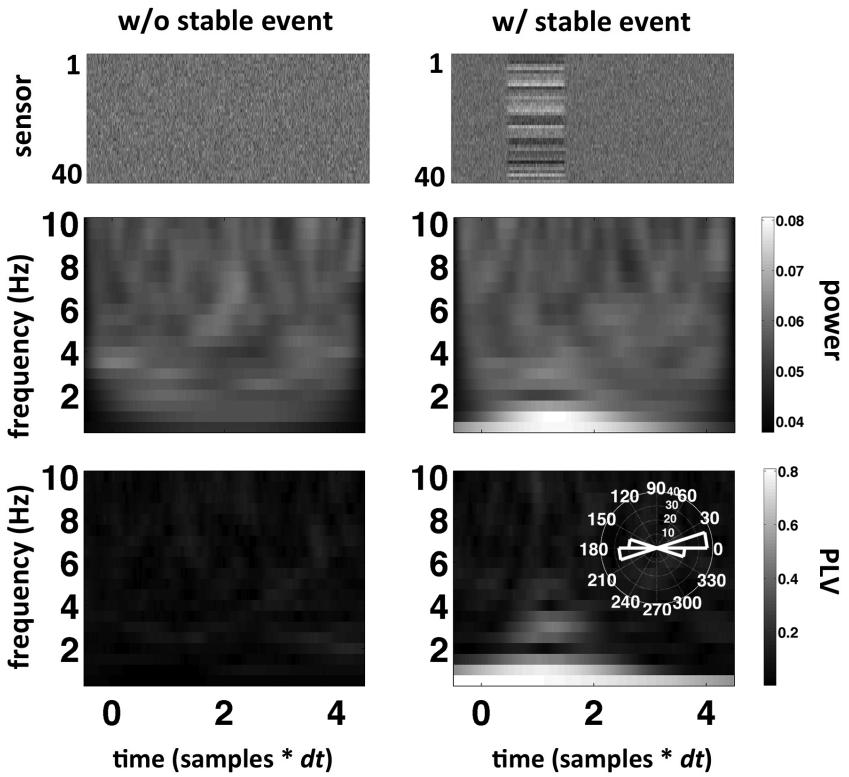

Figure S10: Stability predicts low-frequency power and phase locking

Each column presents the results of a simulation (200 trials), where each trial has 40 channels and 1280 time samples (5 seconds at 256 Hz). The top panel of each column (**A** and **B**) simply shows the average over the 200 simulated trials. For simulation 1 (left column), each trial was generated from Gaussian random noise, and each sample (i.e. each column in the matrix) was divided by its own norm so that all samples of every trial had exactly the same amount of energy. The only difference between the data used for simulation 1 and simulation 2 (right column) was that in simulation 2 the (random) pattern of activity across the 40 simulated sensors was held constant across time, from 0.5 to 1.5 sec, and across trials. However, the norm across the 40 sensors was always made = 1 at every time point, on every trial, just as in simulation 1 (left). From the point of view of a single simulated sensor, this stable and reliable pattern leads to an increase in spectral power at the frequency whose period is twice the duration of the period of stability (in this case 0.5 Hz; **D**). It also yields significant phase locking at around the same frequency (**F**). In addition, the phase of sensors showing a significant effect in the time-locked average tends to cluster around 0 and 180 degrees (**F** inset). The presence of a stable pattern guarantees that this will be the case, because sensors showing a significant effect will tend to be the ones that are most "active" in the pattern, i.e. whose amplitude is farther from zero (whether positive or negative). Thus reproducibility and stability can account for the prior observation of significant low-frequency energy and phase locking associated with conscious perception (Li et al, *J Neurosci*, 2014, 34(12):4382-4395), but provide a different explanatory perspective.

SUPPORTING INFORMATION APPENDIX

Cortical activity is more stable when sensory stimuli are consciously perceived

Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene (2015)

SUPPORTING METHODS

Dichoptic color masking

Stimuli were simple line drawings of faces and houses (fig. S1), 175 x 175 pixels, rendered in exactly two colors using error-diffusion dither: After blurring the initial grayscale images, the images were dithered into different densities of black and white pixels. During each trial of the experiment each black pixel in one of the images was replaced by a shade of pale green and each white pixel replaced by an isoluminant shade of pale orange (or vice-versa). The opposite color assignment was used for the other eye resulting in an orange foreground / green background in one eye, and green background / orange foreground in the other (or vice versa; see fig. S1). Throughout the experiment stimuli were always presented in the opposite color configuration and the color contrast and foreground / background color assignments were varied pseudo-randomly from trial to trial.

The subjective visibility of the stimuli was manipulated by varying the color contrast. At moderately-low color contrast, the two colors fuse cortically resulting in the percept of a uniform yellow square (fig. S1), even though the object would otherwise be visible at the same color contrast if presented in the same color assignment to both eyes, or simply viewed with one eye occluded. At progressively higher color contrast the effect of dichoptic color masking is broken and contours become (faintly) visible. During the experiment we used a range of five different color contrasts chosen such that the middle contrast level was near the average subjective threshold (determined from pilot testing). We have labeled these '1' through '5', with '1' being the lowest color contrast.

Subjective isoluminance within and between each of the five different color pairs was established by using heterochromatic flicker photometry (HcFP) during a pre-experiment calibration performed by the experimenter. In addition to the two colors within each orange / green color pair being isoluminant, the five different color pairs were also isoluminant with each other. This was done by finding five shades of orange, at five linearly increasing values of R in RGB coordinates, with the corresponding values for G chosen using HcFP such that levels 1, 2, 4, 5 were each isoluminant with level 3 (which was in the vicinity of RGB = [0.5, 0.4, 0]). Then for each corresponding shade of green (e.g. [0.4, 0.5, 0] for level 3) we used HcFP to find a coefficient, $\beta < 1$, such that the luminance of β * *green* was equal to the luminance of *orange*. We double-checked the final color set by performing HcFP on each *orange-green* pair (25 comparisons), and made small adjustments accordingly.

HcFP color isoluminance procedure

Approximate isoluminance within and between each of the five different color pairs was established by using heterochromatic flicker photometry (HcFP) during a pre-experiment calibration performed by the experimenter (see Supplementary Material). In addition to the two colors within each orange / green color pair being isoluminant, the five different color pairs were also isoluminant with each other. This was done by finding five shades of orange, at five linearly increasing values of R in RGB coordinates, with the corresponding values for G chosen using HcFP such that levels 1, 2, 4, 5 were each approximately isoluminant with level 3 (which was in the

vicinity of RGB = [0.5, 0.4, 0]). Then for each corresponding shade of green (e.g. [0.4, 0.5, 0] for level 3) we used HcFP to find a coefficient, $\beta < 1$, such that the luminance of $\beta * green$ was equal to the luminance of *orange*. We double-checked the final color set by performing HcFP on each *orange*-*green* pair (25 comparisons), and made small adjustments accordingly. The background of the stimulus display was blue, and adjusted so as to be similar in luminance to the stimuli, so that the onset of stimuli would be less likely to provoke blinks or pupillary reflexes, which might disrupt fusion of the two images.

Simulations

We performed simulations using MatLab (MathWorks, Inc.) in order to validate dva and test its behavior under different controlled conditions (fig. 1 C & D and fig. 7). We generated matrices of Gaussian random data with 40 "sensors" and 256 "samples" (equivalent to 1 second of data at 256 Hz sampling rate) – each matrix simulated a single trial. Beginning at a fixed time point in each simulated trial, a randomly-generated pattern that remained steadily "on" for 25 samples (the equivalent of ~100 ms) was added to the noise, with a Gaussian-smoothed onset and cutoff. The amplitude of the fixed pattern with respect to the noise could be varied for each run of the simulation (with 200 – 1000 trials per run, depending on what was being tested). We then divided each column of the result by its own norm, so that only differences in the pattern, but not the norm, would remain, and then added another layer of Gaussian noise to the matrix. Directional variance and L2 norm were computed on these data in the very same way that they were computed on the empirical MEG data in order to verify that the presence of a very weak, but stable pattern lead to a significant increase in angular coherence even when there was no difference in the mean norm (fig. 1C). Varying the amplitude of the fixed pattern across a range of values allowed us to inspect the way dva changes as a function of SNR (fig. 1C inset).