A novel signature of visual awareness independent of post-perceptual processing

Michael A. Cohen[1,2,5], Cole Dembski[2,3,5], Kevin Ortego[4], Clay Steinhibler[3] & Michael Pitts[3]

[1] McGovern Institute for Brain Research, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology
[2] Department of Psychology and Program in Neuroscience, Amherst College
[3] Department of Psychology, Reed College
[4] Department of Psychological and Brain Sciences, Dartmouth College
[5] Denotes equal contribution

Abstract

What are the neural processes associated with perceptual awareness that are distinct from pre-conscious sensory encoding and post-perceptual processes such as reporting an experience? Here, we used electroencephalography (EEG) and a no-report visual masking paradigm with parametric manipulations of stimulus visibility to search for neural signatures associated with perceptual awareness independent from both early sensory processing and subsequent reporting. Overall, we found only one neural signal that was uniquely associated with perceptual awareness: a fronto-central event-related potential (ERP) that we refer to as the N2. In contrast, earlier ERP signals were linked with the linear stimulus manipulation, while later candidate signatures, such as P3b and temporal generalization of decoding, were present in the report condition but disappeared in the no-report condition. Taken together, these findings challenge several prominent theories of consciousness and offer a new avenue for exploring the neural mechanisms supporting conscious processing.

Introduction

What neural processes differentiate conscious and unconscious processing in the brain? To answer this question, researchers often compare neural responses in cases where observers report being aware vs. unaware of a particular stimulus, using paradigms such as visual masking (1,2), the attentional blink (3,4), or change blindness (5,6). Results from these paradigms have served as the foundation of several influential theories of consciousness, such as those that emphasize the importance of higher-order regions accessing sensory representations (7-12) and those that propose a key transition from information processing in sensory regions to a broader fronto-parietal network (1, 13-15).

However, many researchers disagree with these frameworks, arguing that activation in fronto-parietal regions is associated with post-perceptual processing (i.e., memory, decision-making, reporting a stimulus, etc.), not perceptual awareness per se (16,17). Under this view, conscious processing truly occurs in more posterior, sensory regions of the brain (18,20). The key methodological basis of this argument is that in the vast majority of these prior studies, observers were always asked to explicitly report whether they detected a target stimulus or to report the contents of their perceptual experience (e.g., “Was the target an animal or object?”). However, such judgements cannot be made when a stimulus is not perceived. Thus, in these report-based paradigms, certain post-perceptual processing steps occur only when the observer perceives the stimulus, often leading to an overestimation of the neural differences between conscious and unconscious processing (21).

This particular methodological insight led researchers to develop “no-report” paradigms in which observers are shown visible and invisible stimuli, but do not make any judgements about their perceptual experience (18,22-26). The overarching goal of no-report paradigms is to minimize the extent to which post-perceptual processing contaminates the neural activity associated with perceptual awareness. Indeed, a series of studies using a variety of experimental paradigms has shown that several candidate signatures of conscious processing, such as the P3b or activation in fronto-parietal circuits, disappear in no-report paradigms (27-31).

The findings from these no-report studies raise a natural question: If the P3b and fronto-parietal activation are not signatures of conscious processing, what are? Answering this question requires more than simply adding no-report conditions to the paradigms used in prior studies. For example, during the no-report condition of our prior visual masking study (29), we found a significant difference in the size of the early P1 wave for visible vs. masked stimuli (see Figures 3b and 4 of that study). However, it has been previously shown that the P1 does not track perceptual awareness (32,33). In this case, the difference in P1 amplitude was likely due to stimulus-based differences between the visible and invisible stimuli (i.e., lightly masked vs. heavily masked conditions).

Here, to better isolate the neural signals truly linked with perceptual awareness, we combined a no-report methodology with a parametric manipulation of stimulus strength. Specifically, we manipulated stimulus visibility by linearly increasing the stimulus onset asynchrony (SOA) between the target and the masks (i.e., evenly-spaced SOAs of 17ms, 33ms, 50ms, 67ms, and 83ms). Critically, even though the time between the target and mask

increases linearly (i.e., a stimulus-based difference), prior studies have shown that the rate at which the targets are consciously perceived (i.e., a perception-based difference) increases in a non-linear, bifurcated fashion (34-37). This contrast between linear changes in stimulus strength and non-linear changes in perceptual awareness enables a useful distinction when assessing the corresponding neural signals (34). Any neural signal that displays a non-linear, bifurcated pattern across SOAs is a potential correlate of conscious processing, while neural signals showing non-bifurcated (e.g., linear, logarithmic, etc.) patterns across SOAs can be ruled-out as candidates. Furthermore, by performing this experiment under conditions of no-report, we can also rule-out any of the bifurcated neural signals that are linked with post-perceptual processing, leaving only those that are associated with perceptual awareness per se. Thus, the main result that would link a neural signal to perceptual awareness would be the finding that in the no-report condition, the neural signal in question mirrored the behavioral responses from the report condition.

To preview the results, only one neural signal matched the behavioral results in this critical way: a novel ERP measured over a set of fronto-central electrodes, which we refer to as the N2. In the no-report condition, the N2 increased in a non-linear, bifurcated manner that closely mirrored the psychometric function of behavioral responses from the report condition. This N2 stands in contrast to the P1, which increased in a linear fashion across SOAs. Moreover, two other proposed signatures of conscious processing, the P3b (1) and late, sustained temporal generalization of decoding (38), while clear and robust in the report condition, disappeared entirely under conditions of no-report. Finally, an earlier signal (N170/VAN) approximated the shape of the psychometric function, but not as clearly as the N2 did in the no-report condition. Overall, these findings suggest that the neural processes associated with perceptual awareness independent of post-perceptual processing arise earlier than predicted by several prominent higher-level theories of consciousness (7,10,39), but later than predicted by some influential sensory theories of consciousness (18,33,40,41).

Results

Masking paradigm (report and no-report)

All of the methods and primary analysis procedures were pre-registered and can be found at the following link: osf.io/6d5av. Our masking paradigm was modeled after the experiments by Del Cul et al. (33) and Kouider et al. and used the same stimulus set as Kouider et al. (34). The targets were images of cropped faces on a black background and the masks were scrambled layers of upside-down faces and objects on a black background. On trials in which a target was shown, there was a variable stimulus onset asynchrony (SOA) between the target and the first mask (Figure 1A; see Methods). These SOAs linearly increased from 17ms to 33ms to 50ms to 67ms to 83ms presented in a randomized order. Two consecutive 100ms masks were always presented; on 1/6th of the trials, no target was shown (only the two masks).

During the report condition, participants were instructed to provide a button-press after every trial, indicating whether they did or did not detect a face. During the no-report condition, participants were instructed to press a button any time a green ring appeared on the screen. These green rings appeared on 15% of the trials during both the report and no-

report conditions, and any trial in which a green ring appeared was removed from analysis. It should be stressed that all stimulus parameters were identical across the two conditions. The only difference was the task the participants were instructed to complete (see Methods).

The behavioral results from the report condition are plotted in Figure 1B. Overall, we replicated prior findings showing a non-linear, bifurcated pattern of perceptual reports in which faces were rarely detected at the two shortest SOAs, were detected roughly half the time at the middle SOA, and were frequently detected at the two longest SOAs (34,35; see Supplemental Results for full statistics). This is critical because our EEG analyses explicitly seek to differentiate neural responses that scale in a non-bifurcated manner with SOA duration (i.e., 17ms, 33ms, 50ms, 67ms, and 83ms) from those that bifurcate with perceptual awareness (i.e., behavioral responses in Figure 1B).

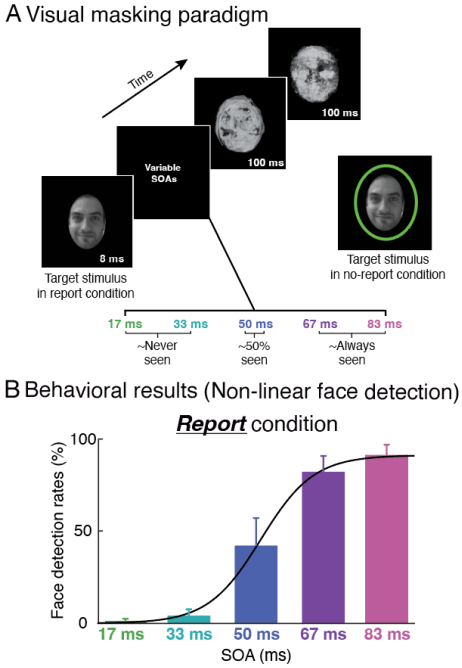

Figure 1. Behavioral paradigms and results. A) Design of visual masking paradigm. The target face (or blank) was presented for 8ms, followed by a variable SOA, and then two successive masks, each presented for 100ms. In the report condition, subjects were tasked with detecting faces on each trial (Y/N). In the no-report condition, the subjects' task was to detect occasional green ring targets randomly presented throughout the stream (i.e., sometimes superimposed on a face and sometimes shown when no face is present). B) Report condition behavioral results. Face detection rates are presented on the y-axis and the different SOAs are presented on the x-axis. Error bars represent the standard error of the mean and the five data points are fit with a non-linear regression.

Distinguishing non-bifurcated from bifurcated EEG signals

Which EEG signals increase in a non-bifurcated manner? In this case, we found that the amplitude of the P1 matched the linear increase in target-mask timing across the five SOAs (i.e., 17ms, 33ms, 50ms, 67ms, and 83ms, Figure 2; see Supplemental Results, and preregistration materials (osf.io/6d5av) for information on pilot testing that identified these electrodes and the time windows of analysis for the P1, P3b, and N2 ERPs). In the no-report condition (Figure 2A), there was a significant increase in the size of the P1 for each successive SOA pairing (17 vs. 33ms, $t(39)=4.24$; $P<0.001$; 33 vs. 50ms, $t(39)=7.86$; $P<0.001$; 50 vs. 67ms, $t(39)=2.10$; $P<0.05$; 67 vs. 83, $t(39)=2.08$; $P<0.05$) . A similar pattern

was found in the report condition (Figure 2B): P1 amplitude was significantly different across all four SOA pairings (17 vs. 33ms, $t(39)=6.30$; $P<0.001$; 33 vs. 50ms, $t(39)=7.76$; $P<0.001$; 50 vs. 67ms, $t(39)=3.06$; $P<0.05$; 67 vs. 83, $t(39)=1.95$; $P<0.05$). This pattern of linear P1 amplitude change across SOA did not change as a function of the act of reporting: compare Figures 2A and B (interaction between no-report and report: $F(4,390)=0.41$, $P=0.80$).

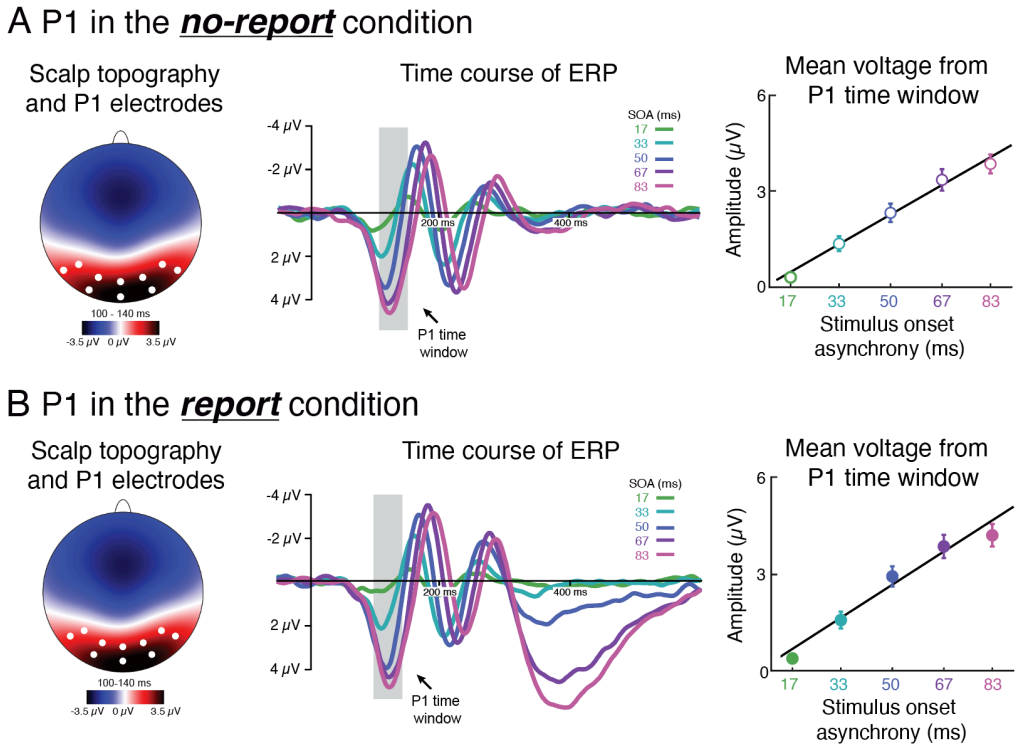

Figure 2. Visualizations of the P1 in the no-report and report conditions. A) No-report and B) report conditions. Left: Topographical voltage distributions between 100-140ms with all electrodes used to analyze the P1 indicated on the scalp map. Center: P1 waveforms for all five SOAs plotted over time. Right: Average amplitudes from electrodes used for the P1 across all five SOAs. Error bars represent the standard error of the mean and the five data points are fit with a linear regression.

Distinguishing bifurcated EEG signals associated with perception from post-perceptual processing

To identify potential correlates of perceptual consciousness independent of report, we then looked for neural responses in the no-report condition that matched participants' bifurcated behavioral responses in the report condition. For a neural signal to bifurcate, it must meet four criteria that we pre-registered in advance: 1) little or no difference between the 17ms and 33ms SOAs, 2) a significant difference between the 33ms and 50ms SOAs, 3) a significant difference between the 50ms and 67ms SOAs, and 4) little or no difference between the 50ms and 67ms SOAs. If any of these benchmarks are not met for a given neural signal in the no-report condition, we would argue that signal does not properly mimic the psychometric function from the behavioral results.

The first candidate signal we examined was the P3b. The P3b has been associated with conscious perception across numerous studies (1, 14, 32, 34, 42-44), and a previous study using the same type of linear/nonlinear logic used here linked perceptual awareness with

the P3b under conditions of report (see Figure 1; 34). We sought to replicate this result in the report condition and examine whether it would hold in the no-report condition.

The results from the P3b analysis (see Methods) are shown in Figure 3. In the report condition, we observed signs of non-linear bifurcation in the P3b, replicating prior work (34). There was no difference in P3b amplitude for the two shortest SOAs (17 vs 33ms, $t(39)=0.62$; $P=0.26$), along with significant increases in amplitude surrounding the middle SOA (33 vs. 50ms and 50 vs 67ms; $t(39)>7.09$; $P<0.001$), consistent with bifurcation dynamics. However, we also found a significant increase in P3b amplitude between the two longest SOAs (67 vs 83ms; ($t(39)=8.01$; $P<0.001$). Most importantly, we found no evidence of a P3b in the no-report condition, during which mean amplitudes never rose significantly above zero at any SOA ($t(39)=1.88$; $P>0.07$ in all cases; Figure 3A; interaction between no-report and report P3b amplitudes: $F(4,390)=226.83$, $P<0.001$). The complete disappearance of these effects in the no-report condition reinforces findings from prior work showing that the P3b is associated with task performance, such as reporting a perceptual experience, rather than with the experience itself (22, 27, 28, 29, 31).

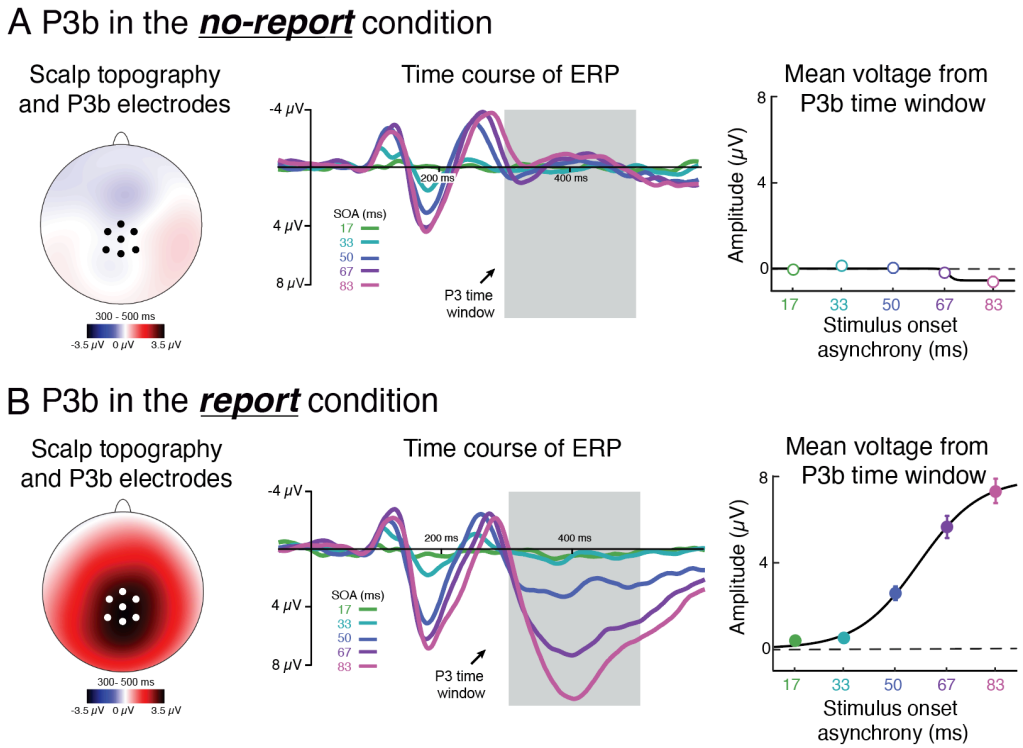

Figure 3. Visualizations of the P3b in the no-report and report conditions. A) No-report and B) report conditions. Left: Topographical voltage distributions between 300-500ms with all electrodes used to analyze the P3b indicated on the scalp map. Center: P3b waveforms for all five SOAs plotted over time. Right are the average amplitudes from all electrodes used for the P3b across all five SOAs. Error bars represent the standard error of the mean and the five data points are fit with a non-linear regression.

N170/VAN

After ruling out the P1 and P3b, we looked next at negative-going waveforms between the P1 and P3b time-windows, which have been previously proposed as potential neural correlates of perceptual awareness (40). For example, the “visual awareness negativity” (VAN) has been defined as a relatively more negative ERP signal for seen compared to

unseen stimuli, with a bilateral posterior distribution, roughly in the time window between 150-300ms (32, 40). When faces are used as stimuli, it is difficult to distinguish the VAN from the N170, and several studies have proposed the latter as a content-specific correlate of conscious face perception (28,45-51). In the current analysis, we refer to this effect as an N170/VAN, and later discuss how future studies might be designed to better distinguish between these highly similar neural signals. Results from the N170/VAN analysis are shown in Figure 4. In both conditions, the N170/VAN followed a pattern of amplitude increase that was neither linear nor bifurcated. In the report condition (Figure 4B), N170/VAN amplitudes increased significantly across the first three SOA pairings (17 vs. 33ms, $t(39)=4.79$; $P<0.001$; 33 vs. 50ms, $t(39)=8.61$; $P<0.001$; 50 vs. 67ms, $t(39)=3.42$; $P<0.001$) and plateaued at the two longest SOAs (67 vs. 83ms; $t(39)=1.02$; $P>0.05$). Similarly, in the no-report condition (Figure 4A), the N170/VAN grew significantly across SOA until SOA 67 (17 vs. 33ms, $t(39)=5.60$; $P<0.001$; 33 vs. 50ms, $t(39)=4.01$; $P<0.001$; 50 vs. 67ms, $t(39)=2.40$; $P<0.05$), with no further increase in amplitude from SOA 67 to 83 ($t(39)=-0.43$; $P>0.05$). While these patterns of N170/VAN amplitude modulation across SOA came close to matching the sigmoid-shaped psychometric function in the behavioral data from the report condition and remained present in the no-report condition, they did not meet the pre-registered criteria for bifurcation. The presence of the N170/VAN for the 33ms SOAs suggests that it may not be directly linked with perceptual awareness since subjects almost never perceived the face stimuli at the 33ms SOA. This finding is in line with previous work reporting an N170 in response to unseen faces (52-54), and we discuss the implications for the VAN below.

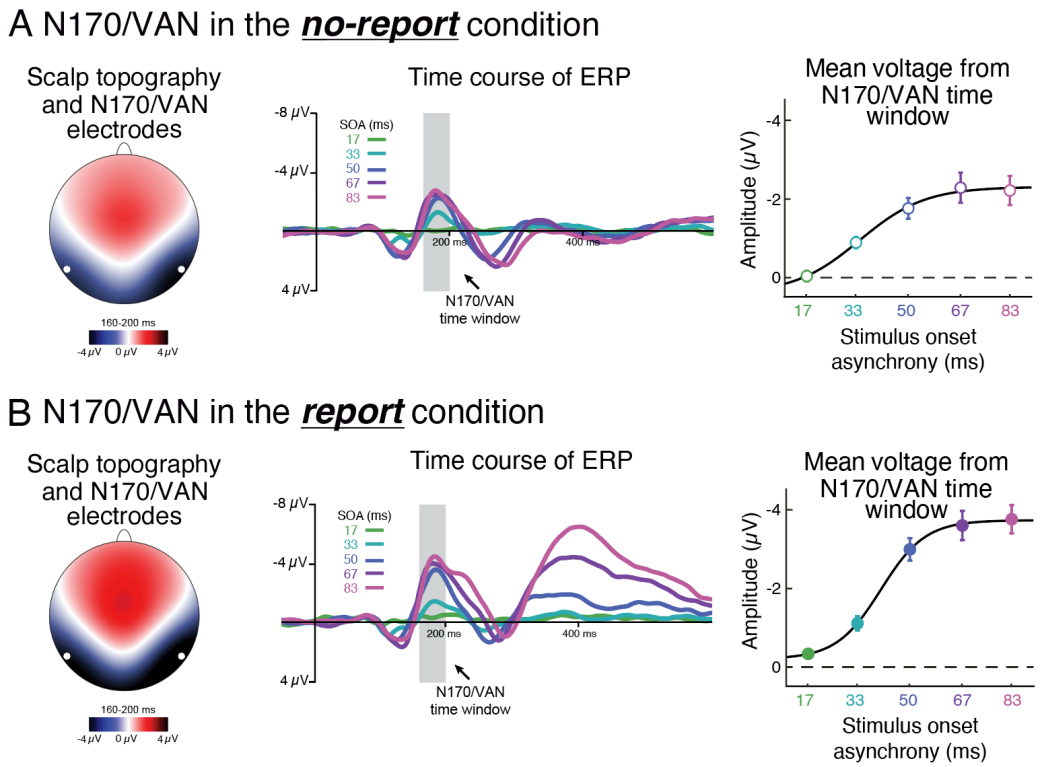

Figure 4. Visualizations of the N170/VAN in the no-report and report conditions. A) No-report and B) report conditions. Left: Topographical voltage distributions between 160-200ms with the electrodes used to analyze the N170/VAN indicated on the scalp map. Center: N170/VAN waveforms for all five SOAs plotted over time. Right: average amplitudes from

electrodes used for the N170/VAN analysis across all five SOAs. Error bars represent the standard error of the mean and the five data points are fit with a non-linear regression.

A novel neural signature of perceptual awareness: The N2

In the time-period between the N170/VAN and P3b, one ERP signal closely matched the nonlinear, bifurcated pattern of participants' perceptual reports: a negativity distributed over fronto-central electrodes which we refer to as the N2 (Figure 5A). Between 250-290ms, this N2 wave in the no-report condition displayed a clear non-linear, bifurcated pattern of amplitude increase that mirrored psychometric functions in the report condition (Figures 4A center and right panels). There was no difference in the amplitude of the N2 between the two shortest SOAs: 17ms and 33ms ($t(39)=0.29$; $P=0.39$). There were significant differences in N2 amplitudes between 33ms and 50ms ($t(39)=5.83$; $P<0.001$), as well as between 50ms and 67ms ($t(39)=4.44$; $P<0.001$). Finally, N2 amplitudes at the two longest SOAs did not differ significantly: 67ms and 83ms ($t(39)=0.68$; $P=0.25$).

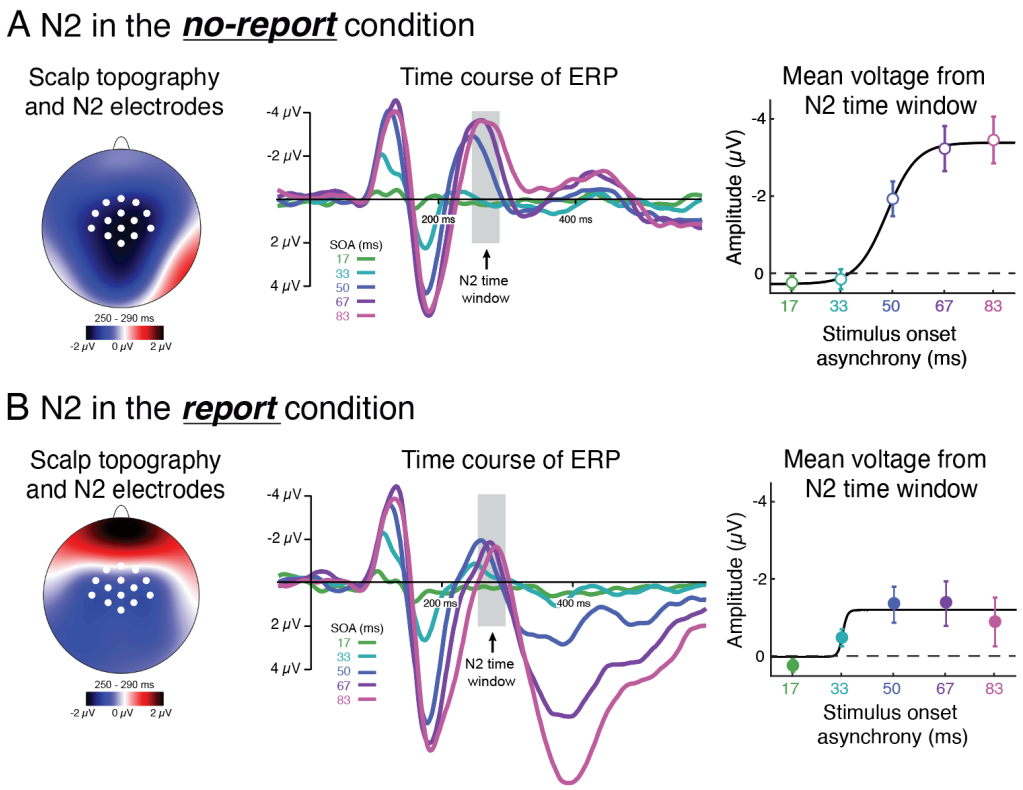

Figure 5. Visualizations of the N2 in the no-report and report conditions. A) No-report and B) report conditions. Left: Topographical voltage distributions between 250-290ms with all electrodes used to analyze the N2 indicated on the scalp map. Center: N2 waveforms for all five SOAs plotted over time. Right: Average amplitudes across all electrodes used for the N2 across all five SOAs. Error bars represent the standard error of the mean and the five data points are fit with a non-linear regression.

It should be emphasized that we measured the N2 in neural responses from the no-report condition, while the perceptual reports it mirrored came from the report condition. As such, the N2 cannot reflect the neural processes underlying the non-linear bifurcation in participants' perceptual judgments or behavioral responses (i.e., post-perceptual bifurcation dynamics). Instead, the N2 is more likely associated with perceptual awareness per se: in

the no-report condition, the stimuli in question were task-irrelevant but otherwise identical to those in the report condition, such that the influence of post-perceptual processing was limited while the stimuli were presumably perceived at similar rates across SOAs compared to the report condition. Of course, we recognize it cannot be definitively stated that awareness of the face stimuli in the report and no-report conditions was identical. This is an inherent limitation of no-report paradigms that all studies must consider. Here, we specifically designed the no-report task (detection of green rings superimposed on the faces) to minimize differences in attention, effort, etc. across tasks. Nevertheless, we openly acknowledge that we cannot state with certainty that the rates of face detection were identical in the report and no-report conditions and instead can only make that assumption.

ERP analyses on the same set of electrodes during the report condition revealed a different pattern of results (Figure 5B). The N2 was partially evident, but its pattern of amplitude change across SOAs was clearly affected by the act of reporting the presence of a target face. As a result, the topographical voltage distributions looked quite different between the report and no-report conditions, and none of the bifurcating characteristics of the N2 observed in no-report were noticeably present in the report condition (interaction between no-report and report: $F(4,390)=34.88$, $P<0.01$). There was a significant difference in amplitude between SOAs 17 vs. 33ms ($t(39)=3.17$; $P<0.01$) and no difference between SOAs 50 vs. 67ms ($t(39)=0.09$; $P=0.47$), neither of which match the trends expected from a non-linear, bifurcated signal. In fact, when comparing the report and no-report conditions directly, there was a significant interaction between the 17 vs. 33ms SOA, the 33 vs. 50ms SOA, and the 50 vs 67ms SOA ($t(39)>2.08$; $P<0.05$ in all cases). This suggests that prior studies may have failed to detect a link between perceptual awareness and the N2 because observers always had to make explicit reports about the contents (or presence) of their perceptual experience.

If the N2 is a true potential signature of perceptual awareness, why was it not found in the report condition? We argue that in the report condition, the spatially and temporally overlapping P3b wave “drowned out” the N2 and prevented it from being detected. To test this possibility, we performed a post-hoc exploratory analysis. In this analysis, we split our participants into two groups based on the strength of the P3b (i.e., largest P3b group and smallest P3b group; N=20 in each group). We predicted that we would see more “hints” of the N2 amongst the smallest P3b group since the P3b would interfere with the N2 less in that group. Overall, we found no hints of the N2 amongst the largest P3b group (Figure 6A) but did find some hints of the N2 amongst the smallest P3b group (Figure 6B). While not definitive, this analysis suggests that interference from the P3b in the report condition prevented us from finding an N2 in that condition. Moreover, we believe this finding highlights the importance of no-report paradigms since the neural processes most associated with perceptual awareness may simply go unnoticed due to interference from the act of report.

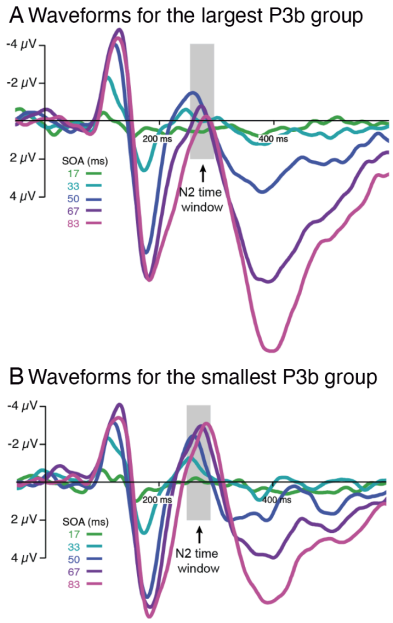

Figure 6. N2 waveforms in the report condition from two groups of participants sorted by P3b peak amplitude (median split, N=20 in each group). A) Waveforms for the largest P3b group and B) waveforms for the smallest P3b group.

Late metastability via temporal generalization

Having found the N2 to be the only ERP that clearly follows a bifurcated response pattern that reflects perception, we proceeded to ask if any other EEG signals might display these non-linear dynamics in the no-report condition. To examine this possibility, we focused on another posited signature of conscious processing, the presence of late metastable activity (4,38,42,55-57). Under this proposal, as information processing transitions from unconscious sensory stages to conscious perception, it is stabilized for a few hundred milliseconds so that conscious content may be used to achieve specific goals (i.e., decision-making, verbal responses, motor outputs, etc.). This metastability can be represented in temporal generalization matrices, which visualize how different processing stages and their underlying neural codes unfold over time (37). Specifically, these matrices are formed by training pattern classifiers to distinguish between two or more stimulus conditions at each time-point and then testing classification performance across all time-points. This process returns a temporal generalization matrix that visualizes the decoding accuracy for every pair of training and testing time points. Several studies have found that perceptual awareness is linked with the transformation of information into a metastable representation (14,34,35,57,58). This late metastability is reflected by an off-diagonal square-shaped pattern in the temporal generalization matrix (black square in right panel of Figure 7).

To look for evidence of a non-linear bifurcation in this late temporal generalization signal, we trained multivariate decoders to distinguish between trials in which a face was present (stimulus trials) versus absent (mask only trials; see Figure 1) and visualized the results in a standard temporal generalization matrix (see Methods). In the report condition, we found strong evidence of a metastable representation associated with conscious perception, indicated by late square-shaped patterns in the temporal generalization matrix (Figure 7B).

Indeed, at the 50ms SOA, in which observers first started to sometimes notice the target face (see Figure 1B), a late square-shaped pattern emerged that increased in power at longer SOAs, when stimuli were reliably visible. However, in the no-report condition, we found no evidence for a late, sustained, metastable representation: decoding accuracy was only significantly above chance when trained and tested on similar time points (Figure 7A). Together, these results suggest that late, sustained metastability as represented by temporal generalization matrices is more associated with post-perceptual processing (i.e., reporting tasks) than perceptual awareness per se.

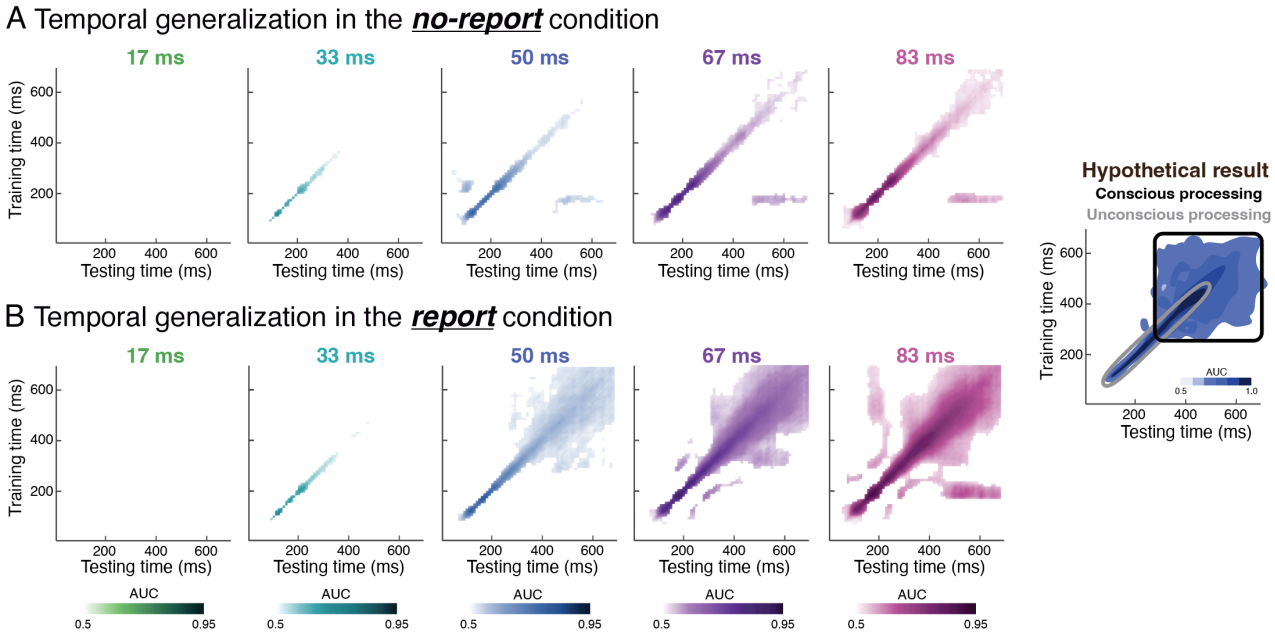

Figure 7. Temporal generalization matrices in no-report and report conditions. A) No-report and B) report conditions. A decoder was trained to discriminate target-present and target-absent trials between 0 and 700ms. In the no-report condition (A), decoding was successful only along the temporal diagonal, suggesting a sequential activation of different brain regions. However, in the report condition (B), while the 33ms SOA only allowed for weak decoding along the diagonal, the longer SOAs elicited a combination of diagonal decoding and late sustained generalization that begins roughly 300ms after stimulus onset. The right panel represents a hypothetical result interpreted under the proposal that temporal generalization reflects conscious perception, depicting the diagonal associated with unconscious processing and the square associated with conscious processing. The current results challenge this proposal.

Inter-trial EEG/ERP variability

Another recently proposed signature of conscious processing independent of report is a spike in inter-trial EEG/ERP variability for stimuli presented at threshold (35). The rationale for this measure revolves around the idea that a threshold stimulus will yield a mixture of high-activity (aware) and low-activity (unaware) trials. This burst in inter-trial variability will not be seen for stimuli that are well above or well below threshold since neural activity should be more homogenous across such trials. Thus, this proposal predicts that we should see an increase in inter-trial variability in the 50ms SOA condition, when the target stimulus is at threshold, along with lower variability at the other SOAs which were well below or well above threshold. However, we found minimal evidence of such variability in the no-report condition and no evidence for this variability pattern in the report condition (Figure 8; see Methods). In the no-report condition, we observed some signs of a peak of inter-trial EEG/ERP variability at the threshold SOA (50ms) during the 250-300ms and 300-350ms time windows,

overlapping the time window of the N2 wave (SOAs 33 vs. 50ms: 250-300 ms, $t(39)=1.76$; $p < 0.05$; 300-350 ms, $t(39)=1.68$; $p=0.05$), but the difference in variability on SOA 50 vs. 67 trials did not reach significance in either time window ($t(39)<0.56$; $p>0.32$). Inconsistent with previous reports (35), no such pattern was evident during any time window in the report condition (see Supplemental Figure 5 for results from exploratory analyses at different electrode clusters).

Inter-trial EEG variability, centro-parietal cluster

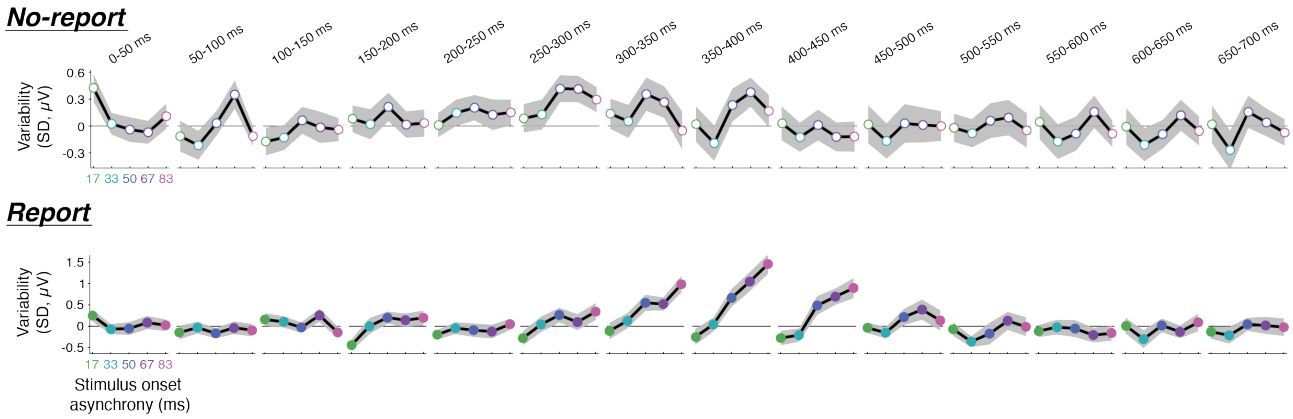

Figure 8. Inter-trial EEG variability in no-report and report conditions at centro-parietal electrodes. Top row corresponds to the no-report condition and the bottom row corresponds to the report condition. On the horizontal axes are the five SOAs (ms) and on the vertical axes is the inter-trial variability of evoked activity as a function of SNR at these different time windows. Note that the electrode cluster used in this analysis is the same one used for the P3b in the present study, which is highly similar to the group of electrodes used by Sergent et al. (35) in their previous variability analysis.

Discussion

In the present study, we sought to identify neural signatures uniquely associated with perceptual awareness in a no-report visual masking paradigm. Overall, we found a novel EEG signal in the no-report condition — the N2 — that displayed bifurcation dynamics (Figure 5A) that closely matched the psychometric function of behavioral responses in the report condition (Figure 1B). The close correspondence between N2 amplitude patterns and reported perception across the five levels of stimulus visibility suggests that the N2 reflects perceptual awareness in the absence of report. This particular component stood in contrast to the P1, which increased linearly regardless of the task condition (Figure 2), likely reflecting the linear increase in the SOAs between the target and masks. We were also able to rule-out two other candidate signatures of conscious processing: the P3b and late temporal generalization of decoding (38,41). Although we found evidence for non-linear, bifurcation dynamics in the report condition for both the P3b (Figure 3B) and late temporal generalization (Figure 6B), both of these neural signals entirely disappeared in the no report condition (Figure 3A and 7A) suggesting an association with post-perceptual processing, not perceptual awareness. Meanwhile, in both the report and no-report conditions, the N170/VAN was somewhat similar to the sigmoid-shaped psychometric function. However, it did not fully meet the criteria for bifurcation, unlike the N2. Finally, although the inter-trial EEG variability analysis showed hints of a peak in variability around the threshold SOA at certain time-points in the no-report condition, this result was inconclusive and did not clearly serve as a signature of conscious processing.

The N2 and other proposed neural signatures of consciousness

The utility of experiments aimed at isolating spatio-temporal brain signals linked with conscious perception lies in their ability to elucidate stages of processing and critical network dynamics that will inform our broader understanding of the neural basis of consciousness. Contemporary consciousness researchers generally agree that there is not one time or location in the brain in which consciousness "arises" from unconscious or preconscious processing (18,40,59-61). Indeed, the goal of the present study and many previous studies in the literature is not to identify a single critical time-point at which conscious perception "happens" in the brain (62), but rather to isolate the key neural patterns and interactions between neural systems that are most closely associated with our perceptual experience. In other words, the key neural signal we identified in this experiment — the N2 — needs to be considered in terms of its relationship to other neural and psychological events in the complex chain of processing that occurs from sensation to cognition and action.

One particularly noteworthy aspect of the N2 is its temporal relationship to both the P3b and late temporal generalization. Even though both of those neural patterns appear to be associated with post-perceptual processing, not perceptual awareness, the N2 appears to end right as the P3b and late temporal generalization begin. As can be seen in Figure 5A, the N2 terminates at around 300ms, which is precisely when the P3b becomes positive (Figure 3B) and late temporal generalization begins (Figure 6B). As noted in the results section, the initial phase of the P3b appears to overlap the N2 in the report condition, reducing its amplitude and destroying its bifurcation pattern (Supplemental Figure 6). Thus, an intriguing possibility is that the N2 is truly associated with perceptual awareness per se (i.e., independent of report) and in situations where these perceptual experiences are relevant to the task-at-hand, this information is then stabilized and fed into other networks associated with maintaining and using those perceptual experiences to guide behavior.

Recently, Sergent et al. (35) provided evidence suggesting that inter-trial EEG/ERP variability may be a signature of conscious processing independent of report. One of their main findings was a distinct burst of inter-trial variability that emerged in their passive (no-report) condition for stimuli presented around the threshold of conscious perception (i.e., the 50ms SOA, in the present study). This particular variability showed non-linear bifurcation dynamics that matched the perceptual reports in the report condition. Strikingly, the time window in which this bifurcated EEG/ERP variability emerged was almost identical to the time window of the N2 bifurcation in the present study (250-300ms). Therefore, even though the Sergent et al. (35) study and the present study examined different sensory modalities (i.e., auditory vs. visual), manipulated perceptual awareness in different ways (i.e., variable stimulus intensity vs. backwards masking at variable SOAs), and used different measurements (i.e., inter-trial EEG/ERP variability vs. ERP amplitudes), the findings from both studies converge about the timeframe of neural dynamics linked with perceptual awareness. Together, these results form the foundation of a new hypothesis that the true neural mechanisms supporting conscious perception begin roughly 200-300ms after stimulus onset, during a brief stage of processing in-between low-level sensory encoding and high-level broadcasting across more global neuronal networks.

While the current study was unable to replicate the finding of enhanced inter-trial EEG/ERP variability at perceptual threshold reported by Sergent et al. (35), we observed

some signs of such a pattern between 250-350ms in the no-report condition. While it is unclear why there was no trace of this pattern in the report condition, several viable explanations exist, including the different stimulus modality (vision vs. audition) and the different method used to render stimuli invisible (backwards masking vs. weak stimulus energy embedded in noise). In line with the latter explanation, the current results appeared to show a systematic increase in inter-trial EEG/ERP variability during most of the P3b time window (300-450ms), which was not observed in Sergent et al.'s (35) data. Such an increase in variability that continues beyond the threshold SOA could be explained by the additional cognitive demands of perceiving two competing stimuli (stimulus and mask) at longer SOAs compared to perceiving only one stimulus (the mask) at shorter SOAs. Such differential cognitive demands across SOAs would only be expected to come into play in conditions that require trial-by-trial reports. Ongoing experiments in our group are aimed at exploring this interpretation by measuring inter-trial EEG/ERP variability across different levels of stimulus visibility without using backward masks.

Another previously proposed neural signature of visual consciousness is the visual awareness negativity, or VAN (32,39,40,63-65). While we were able to analyze two different negative-going waveforms in the VAN time-window, ~150-300ms (N170 and N2), it is unclear whether either of these two neural signals should be considered a variant of the VAN. The N170 showed a similar bi-lateral posterior distribution to the VAN, and was larger for seen versus unseen stimuli. However, as mentioned above, the presence of a significant N170 in the 33ms SOA condition could be considered evidence against its association with visual awareness, as stimuli presented at this SOA were hardly ever consciously perceived. The N2 appears distinct from the VAN for several reasons. First, the electrodes associated with the N2 (i.e., fronto-central) are more anterior than those associated with the VAN (i.e., occipito-parietal). Second, in most cases, the VAN appears earlier than the N2, beginning as early as 150ms in some studies (32; but see 28), with the N2 not emerging until around 250ms. Finally, some researchers have proposed that the VAN might be more associated with attention than awareness (66), and indeed, it is likely that a key factor that renders a threshold stimulus perceived vs. not-perceived is trial-by-trial fluctuations in attention (67-69). If the N170 measured here is indeed a content-specific variant of the VAN, the current results support this notion, i.e., this face-specific signal itself may be elicited regardless of awareness (e.g., being present in the 33ms SOA condition) while also being modulated by attention, thus increasing in amplitude for the 50ms and 67ms conditions in which the stimuli were consciously perceived at increasing rates. Consistent with this interpretation, the N170/VAN amplitudes were larger in the report compared to the no-report condition, a possible task-based attention modulation. An intriguing possibility, therefore, is that the VAN is more linked with a type of selective attention that is a necessary pre-cursor for visual awareness, while the N2 reflects the neural processes underlying perceptual awareness per se. Future studies will be needed to evaluate this emerging hypothesis.

The N2 and theories of consciousness

How do these findings relate to the longstanding debate between "cognitive" and "sensory" theories of consciousness? On one side of the debate, cognitive theories like the Global Neuronal Workspace or Higher-Order Thought Theory claim that information must reach a fronto-parietal network in order to be consciously experienced, with a specific focus on certain regions within the prefrontal cortex (7,10,11,12,38,41). These particular theories

have generally been more closely linked with electrophysiological signatures such as the P3b and late patterns of temporal generalization of decoding (i.e., 300-600ms post-stimulus). On the other side of this debate, more sensory-leaning theories like Recurrent Processing and Integrated Information Theory maintain that the neural regions supporting conscious processing are located in more posterior, sensory areas (16-20). This set of theories has primarily cited support from electrophysiological signatures much earlier in time, including the VAN.

The current results, however, do not fit neatly into either of these popular frameworks. In terms of both time and space, the N2 appears to index a stage of processing in between the N170/VAN and the P3b. Two hallmarks of the P3b are its sustained response (i.e., several hundreds of milliseconds) and its broad spatial distributions (i.e., covering nearly the entire scalp). These features stand in contrast to the N2, which is much more temporally truncated (lasting roughly 50-100ms) and spatially circumscribed (with bilateral posterior/inferior positive foci distinct from the P3b, potentially indicating the opposite end of the N2 dipoles). Meanwhile, compared to the VAN, the N2 scalp topography is more anterior, and its timing is later (250ms to 300ms for the N2 and anywhere between 150 to 300ms for the VAN). Interestingly, the N2 reported by Sergent et al. (35) in the no-report condition of their auditory threshold experiment had a very similar timing as the N2 recorded here for visual masking. Thus, the N2 appears to be situated in an intermediate stage of processing between the neural markers most often cited as supporting these two classes of theories, with neither framework being obviously supported by the dynamics of the N2. Several alternative theories, particularly those that posit a key role for interactions between attention and perception, are consistent with the current N2 results (70-75).

Identifying where the N2 originates in the brain will further elucidate how it fits into these theoretical frameworks. Unfortunately, we cannot definitively establish its source using EEG alone and would need a method with better spatial resolution such as functional magnetic resonance imaging (fMRI) or intracranial electrocorticography (ECoG). However, a few possible inferences can be made from the present set of results. Given its location in space and time and its relation to other established ERPs, one possibility is that the N2 originates primarily from within parietal regions (or interactions between parietal and occipito-temporal regions). Of course, if this proved to be the case, such a result would not clearly support one of the dominant competing theories over the other, since both can integrate such a finding into their framework. For example, cognitive theorists might see parietal activation as part of a “parietal workspace”, or a so-called “global playground” in which information is made conscious by activating the regions of the parietal lobe associated with higher-order functions such as attention, working memory, and decision-making (14,76-80). Conversely, sensory theories might relate this activity to the previously described posterior cortical “hot zone” that is separate from regions involved in higher-level functions (18,19). Nevertheless, despite the challenges involved in relating such findings to different theories, understanding the regions from which the N2 originates will be an important step in the process of triangulating the regions containing the key neural networks that truly support perceptual awareness independent of reports.

The N2 and post-perceptual processing

It is worth emphasizing that while the no-report condition was created to limit post-perceptual processes, we cannot say with certainty that such post-perceptual processes were entirely eliminated. Indeed, it is quite possible that with minimal task demands in the no-report condition, participants may spontaneously attend to the faces, remember seeing certain faces, or think about the faces (81). Indeed, prior studies have suggested that the N2 is directly involved in the detection of novelty and attention orienting (82). Thus, the N2 may reflect certain post-perceptual processes that are present even in no-report paradigms. This possibility is an inherent limitation of no-report paradigms and is in no way unique to our methodology. However, the fact that only the N2 followed the same bifurcation dynamics as behavior in the report condition suggests that it may be intimately linked with conscious processing and warrants further examination, e.g., with different types of stimuli and/or different manipulations of awareness other than backward masking.

Conclusion

By studying EEG bifurcation dynamics in a no-report paradigm, we identified a new ERP wave that may be uniquely associated with perceptual experience, independent from the act of reporting such experiences and decoupled from stimulus-based differences often associated with visible and invisible stimuli. Critically, the characteristics of this specific ERP, which we call the N2, do not neatly align with either sensory or cognitive theories. While the N2 appears too late in time and is too anterior in its spatial location to naturally support sensory theories, it is also too early in time and too short-lived to be readily accommodated by cognitive theories as they currently exist. Indeed, these findings challenge multiple prominent theories of conscious awareness and lend support to a more intermediate view of perceptual experience in the brain. Going forward, it will be imperative to continue using no-report paradigms and bifurcation dynamics to identify the precise regions involved in perception and further investigate the stages of conscious processing.

Methods

Pre-registration

The methods and planned analyses were pre-registered and available on the Open Science Framework: https://osf.io/6d5av/. Time windows and electrodes for the ERP analyses of the P1, N2, and P3b were established based on data from an initial set of 15 participants and included in the pre-registration before the remaining data acquisition began (see ERP Analysis, below).

Participants

56 participants between the ages of 18-29 participated in the experiment. Data from 8 participants were unusable due to technical issues during EEG recording, and 8 participants were excluded for failing to meet the established inclusion criteria (see Exclusion Criteria, below), resulting in a total of 40 participants in the final sample. All participants had normal or corrected-to-normal vision, no known neurological conditions, and no concussions in the preceding year. The experimental procedure was approved by the Institutional Review Board at Reed College and informed consent was obtained from each participant before the experiment, who were compensated for their time.

Stimuli

Stimuli consisted of 40 grayscale face images (120 x 140 px; 2.7° x 3.15°) and 40 masks embedded in a black background adapted from Kouider et al. (34). Masks were composed of scrambled, inverted faces and oval objects, and were similar to the face stimuli in size and shape. The contrast level of the face stimuli was

adjusted to match each participant's perceptual threshold as determined by a staircase procedure completed prior to beginning the main experiment (see below).

The green rings (RGB 0, 255, 0) that served as targets in the no-report condition were 15 px wide with a horizontal diameter of 120 px and a vertical diameter of 140 px, measured from outer edge to outer edge. These dimensions ensured each face was fully encircled and minimized the negative space between the stimuli and ring without occluding facial features.

Stimuli were presented centrally against a black background at 120 Hz on a 1920 x 1080 px BenQ monitor, positioned at eye level approximately 70 cm in front of the seated participant. A small red fixation dot (RGB 255, 0, 0; 5 x 5 px; 0.23°) was centered on the screen throughout trial presentation. Neurobehavioral Systems' Presentation 19.0 was used to present the stimuli in the main experiment, and stimuli were controlled using Psychtoolbox 3.0.14 (83) for MATLAB (v. 2017a) during the thresholding procedure.

Experimental design

Each participant first completed a staircase procedure to determine their perceptual threshold (see below), followed by a brief demonstration of the stimuli, and then the main experiment. The primary experimental procedure consisted of 32 blocks of 60 trials, divided into two 16-block sections with one section being the report condition and another being the no-report condition. All participants completed both the report and the no-report conditions and the order of the conditions was counterbalanced across participants. Participants were given written instructions on the screen at the start of each condition.

Each block began with the appearance of a central red fixation dot that remained on-screen throughout the block. Participants were instructed to keep their eyes on the dot and avoid eye movements and excessive blinking. On trials in which a face target was present, the a face stimulus was presented in the middle of the screen for 8.33ms, then followed by a blank screen for a variable delay, then followed by two consecutive masks shown for 100ms, and then finally a randomly jittered inter-trial interval of 1050-1450ms, which served as a response window in the report condition.

On each trial, the stimulus-onset asynchrony (SOA) between the face stimulus and the first mask could be 16.67ms, 33.33ms, 50ms, 66.67ms, or 83.33ms (SOAs '17,' '33,' '50,' '67,' and '83,' respectively). The SOA on any given trial was selected randomly, but all SOAs appeared the same number of times in a block. These SOA durations were chosen after pilot testing demonstrated that, when stimulus contrast was adjusted according to each participants' perceptual threshold so that they could detect the stimulus half of the time on SOA 50 trials, participants were almost never able to see the stimuli on SOA 17 and 33 trials, and almost always saw them on SOA 67 and 83 trials. Linearly manipulating the SOA duration across these five intervals led to a clear sigmoidal psychophysical function, with minimal differences in visibility between the two shortest SOAs (17 and 33ms) or the two longest (67 and 83ms), and maximal differences around the threshold SOA (33 vs. 50ms and 50 vs. 67ms).

On 17% of trials, no face stimulus was presented at all. This provided a set of 'mask-only' trials that were used in analysis to isolate stimulus-related from mask-related brain activity (see ERP Analysis, below). Independently, on 15% of trials a green elliptical ring appeared with the same timing as the face stimulus, centrally located and sized so as to encircle the same area where the face stimuli were presented ('green ring' trials). These were used as targets in the no-report condition and excluded from all further analyses.

In the report condition, participants indicated via button press after every trial if they had or had not seen a face. These results were used to confirm that the SOA manipulation successfully led to bifurcated perceptual reports, and as a replication of Del Cul et al. (33).

For the no-report portion of the experiment, participants were instructed to press a button whenever they saw a green ring that randomly appeared in roughly the same location as where the face stimuli would be shown. Importantly, the green ring did not only appear superimposed around a face; rather, it was presented at the same rate (15%) on trials when no face was shown. As a result, the face stimuli were spatially and temporally attended but were task-irrelevant and were not followed by a response. The task was designed to be challenging enough to keep the participant engaged but not so demanding that they might become intentionally blind to the task-irrelevant face stimuli.

After each 60-trial block, participants received feedback on their task performance. In the report condition, they were given the percentage of all faces they had detected and their false alarm rate (i.e., if they had any “seen” responses to a mask-only trial). Since participants were not supposed to see any of the faces at the shortest SOAs (and only half the trials at SOA 50) but those trials still counted as “misses” for the purposes of scoring, participants were informed that the expected detection rate was around 50%. This was so participants did not become frustrated with the task due to perceived low-performance. In other words, if participants believed they were expected to see 100% of the faces but were unable to score above 50%, they were more likely to feel discouraged and lose engagement with the task or alternatively adjust their response criterion which would increase the rate of false alarms.

Feedback in the no-report condition consisted of a count of the number of rings they had missed and it was emphasized to the participants that they should not be missing any of the rings. However, in pilot testing, we found that some participants were unable to see the green rings at all on SOA 17 trials, thus these most difficult trials were not counted as misses. Participants were informed of this at the start.

The experiment began with a staircase procedure manipulating the contrast of the face stimuli to establish each participant’s perceptual threshold (QUEST: 84). The trial design and task were almost identical to the report condition of the main experiment except all trials used a stimulus-mask SOA of 50ms and no green rings were presented. Additionally, the program paused after each trial to wait for the participant’s response. The results of this procedure were used to set the contrast of the stimuli for the rest of the experiment.

This was followed by a short ‘training’ session to further familiarize participants with the stimuli and show them the green rings. Participants were shown 25 trials at each SOA, in order from longest to shortest. The last five trials in the sets were always green ring trials, and participants were given the option of re-watching each set before going on to the next.

EEG acquisition

Scalp EEG was recorded using custom 64-channel electrode caps with an equidistant M10 layout (EasyCap, Herrsching, Germany). Signals were amplified via two 32-channel Brain Amp Standard amplifiers (Brain Products, Gilching, Germany), online filtered from 0.1 to 150 Hz, and digitized at 500 Hz. All channels were referenced to an electrode on the right mastoid during recording and electrode impedance was kept under 10 kΩ. Two horizontal EOG channels positioned laterally to each eye and one vertical EOG channel under the left eye were used to monitor blinks and eye movements.

Exclusion

Participants were excluded from analysis if over 50% of trials were rejected due to EEG artifacts or if their false alarm rate was higher than 20%. Six participants were excluded due to excessive artifacts and two were excluded for false alarms.

Data preprocessing

EEG preprocessing was performed in BrainVision Analyzer 2.2 following the steps below:

1. Visual identification of any noisy electrodes and interpolation of surrounding electrodes to reconstruct the noisy channel, if necessary.
2. Re-referencing to the average of the left and right mastoid channels.
3. Bipolar HEOG and VEOG channels constructed by re-referencing the left and right horizontal EOG channels and the vertical EOG channel and electrode 50 (Fp1) as bipolar pairs.
4. Segmentation of trials around the stimulus (or the first mask in mask-only trials) and removal of all green-ring trials.
5. Baseline correction using the time window from -200 to 0ms pre-stimulus.
6. Semi-automatic rejection of trials with artifacts using peak-to-peak amplitude thresholds (minimum to maximum) of 50 μV for eye movements and 150 μV for blinks and other artifacts (e.g. electrode drift, muscle noise), with adjustments to these thresholds made as-needed at the participant level.

ERPs and mask subtraction

Trials were sorted by condition and trial type (SOA or mask only). Each group of trials was averaged at the participant-level and low-pass filtered at 25 Hz. To isolate stimulus-related neural activity from mask-evoked

activity, the mask-only ERPs were time-shifted so that the timing of the masks were aligned with mask onset in the stimulus-mask trials (following the same procedure as in Del Cul et al. (33)). This was done for each SOA and these time-shifted mask ERPs were subtracted from the corresponding stimulus-mask ERPs, in the following procedure:

1. The mask-only ERP is initially time-locked to mask onset (i.e. time 0 is when the first mask was presented).
2. Time 0 of the mask-only ERP is shifted by -17, -33, -50, -67, or -83ms to match stimulus onset in each SOA. In other words, the mask-only ERP is shifted forwards in time relative to time 0 so that mask onset is aligned with mask onset in each of the stimulus-mask ERPs.
3. The time-shifted mask ERPs are subtracted from the corresponding stimulus-mask ERPs.

ERP analysis

Electrode and time window selection

Grand-average data from an initial group of 15 participants was used to determine electrodes and time windows for ERP components. Three ERPs of interest were identified based on data from this initial group: the P1, displaying a linear increase in amplitude with SOA in both the report and no-report conditions; the N2, displaying bifurcation dynamics in no-report; and the P3b, displaying bifurcation dynamics in the report condition. Based on the data from these 15 participants, the following time windows and electrodes were selected for each component and pre-registered on OSF (nearest channels of the international 10-20 system reported in parentheses):

P1: 100-140ms, electrodes 41-45 and 53-57 (Iz, Oz, O1, O2, O9, O10, PO7, PO8, PO9, PO10).
P3b: 300-500ms, electrodes 1, 4-6, and 13-15 (Pz, P1, P2, CPz, CP1, CP2, Cz).
N2: 250-290ms, electrodes 1-11 and 17-19 (CPz, CP1, CP2, Cz, C1, C2, C3, C4, FCz, FC3, FC4, Fz, F1, F2).

These time windows and electrodes were used for all analyses of these three ERPs, and tested on an independent group of 25 additional participants to confirm replicability. The results of the two groups separately can be found in the supplement.

The N170/VAN analysis was added post-hoc and measured from 160-200ms at electrodes 58 and 64 (bilateral mastoid electrodes). Since the N170/VAN signal was maximal over these mastoid electrodes, EEG data was re-referenced to the average across electrodes for the N170/VAN analyses only.

Assessment of linearity & bifurcation dynamics

Following the methods of Kouider et al. (34), ERP linearity and bifurcation was assessed by comparing average ERP amplitudes across SOAs. An ERP that scales linearly with stimulus strength will exhibit steady and significant increases in amplitude as SOA increases. In contrast, the amplitude of an ERP displaying bifurcation dynamics will not differ significantly between SOAs 17 and 33 or between SOAs 67 and 83, but there will be a significant increase in amplitude from SOAs 33 to 50 and SOAs 50 to 67.

Multivariate pattern analysis

All decoding was conducted using the Amsterdam Decoding and Modeling (ADAM) toolbox (85) for MATLAB. Data was exported from BrainVision Analyzer, formatted for use with the ADAM toolbox using EEGLAB (86) and down sampled to 20Hz. For each SOA and task condition, a linear classifier was trained to distinguish between stimulus-present and mask-only trials using a 10-fold cross-validation procedure. At each timepoint, the decoder was first trained on 9/10ths of the trials and then tested on the remaining subset of trials. This was repeated 10 times such that the decoder was both trained and tested on the entire dataset over the course of the procedure, but the classifier was never tested on data used for training when assessing classifier performance. Trial numbers were balanced across classes so that there were equal numbers of stimulus-present and mask-only trials in each fold. Classifier sensitivity, summarized as the area under the receiver operating curve (AUC), was averaged across the folds to calculate overall classifier performance.

Temporal generalization

The temporal generalization analysis technique can be used to characterize neural dynamics of perceptual representations by testing if a classifier trained on data from one point in time can successfully discriminate between classes when applied to data from different timepoints. If a classifier trained at time T1 is able to

decode stimulus presence at a level significantly above chance when tested at time T2, it suggests that the neural dynamics supporting above-chance decoding at T1 are recurring at T2. Temporal generalization matrices were computed by training classifiers on one timepoint and testing each at every other timepoint. Classification sensitivity was statistically tested against chance and corrected for multiple comparisons using 10,000-permutation cluster-based permutation testing.

Variability analysis

Following the methods of Sergent et al. (35), the inter-trial variability profiles of EEG amplitudes were analyzed at the participant level for three electrode clusters. Trials were sorted by SOA and condition, and mask-only trials were separated according to condition.

Separately for each participant, trial-by-trial EEG amplitudes were averaged across each of the clusters of electrodes used in the ERP analyses. Data were then averaged across 50ms time windows from 0 to 700ms post-stimulus, and the standard deviation across trials of this spatially and temporally averaged EEG signal was calculated for each trial group.

To account for individual differences in variability and mask-related effects, the standard deviation of mask-only trials was subtracted from the standard deviation of stimulus-present trials in the same condition at the participant level. These adjusted results were averaged across participants.

Acknowledgements

Thanks to Sarah Cormiea, Stanislas Dehaene, Steven Hillyard, Nancy Kanwisher, Ken Nakayama, N. Apurva Ratan Murty, Caroline Robertson, and Claire Sergent for helpful discussions. This work was supported by National Science Foundation Grant BCS-1829470 to M.A.C. and M.P., Templeton World Charity Foundation Grant TWCF-2022-30267 to M.A.C. and M.P., the Canadian Institute for Advanced Research to M.A.C., the American Psychological Foundation to M.A.C., and the Stillman Drake Fund for Faculty Development to M.P.

References

1. Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J-F., Poline, J-B., & Riviere, D. Cerebral mechanisms of word masking and unconscious repetition priming. Nat. Neurosci. 4, 752- 758. (2001).
2. Gaillard, R., Dehaene, S., Adam, C., Clemenceau, S., Hasboun, D., Baulac, M., Cohen, L., & Naccache, L. Converging intracranial markers of conscious access. PLoS Biol. 7, 0472-0492. (2009).
3. Gross, J., Schmitz, F., Schnitzler, I., Kessler, K., Shapiro, K., Hommel, B., & Schnitzler, A. Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. Proc. Natl. Acad. Sci. USA. 101, 13050-13055. (2004).
4. Marti, S., & Dehaene, S. Discrete and continuous mechanisms of temporal selection in rapid visual streams. Nat. Commun. 8, 1955. (2017).
5. Pessoa, L., & Ungerleider, L.G. Neural correlates of change detection and change blindness in a working memory task. Cereb. Cortex. 14, 511–520. (2004).
6. Quiroga, R.Q., Mukamel, R., Isham, E.A., Malach, R., Fried. I. Human single-neuron responses at the threshold of conscious recognition. Proc. Natl. Acad. Sci. USA. 105, 3599-3604. (2008).
7. Lau, H., & Rosenthal, D. Empirical support for higher-order theories of conscious awareness. Trends Cogn. Sci. 15, 365-373. (2011).
8. Cohen, M.A., Cavanagh, P., Chun, M.M., & Nakayama, K. The attentional requirements of consciousness. Trends Cogn. Sci. 16, 411-417. (2012).
9. Cohen, M.A., & Dennet, D.C. Consciousness cannot be separated from function. Trends Cogn. Sci. 15, 358-264. (2011).
10. Odegaard, B., Knight, R.T., & Lau, H. Should a few null findings falsify prefrontal theories of conscious perception? J. Neurosci. 37, 9593-9602. (2017).
11. Brown, R., Lau, H., & LeDoux, J.E. Understanding the higher-order approach to consciousness. Trends Cogn. Sci. 23, 754-768. (2019).
12. Mashour, G.A., Roelfsma, P., Changeux, J.P., & Dehaene, S. Conscious processing and the global neuronal workspace hypothesis. Neuron. 105, 776-798. (2020).
13. Beck, D.M., Rees, G., Frith, C.D., & Lavie, N. Neural correlates of change detection and change blindness. Nat. Neurosci. 4, 645–650. (2001).
14. Sergent, C., Baillet, S., & Dehaene, S. Timing of the brain events underlying access to consciousness during the attentional blink. Nat. Neurosci. 8, 1391–1400. (2005).
15. van Vugt, B., Dagnino, B., Vartak, D., Safaai, H., Panzeri, S., Dehaene, S., & Roelfsema, P.R. The threshold for conscious report: signal loss and response bias in visual and frontal cortex. Science. 360, 537–542. (2018).
16. Block, N. The higher order approach to consciousness is defunct. Analysis. 71, 419 - 431. (2011).
17. Lamme, V.A.F. Challenges for theories of consciousness: seeing or knowing, the missing ingredient and how to deal with panpsychism. Philos. Trans. R. Soc. B. 373, 1755. (2018).
18. Koch, C., Massimini, M., Boly, M., & Tononi, G. Neural correlates of consciousness: progress and problems. Nat. Rev. Neurosci. 17, 307-322. (2016).
19. Tononi, G., Boly, M., Massimini, M. & Koch, C. Integrated information theory: from consciousness to its physical substrate. Nat. Rev. Neurosci. 17, 450-461. (2016).
20. Boly, M., Massimini, M., Tsuchiya, N., Postle, B.R., Koch, C., & Tononi, G. Are the neural correlates of consciousness in the front of or in the back of the cerebral cortex? Clinical and neuroimaging evidence. J. Neurosci. 37, 9603-9613. (2017).
21. Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V.A.F. No-report paradigms: extracting the true neural correlates of consciousness. Trends Cogn. Sci. 19, 757-770. (2015).
22. Pitts, M.A., Martinez, A., & Hillyard, S. Visual processing of contour patterns under conditions of inattentional blindness. J. Cogn. Neurosci. 24, 287-303. (2012).
23. Frässle, S., Sommer, J., Jansen, A., Naber, M., & Einhäuser, W. Binocular rivalry: frontal activity relates to introspection and action but not to perception. J. Neurosci. 34, 1738-1747. (2014).

24. Wiegand, K., Heiland, S., Uhlig, C.H., Dykstra, A.R., & Gutschalk, A. Cortical networks for auditory detection with and without informational masking: task effects and implications for conscious perception. Neuroimage. 167, 178–190. (2018).
25. Farooqui, A.A., & Manley, T. When attention and conscious perception deactivates fronto-parietal regions. Cortex. 107, 166-179. (2018).
26. Kapoor, V., Dwarakanath, A., Safavi, S. Werner, J., Besserve, M., Panagiotaropoulos, T.I., & Logothetis, N.K. Decoding internally generated transitions of conscious contents in the prefrontal cortex without subjective reports. Nat. Commun. 13, 1535. (2022).
27. Pitts, M.A., Padwal, J., Fennelly, D., Martínez, A., & Hillyard, S.A. Gamma-band activity and the P3 reflect post-perceptual processes, not visual awareness. Neuroimage. 101, 337–350. (2014).
28. Shafto, J., & Pitts, M.A. Neural signatures of conscious face perception in an inattentional blindness paradigm. J. Neurosci. 35, 10940-10948. (2015).
29. Cohen, M.A., Ortego, K., Kyroudis, A., & Pitts, M.A. Distinguishing the neural correlates of perceptual awareness and post-perceptual processing. J. Neurosci. 40, 4925-4935. (2020).
30. Hatamimajoumerd, E., Ratan Murty, N.A., Pitts, M.A., & Cohen, M.A. Decoding perceptual awareness across the brain with a no-report fMRI masking paradigm. Curr. Biol. 32, 4139-4149. (2022).
31. Chen, Y.K., Cheng, T., & Hsieh, P.J. P3b does not reflect perceived contrasts. eNeuro. 9. (2022).
32. Förster, J., Koivisto, M., & Revonsuo, A. ERP and MEG correlates of visual consciousness: the second decade. Conscious. Cogn. 80, 102917. (2020).
33. Del Cul, A., Baillet, S., & Dehaene, S.. Brain dynamics underlying the nonlinear threshold for access to consciousness. PLoS Biol. 5, 2408-2423. (2007).
34. Kouider, S., Stahlhut, C., Gelskov, S.V., Barbosa, L.S., Dutat, M., de Gardelle, V., Cristophe, A., Dehaene, S., & Dehaene-Lambertz, G. A neural marker of perceptual consciousness in infants. Science. 340, 376-380. (2013).
35. Sergent, C., Corazzol, M., Labouret, G., Stockard, F., Wexler, M., King, J.R., Meyniel, F., & Fressnitzer, D. Bifurcation in brain dynamics reveals a signature of conscious processing independent of report. Nat. Commun. 12, 1149. (2021).
36. Charles, L., Van Opstal, F., Marti, S., & Dehaene, S. Distinct brain mechanisms for conscious versus subliminal error detection. Neuroimage. 73, 80–94. (2013).
37. King, J.R., & Dehaene, S. Characterizing the dynamics of mental representations: the temporal generalization method. Trends Cogn. Sci. 18, 203-210. (2014).
38. Dehaene, S. Consciousness and the brain: deciphering how the brain codes our thoughts. New York, NY: Penguin Books. (2014).
39. Koivisto, M., & Revonsuo, A. Electrophysiological correlates of visual consciousness and selective attention. Neuroreport. 18, 753–756. (2007).
40. Dembski, C., Koch, C., & Pitts, M.A. Perceptual awareness negativity: a physiological correlate of sensory consciousness. Trends Cogn. Sci. 25, 660-670. (2021).
41. Dehaene, S., & Changeux, J.P. Experimental and theorietcal approaches to conscious processing. Neuron. 70, 200-227. (2011).
42. Derda, M., Koculak, M., Windey, B., Gociewicz, K., Wierzchón, M., Cleeremans, A., & Binder, M. The role of levels of processing in disentangling the ERP signatures of conscious visual processing. Conscious. Cogn. 73, 102767. (2019).
43. Ye, M., Lyu, Y., Sclodnick, B., & Sun, H-J. The P3 reflects awareness and can be modulated by confidence. Front. Neurosci. 13, 519. (2019).
44. Dehaene, S., Sergent, C., & Changeux, J.P. A neuronal network model linking subjective reports and objective physiological data during conscious perception. Proc. Natl. Acad. Sci. USA. 100, 8520–8525. (2003).
45. Reiss, J.E. & Hoffman, J.E. Disruption of early face recognition processes by object substitution masking. Vis. Cogn. 15, 789-798. (2007).
46. Fisch, L., Privman, E., Ramot, M., Harel, M., Nir, Y., Kipervasser, S., Andelman, F., Neufeld, M. Y., Kramer, U., Fried, I., & Malach, R. Neural "ignition": enhanced activation linked to perceptual awareness in human ventral stream visual cortex. Neuron. 64, 562–574. (2009).

47. Babiloni, C., Vecchio, F., Buffo, P., Buttiglione, M., Cibelli, G., & Rossini, P.M. Cortical responses to consciousness of schematic emotional facial expressions: a high-resolution EEG study. Hum. Brain Mapp. 31, 1556-1569. (2010),
48. Harris, J.A., Wu, C.T., & Woldorff, M.G. Sandwich masking eliminates both visual awareness of faces and face-specific brain activity through a feedforward mechanism. J. Vis. 11, 3. (2011).
49. Rodríguez, V., Thompson, R., Stokes, M., Brett, M., Alvarez, I., Valdes-Sosa, M., & Duncan, J. Absence of face specific cortical activity in the complete absence of awareness: converging evidence from functional magnetic resonance imaging and event-related potentials. J. Cogn. Neurosci. 24, 396-415. (2012).
50. Navajas, J., Ahmadi, M., & Quiroga, R.Q. Uncovering the mechanisms of conscious face perception: a single-trial study of the N170 responses. J. Neurosci. 33, 1337-1343. (2013).
51. Sandberg, K., Bahrami, B., Kanai, R., Barnes, G.R., Overgaard, M., & Rees, G. Early visual responses predict conscious face perception within and between subjects during binocular rivalry. J. Cogn. Neurosci. 25, 969-985. 2013.
52. Sterzer, P., Jalkanen, L., & Rees, G. Electromagnetic responses to invisible face stimuli during binocular suppression. Neuroimage. 46, 803-808. (2009).
53. Harris, J.A., McMahon, A.R., & Woldorff, M.G. Disruption of visual awareness during the attentional blink is reflected by selective disruption of late-stage neural processing. J. Cogn. Neurosci. 25, 1863-1874. (2013).
54. Suzuki, M. & Noguchi, Y. Reversal of the face-inversion effect in N170 under unconscious visual processing. Neuropsychologia. 51, 400-409. (2013).
55. Bekinschtein, T.A., Dehaene, S., Rohaut, B., Tadel, F., Cohen, L., & Naccache, L. Neural signature of the conscious processing of auditory regularities. Proc. Natl. Acad. Sci. USA. 106, 1672–1677. (2009).
56. King, J.R., Faugeras, F., Gramfort, A., Schurger, A., El Karoui, I., Sitt, J.D., Rohaut, B., Wacongne, C., Labyt, E., Bekinschtein, T., Cohen, L., Naccache, L., & Dehaene, S. Single-trial decoding of auditory novelty responses facilitates the detection of residual consciousness. Neuroimage. 83, 726–738. (2013).
57. Schurger, A., Sarigiannidis, I., Naccache, L., Sitt, J.D., & Dehaene, S. Cortical activity is more stable when sensory stimuli are consciously perceived. Proc. Natl. Acad. Sci. USA. 112, E2083–E2092. (2015).
58. Salti, M., Monto, S., Charles, L., King, J.R., Parkkonen, L., & Dehaene, S. Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. eLife. 4, e05652. (2015).
59. Aru, J., Bachmann, T., Singer, W., & Melloni, L. Distilling the neural correlates of consciousness. Neurosci. Biobehav. Rev. 36, 737-746. (2012).
60. Seth, A.K., & Bayne, T. Theories of consciousness. Nat. Rev. Neurosci. 23, 439–452. (2022).
61. Rahimian, S. The myth of when and where: how false assumptions still haunt theories of consciousness. Consc. Cogn. 97, 103246. (2022).
62. Dennett, D. Consciousness explained. Boston, MA: Little, Brown and Co. (1991).
63. Koivisto, M., & Revonsuo, A. An ERP study of change detection, change blindness, and visual awareness. Psychophysiology. 40, 423-429. (2003).
64. Koivisto, M., & Grassini, S. Neural processing around 200 ms after stimulus-onset correlates with subjective visual awareness. Neuropsychologia. 84, 235–243. (2016).
65. Eklund, R., & Wiens, S. Auditory awareness negativity is an electrophysiological correlate of awareness in an auditory threshold task. Conscious. Cogn. 71, 70–78. (2019).
66. Bola, M., & Doradzinska, (2021). Perceptual awareness negativity--does it reflect awareness or attention? Front. Hum. Neurosci. 15, 742513. (2021).
67. Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., & Thut, G. Spontaneous fluctuations in posterior alpha-band EEG activity reflect variability in excitability of human visual areas. Cereb Cortex. 18, 2010-2018. (2008).
68. Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C.S., & Bäuml, K.H. Prestimulus oscillations predict visual perception performance between and within subjects. Neuroimage. 37, 1465-1473. (2007).

69. Womelsdorf, T., & Fries, P. The role of neuronal synchronization in selective attention. Curr. Opin. Neurobiol. 17, 154-160. (2007).
70. Graziano, M.S.A. A conceptual framework for consciousness. Proc. Natl. Acad. Sci. USA. 119, e2116933119. (2022).
71. Markkula, G. Answering questions about consciousness by modeling perception as covert behavior. Front. Psychol. 6, 803. (2015).
72. Malach, R. Local neuronal relational structures underlying the contents of human conscious experience. Neurosci. Conscious. 2, niab028. (2021).
73. Hohwy, J. Attention and conscious perception in the hypothesis testing brain. Front. Psychol. 3, 96. (2012).
74. Whyte, C.J., & Smith, R. The predictive global neuronal workspace: a formal active inference model of visual consciousness. Prog Neurobiol. 199, 101918. (2021).
75. Dolega, K. & Dewhurst, J. Bayesian frugality and the representation of attention. J. Conscious. Stud. 26, 38-63. (2019).
76. Fedorenko, E., Duncan, J., & Kanwisher, N. Broad domain generality in focal regions of frontal and parietal cortex. Proc. Natl. Acad. Sci. USA. 110, 16616-16621. (2013).
77. Woolgar, A., Duncan, J., Manes, F., & Fedorenko, E. Fluid intelligence is supported by the multiple-demand system not the language system. Nat. Hum. Behav. 2, 200–204. (2018).
78. Luck, S.J., & Vogel, E.K. Visual working memory capacity: from psychophysics and neurobiology to individual differences. Trends. Cogn. Sci. 17, 391-400. (2013).
79. Emrich, S.M., Riggall, A.C., LaRocque, J.J., & Postle, B.R. Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. J. Neurosci. 33, 6516-6523. (2013).
80. Todd, J.J., & Marois, R. Capacity limit of visual short-term memory in human posterior parietal cortex. Nature. 428, 751-754. (2004).
81. Block, N. What is wrong with the no-report paradigm and how to fix it. Trends Cogn. Sci. 23, 1003-1013.
82. Folstein, J.R. & Petten, C.V. Influence of cognitive control and mismatch on the N2 component of the ERP: A review. Psychophysicology. 45, 152-170.
83. Kleiner, M., Brainard, D., & Pelli, D. What's new in Psychtoolbox-3. Perception. 36, 1-16. (2007).
84. Watson, A.B., & Pelli, D.G. Quest: a Bayesian adaptive psychometric method. Percept. Psychophys. 33, 113-120. (1983).
85. Fahrenfort, J.J., van Driel, J., van Gaal, S. & Olivers, C.N.L. From ERPs to MVPA using the Amsterdam Decoding and Modeling Toolbox (ADAM). Front. Neurosci. 12, 368. (2018).
86. Delorme, A., & Makeig, S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. Neurosci. Methods. 134, 9-21. (2004).

Supplementary material for

A novel signature of visual awareness independent of post-perceptual processing

Michael A. Cohen[1,2,5], Cole Dembski[2,3,5], Kevin Ortego[4], Clay Steinhibler[3] & Michael Pitts[3]

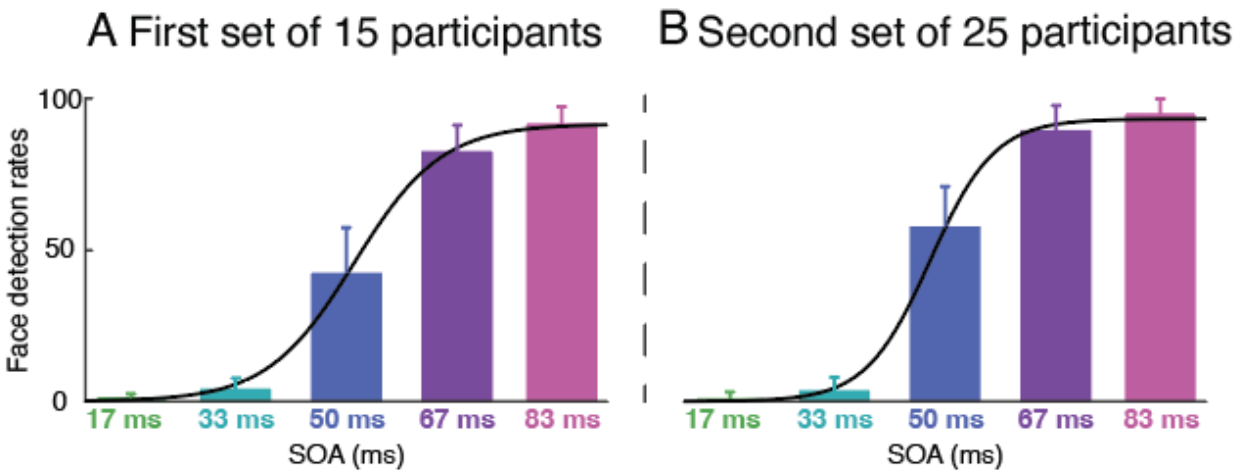

Supplemental Figure 1. Behavioral results from the report condition (i.e., face detection rates) for A) the initial set of 15 participants used to identify the time windows and electrode locations for the ERPs and B) the second set of 25 participants.

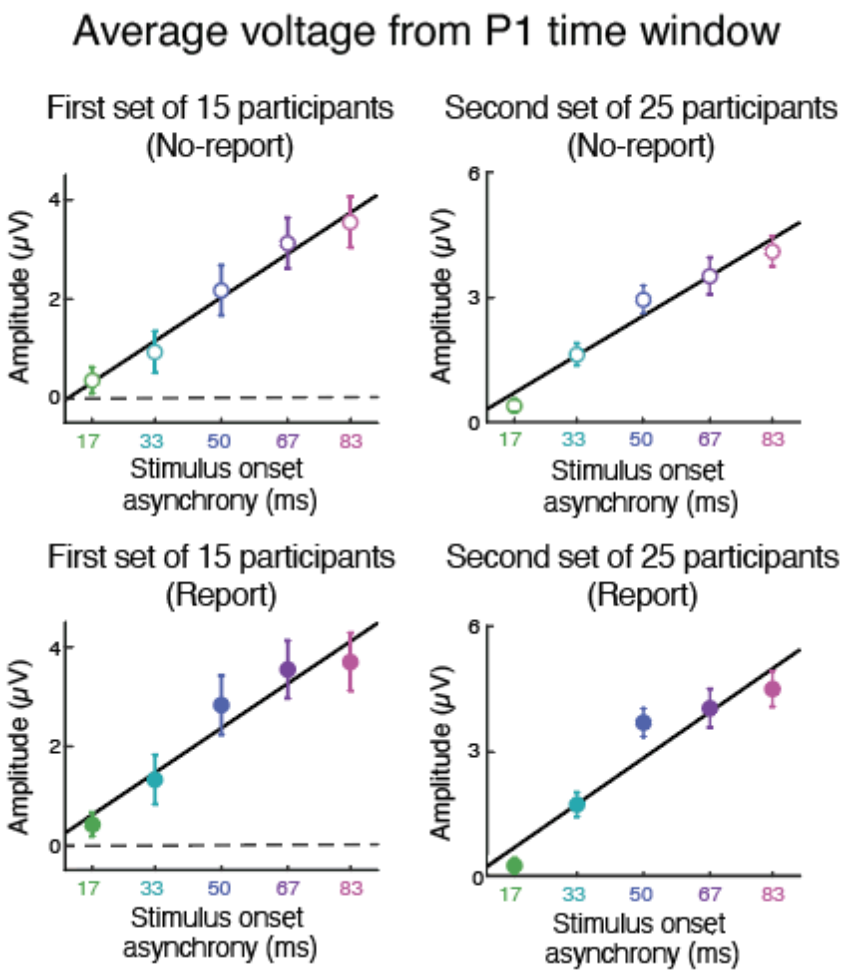

Supplemental Figure 2. Results from the P1 in A) the initial set of 15 participants used to identify the time windows and electrode locations for the ERPs and B) the second set of 25 participants.

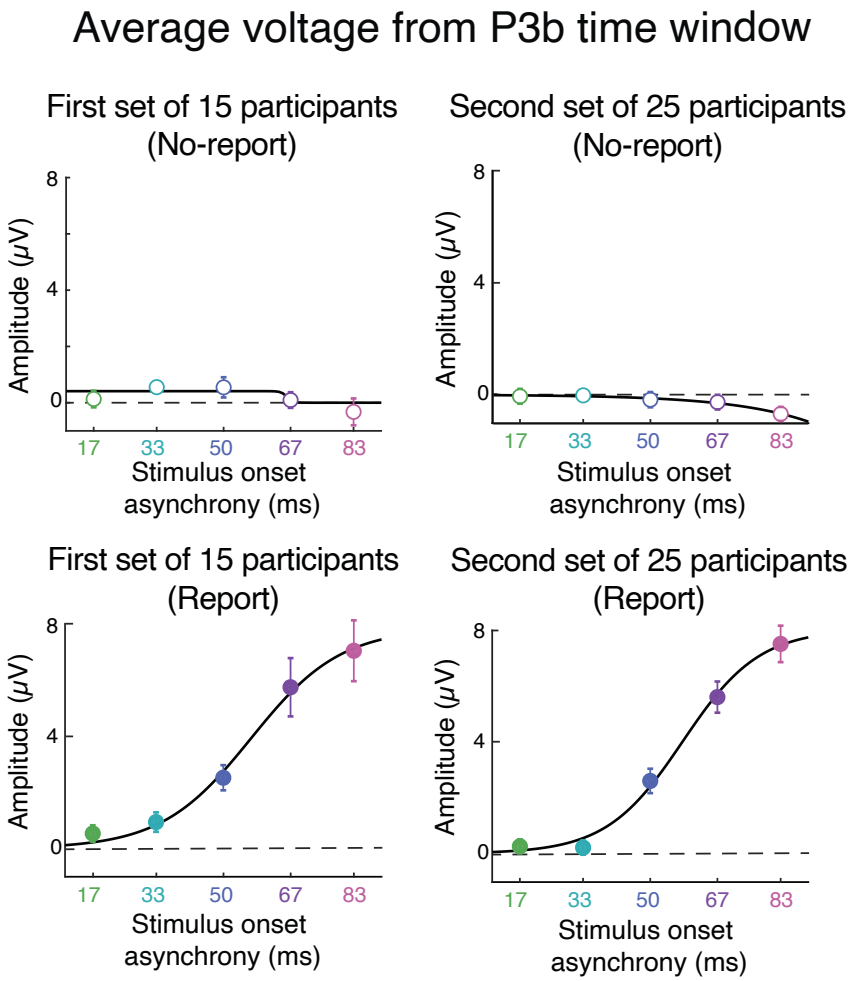

Supplemental Figure 3. Results from the P3b in A) the initial set of 15 participants used to identify the time windows and electrode locations for the ERPS and B) the second set of 25 participants.

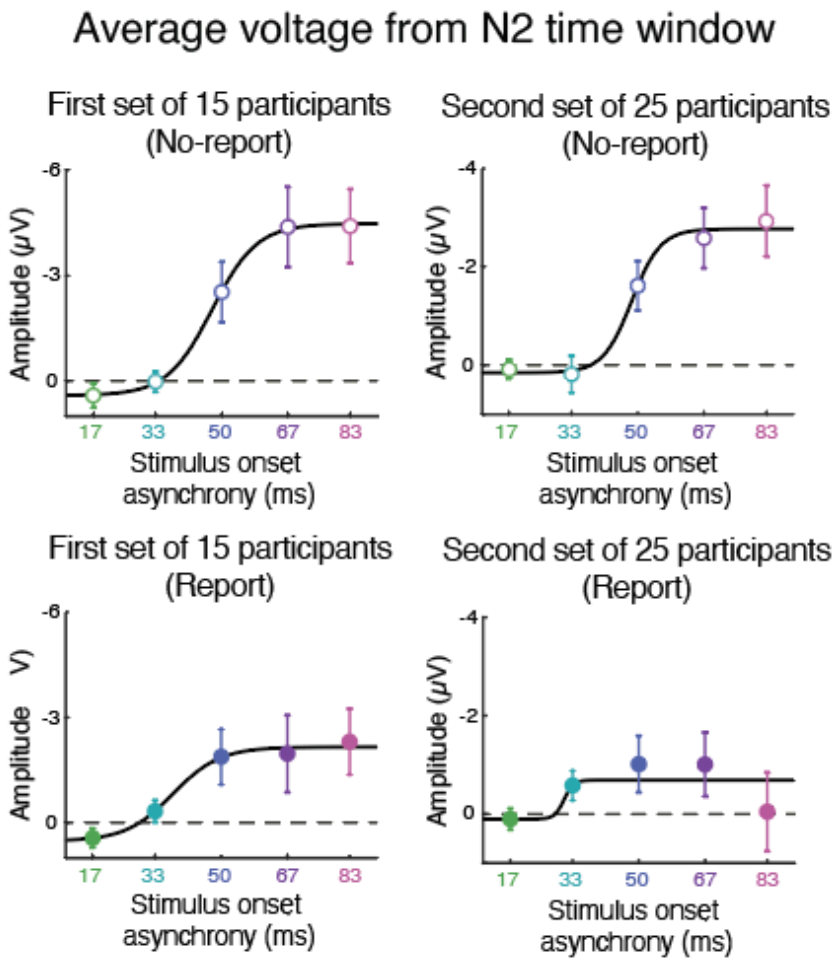

Supplemental Figure 4. Results from the N2 in A) the initial set of 15 participants used to identify the time windows and electrode locations for the ERPs and B) the second set of 25 participants.

Inter-trial EEG variability

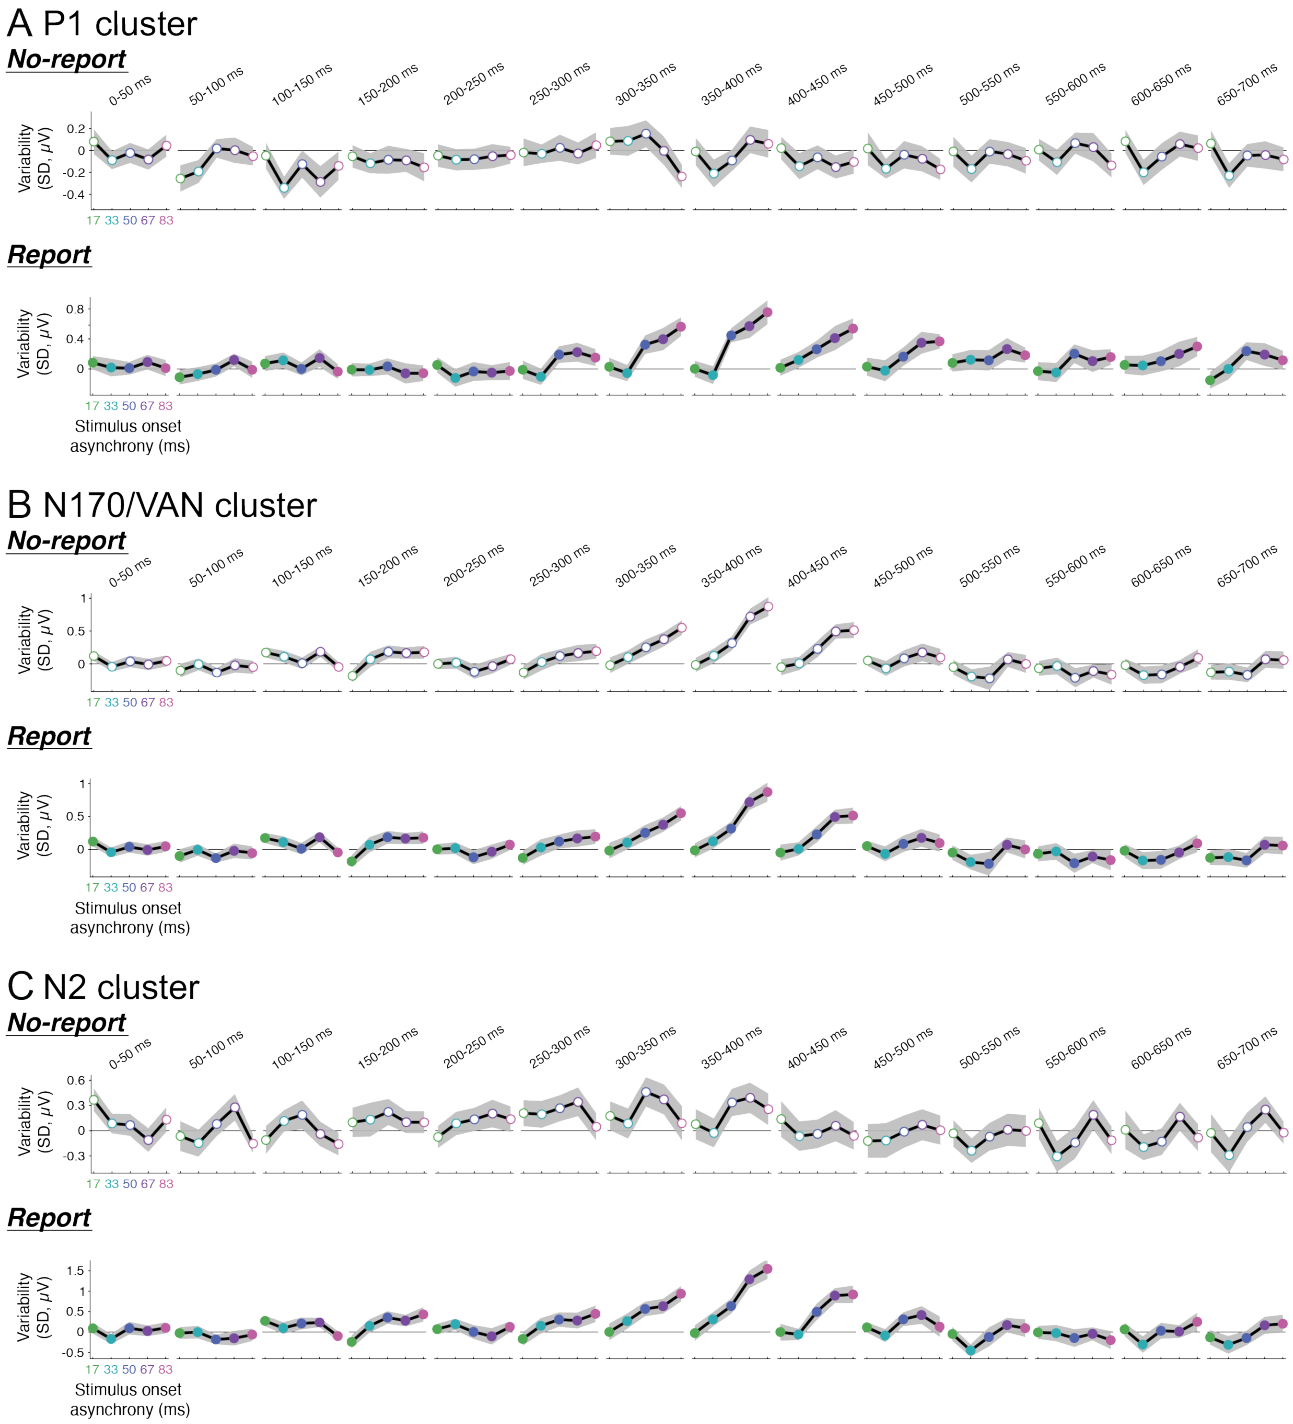

Supplemental Figure 5s. Mean inter-trial EEG variability over time at A) P1 electrodes; B) N170/VAN electrodes; and C) N2 electrodes.

Behavior

Comparison	17 vs. 33	33 vs. 50	50 vs. 67	67 vs. 83
d; p =	0.80; 3.96E-07	1.80; 7.63E-23	1.61; 5.98E-24	0.85; 4.70E-10

Report

	Comparison	17 vs. 33	33 vs. 50	50 vs. 67	67 vs. 83
P1	d; p =	0.85; 9.95E-08	0.89; 9.91E-10	0.23; 0.002	0.16; 0.029
P3b	d; p =	0.10; 0.269	1.05; 8.02E-09	0.99; 2.96E-10	0.48; 4.72E-10
N170/VAN	d; p =	0.78; 1.20E-05	1.06; 7.46E-11	0.29; 7.46E-04	0.07; 0.156
N2	d; p =	0.55; 0.001	0.37; 0.016	0.01; 0.465	0.13; 0.901

No-report

	Comparison	17 vs. 33	33 vs. 50	50 vs. 67	67 vs. 83
P1	d; p =	0.76; 6.73E-05	0.79; 7.25E-10	0.30; 0.021	0.26; 0.022
P3b	d; p =	0.17; 0.182	0.09; 0.687	0.17; 0.920	0.30; 0.984
N170/VAN	d; p =	0.88; 9.31E-07	0.62; 1.34E-04	0.25; 0.011	0.03; 0.664
N2	d; p =	0.06; 0.387	0.82; 4.45E-07	0.39; 3.61E-05	0.06; 0.249

Supplemental Table 1. Statistical result s (i.e., Cohen's *d* and *P*-values from *t*-tests) for behavioral comparisons (top), report condition ERPs (left), and no-report condition ERPs (right). Note that all *t*-tests are uncorrected.