Title

Predicting the unpredictable: critical analysis and practical implications of predictive anticipatory activity

Permalink

Journal

Frontiers in Human Neuroscience, 8

ISSN

Authors

Mossbridge, Julia A
Tressoldi, Patrizio
Utts, Jessica
et al.

Publication Date

2014-03-25

DOI

Copyright Information

Peer reviewed

Predicting the unpredictable: critical analysis and practical implications of predictive anticipatory activity

Julia A. Mossbridge[1] *, Patrizio Tressoldi[2], Jessica Utts[3], John A. Ives[4], Dean Radin[5] and Wayne B. Jonas[4]

[1] Department of Psychology, Northwestern University, Evanston, IL, USA
[2] Dipartimento di Psicologia Generale, Universita di Padova, Padova, Italy
[3] Department of Statistics, University of California at Irvine, Irvine, CA, USA
[4] Samueli Institute, Alexandria, VA, USA
[5] Consciousness Research Laboratory, Institute of Noetic Sciences, Petaluma, CA, USA

A recent meta-analysis of experiments from seven independent laboratories ($n = 26$) indicates that the human body can apparently detect randomly delivered stimuli occurring 1–10 s in the future (Mossbridge et al., 2012). The key observation in these studies is that human physiology appears to be able to distinguish between unpredictable dichotomous future stimuli, such as emotional vs. neutral images or sound vs. silence. This phenomenon has been called *presentiment* (as in "feeling the future"). In this paper we call it *predictive anticipatory activity* (PAA). The phenomenon is "predictive" because it can distinguish between upcoming stimuli; it is "anticipatory" because the physiological changes occur before a future event; and it is an "activity" because it involves changes in the cardiopulmonary, skin, and/or nervous systems. PAA is an unconscious phenomenon that seems to be a time-reversed reflection of the usual physiological response to a stimulus. It appears to resemble precognition (consciously knowing something is going to happen before it does), but PAA specifically refers to *unconscious physiological* reactions as opposed to *conscious premonitions*. Though it is possible that PAA underlies the conscious experience of precognition, experiments testing this idea have not produced clear results. The first part of this paper reviews the evidence for PAA and examines the two most difficult challenges for obtaining valid evidence for it: expectation bias and multiple analyses. The second part speculates on possible mechanisms and the theoretical implications of PAA for understanding physiology and consciousness. The third part examines potential practical applications.

Keywords: presentiment, predictive coding, anticipatory activity, neural prediction, temporal processing
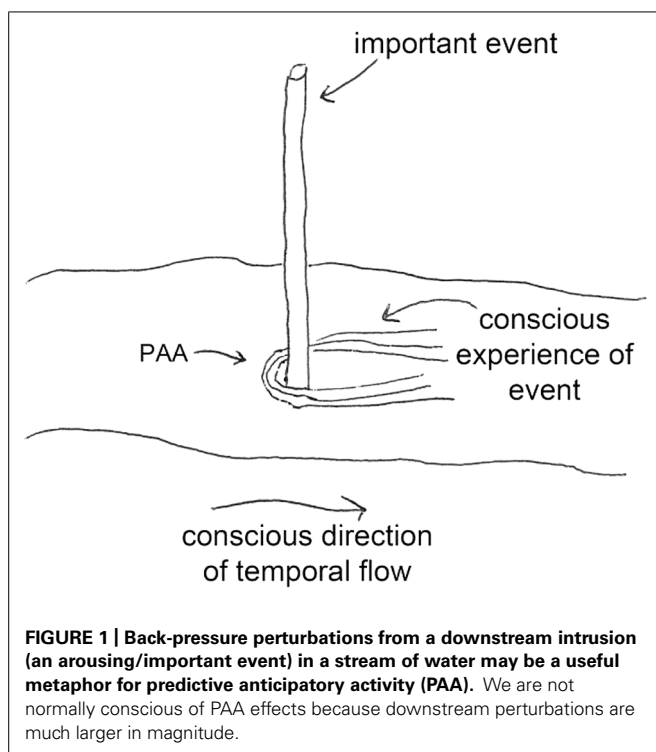
PART 1. THE EVIDENCE FOR PREDICTIVE ANTICIPATORY ACTIVITY

Predictive anticipatory activity (PAA) is defined here as statistically reliable differences between physiological measures recorded seconds before an unpredictable emotional event occurs vs. seconds before an unpredictable neutral event occurs. An emotional or arousing event is defined as an event that activates the sympathetic nervous system; while a neutral event activates the sympathetic nervous system to a lesser degree or not at all. A colloquial definition of PAA is "sensing the future," or *presentiment* (e.g., Radin, 1997, 2004; Radin and Borges, 2009). Here we use the more descriptive term PAA to indicate that this phenomenon is *predictive* of randomly selected future events, *anticipates* these events more often than chance, and is based on *physiological activity* in the autonomic and central nervous systems. In this section, we present the evidence for this phenomenon. We then discuss its implications (Part 2) and potential applications (Part 3).

Predictive anticipatory activity is postulated to be an unconscious physiological phenomenon that may be thought of as a preview of our conscious awareness of future emotional or arousing events. A metaphor may help to provide an intuitive feel for this effect – watching a river move past a stick. The metaphor works as follows (**Figure 1**): imagine that the direction of the water's current is the conscious experience of the flow of time (temporal flow), and imagine that an intrusion in the flow (the stick) is an emotional, arousing, or otherwise important event. The largest disturbance in the water made by the intrusion is downstream (in the "forward" time direction), which is analogous to our conscious reaction to experiencing the important event. But if one examines the flow of water near the stick, one will also see a small perturbation *upstream*, anticipating the intrusion in the water downstream due to the back pressure. Similar to PAA, this upstream perturbation is a hint of things to come. It is not normally part of our conscious awareness and, as with disturbances in a flow of water, the majority of the effect of an intrusion is downstream of the intrusion.

In contrast to PAA, *precognition* may be defined as a perception or a behavior (not a physiological measure) that is influenced by future events. For example, a recent series of experiments published in the *Journal of Personality and Social Psychology* suggested that perception of a future event may influence decisions and memory in the present (Bem, 2011; also see Maier et al., in press;

FIGURE 1 | Back-pressure perturbations from a downstream intrusion (an arousing/important event) in a stream of water may be a useful metaphor for predictive anticipatory activity (PAA). We are not normally conscious of PAA effects because downstream perturbations are much larger in magnitude.

Ritchie et al., 2012). Though it seems plausible that precognition is related to PAA, examination of that possibility is beyond the scope of this article.

Experimental tests of PAA generally use one of two designs, both of which involve a series of randomly interspersed emotional and neutral events. The most common paradigm is one in which participants passively view and/or listen to a series of stimuli that are randomized in terms of stimulus type (e.g., emotional vs. neutral). A less common paradigm is one in which participants actively guess the outcome of each in a series of future events. In both paradigms, care must be taken to ensure that a truly random series of events is generated and that participants or experimenters cannot infer the upcoming event type through usual sensory means. Physiological data (skin conductance, heart rate, respiration rate, EEG activity, etc.) are recorded continuously during the experiment (**Figure 2A**) Each trial may be assessed on a trial-by-trial basis (see Recommendations for Designing Reliable PAA Sensing Tools, below), but more typically Tpre is evaluated by averaging it across multiple trials of similar types (**Figure 2B**; e.g., emotional vs. neutral) in the series.

More than 40 experiments investigating PAA in humans have been published over the past 36 years (including: Hartwell, 1978; Radin et al., 1995, 2011; Bierman and Radin, 1997; Radin, 1997, 2004; Don et al., 1998; Bierman, 2000; Bierman and Scholte, 2002; McDonough et al., 2002; Spottiswoode and May, 2003; McCraty et al., 2004a,b; Sartori et al., 2004; May et al., 2005; Tressoldi et al., 2005, 2009, 2011; Radin and Borges, 2009; Bradley et al., 2011).

This literature prompted a meta-analytic examination of PAA to assess the combined evidence and repeatability of the phenomenon (Mossbridge et al., 2012). The meta-analysis tested the hypothesis that the difference between physiology preceding

emotional and neutral events is in the same direction as the difference after those same events; in other words, it tested the hypothesis that the pre- and post-event physiological differences have the same sign (positive or negative). Using statistically conservative methods, the analysis revealed a small but highly statistically significant effect size in support of the hypothesis [fixed effect: overall ES = 0.21, 95% CI = 0.15–0.27, $z = 6.9$, $p < 2.7 \times 10^{-12}$; random effects: overall (weighted) ES = 0.21, 95% CI = 0.13–0.29, $z = 5.3$, $p < 5.7 \times 10^{-8}$]. Higher quality studies produced a larger overall effect size and greater level of significance, indicating that lack of quality was not responsible for the significant result (Mossbridge et al., 2012)[1].

It is important to note that a meta-analysis is only as good as the data that it examines. Both questionable research practices (QRP) and physiological artifacts have the potential to produce results that mimic a PAA effect. If QRPs are sufficiently widespread, they could potentially be responsible for the highly significant meta-analytic results discussed here. Possible bias can be introduced by experimenter fraud, selective reporting, filter artifacts imposed on the raw data, multiple analyses, and anticipatory and order effects. These and other possible explanations for PAA have been critically examined and found to be lacking in explanatory power (Mossbridge et al., 2012). Here we briefly examine two of the more important criticisms of the evidence for PAA: multiple analyses and order effects, with a focus on expectation bias.
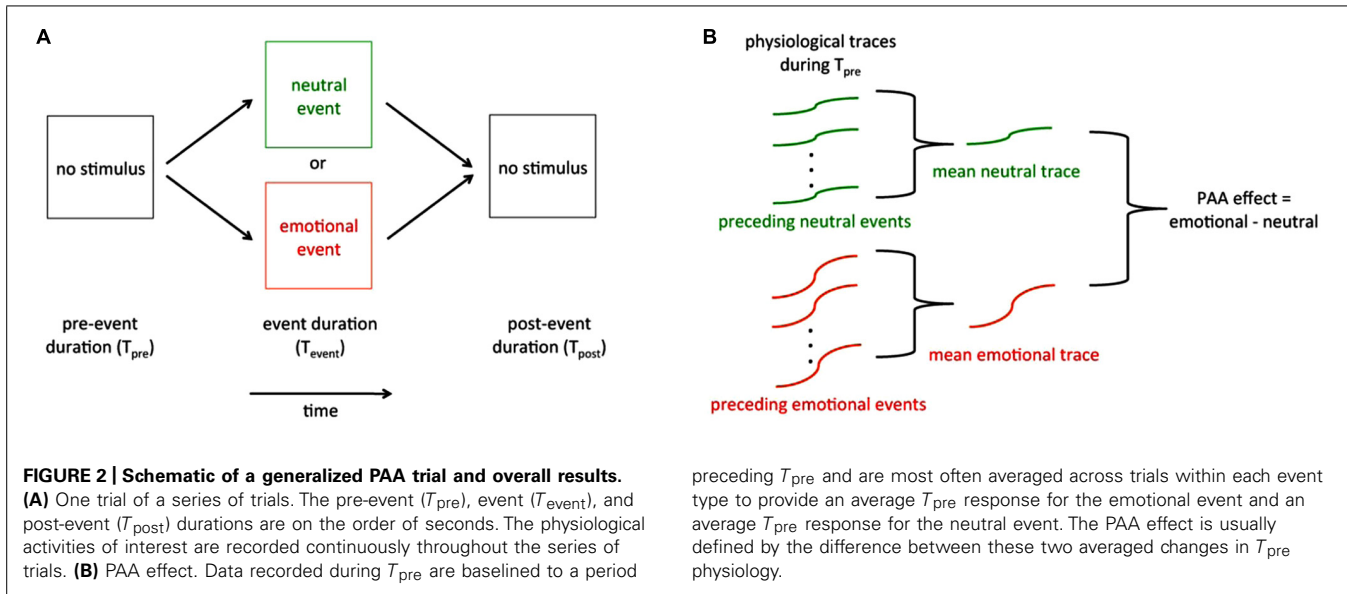
CRITICISMS OF THE EVIDENCE FOR PAA
p-Hacking
One QRP that appears to be common throughout behavioral, social and medical research is to try to create a statistically significant effect using alternative analyses when the originally planned analysis did not find a significant effect (Simmons et al., 2011). If some of the PAA researchers used this technique to produce a significant PAA result, without noting the results as *post hoc*, then a meta-analysis based on those results would have an inflated outcome. This activity, dubbed "*p*-hacking" because it involves cutting through the data in many ways to obtain a *p*-value low enough to declare statistical significance, is a concern when determining the validity of any reported experimental phenomenon. The question is whether *p*-hacking can explain the significant meta-analytic results for PAA.

To answer this question the authors performed a meta-analysis on a subset of the data. This subset consisted of studies using electrodermal activity as the measure of interest, and the meta-analysis of these studies confirmed the presence of a highly significant overall effect (Mossbridge et al., 2012). Electrodermal activity measures have fewer parameters than EEG and fMRI data, so a researcher who wished to perform *p*-hacking would have few parameter choices and thus fewer chances to get a significant effect. One critical parameter in electrodermal studies is the duration of the pre-stimulus activity under examination (Tpre). Interestingly, within each publication and more often than not between publications, experimenters who performed multiple studies using

---

[1]Another meta-analysis using Bayesian statistics and including tests of PAA as well as tests for several other non-ordinary experiences revealed similar results (Tressoldi, 2011), but because that meta-analysis was not limited to PAA alone, it is beyond the scope of this review.

FIGURE 2 | Schematic of a generalized PAA trial and overall results.
(A) One trial of a series of trials. The pre-event (T_{pre}), event (T_{event}), and post-event (T_{post}) durations are on the order of seconds. The physiological activities of interest are recorded continuously throughout the series of trials. **(B)** PAA effect. Data recorded during T_{pre} are baselined to a period preceding T_{pre} and are most often averaged across trials within each event type to provide an average T_{pre} response for the emotional event and an average T_{pre} response for the neutral event. The PAA effect is usually defined by the difference between these two averaged changes in T_{pre} physiology.

electrodermal activity as a measure of interest consistently used the same T*pre* duration (most often 3 s before the stimulus), even when significant results were not found (see Anticipatory Period column in Appendix table A1; Mossbridge et al., 2012). This is not behavior that is consistent with *p*-hacking strategies.

Another feature of the same PAA meta-analysis was that it tested a hypothesis that differed from any explicit hypothesis stated by the authors of the original studies. That is, for the studies included in the meta-analysis, when a hypothesis was formally stated it was of the form that there would be a statistically reliable *difference* between the pre-event physiology for an emotional upcoming event as compared to another (relatively neutral) upcoming event. The direction of this difference was usually not predicted. However, the hypothesis of the PAA meta-analysis was that the direction of the pre-event difference would explicitly mirror the direction of the post-event difference (Mossbridge et al., 2012). In other words, if the direction of the pre-event difference in a physiological measure was the same as the direction of the post-event difference in the same physiological measure, the effect size for the study was given a positive sign (in support of the hypothesis), and if the direction of the pre- and post-event differences were not the same, then the effect size was given a negative sign (in contradiction to the hypothesis). Thus, even if researchers *p*-hacked to get a significant effect in an individual study, it would not necessarily be an effect in the same direction as that tested by the meta-analysis, and therefore it would not necessarily support the hypothesis of the meta-analysis. In fact, several studies examined in the meta-analysis did show effects in directions that opposed the hypothesis of the meta-analysis itself. In this way the meta-analysis was decoupled from one possible source of *p*-hacking. Nevertheless, the cumulative results remained highly significant.

From these analyses, it appears that *p*-hacking and other forms of unreported multiple analyses are not a compelling explanation for PAA. However, it is always possible that the evidence for PAA is influenced by more subtle variations in analyses. Until there is a gold standard experiment that is replicated across laboratories using exactly the same experimental procedure, physiological measures, and statistical analyses, there remains the possibility that multiple analyses could influence the body of evidence supporting PAA. Toward this end, we recommend that all researchers who investigate PAA register their experiments in advance, at any of several registries designed for experiments examining exceptional experiences[2] or at a general experimental research registry[3].

Order effects and expectation bias

Order effects may occur in any experiment with multiple sequential trials, including PAA studies. For example, *forward priming* describes a situation in which previous events influence responses to future events. Thus, responses to the word "flower" are faster if the word is preceded by the word "tree" vs. the word "knife" (Meyer and Schvaneveldt, 1971). Psychophysiologists who wish to avoid priming effects typically use randomization methods to help ensure that it is unlikely for a majority of the participants to systematically experience the same trial order. This increases the likelihood that spurious order effects will average out over participants and will not influence the cumulative results. In addition, the greater the number of trials in any given experiment, the less likely that similar trial orders may occur. If order effects are largely responsible for PAA, there should be a significant negative correlation between study effect size and the number of trials performed. However, this is not the case (Mossbridge et al., 2012).

While order effects do not appear to be a problem in these data, expectation bias is a more subtle effect that requires closer examination. Expectation bias is related to the human propensity to expect a "tail" in a coin toss after observing a series of "head" outcomes (the gambler's fallacy). The reason expectation bias can potentially explain PAA is that a series of (randomly selected)

---

neutral stimuli may produce a physiological shift toward excitement as the presumably imminent emotional trial approaches. In a sequence of trials with several such series of neutral events preceding emotional events, simulations suggest that the resulting physiological data could mimic a PAA effect (Dalkvist et al., 2002; Wackermann, 2002). Thus to understand the mechanisms underlying PAA, it is crucial to determine for each PAA experiment whether expectation bias was a potential explanation for the reported outcome.

There are several ways to quantify expectation bias. For example, one can examine a plot of the physiological measure of interest during T*pre* for an emotional event vs. the number of neutral trials preceding that emotional event. If expectation bias is a viable explanation for the PAA effect in a given experiment, then the activity during T*pre* for emotional events with greater numbers of neutral events preceding them will be larger than for those with fewer neutral events preceding them (Radin, 2004). Of the 26 studies examined in the recent PAA meta-analysis, 19 of them used this method or similar methods to empirically determine whether expectation bias could explain the PAA effect. None of them found that it could. Further, the overall effect size of the subset of studies that performed expectation bias analyses was greater than the effect size of the seven studies that did not perform such analyses, lending little support to the idea that expectation bias creates the PAA effect in general (Mossbridge et al., 2012). However, it is worth noting that studies revealing larger effects usually include more details about attempts to account for mundane explanations such as expectation bias, so there may be an inherent bias here.

Several statistical methods can be used to correct for expectation bias if evidence for that bias is found (e.g., Dalkvist et al., 2013). At least one of us (Mossbridge) has found an expectation bias effect in one PAA study, but because removing the bias produced a larger PAA effect, expectation was not a viable explanation. The bias was removed by discarding data from all but the first trial of the experimental session, because any significant PAA effect on the first trial could not be explained by expectations produced by preceding trials (Mossbridge et al., 2011). Importantly, this method revealed a larger PAA effect than the traditional trial-averaging method (compare Figures 1–6 in Mossbridge et al., 2012). This stronger effect could be due to the reduction of "temporal blurring" when physiological measures preceding only the first trial are examined (see Implications of PAA for Physiology and Consciousness Research, below). Based on this larger effect when only the first trial is considered, experiments are underway in which each individual performs only a single trial[4]. Despite obvious drawbacks due to the increase in inter-individual noise, and the effort involved in collecting data, this approach guarantees that expectation bias is not a viable explanation for any observed PAA effect.

The remainder of this article is based on the assumption that PAA reflects a true anticipatory prediction rather than being a physiological artifact or the result of bias and QRP. This assumption is made to allow us to explore the concept of PAA beyond the initial existential question.

---

[4]Dick Bierman has also performed a slightly different protocol of the single-trial study and found a presentiment effect, but this has not been published.

TYPES OF ANTICIPATORY PHYSIOLOGICAL ACTIVITY

Three broad categories of anticipatory physiological effects are well established in neuroscience and psychophysiology: anticipation of intentional motor activity, anticipation of stimulus detection, and anticipation of a complex firing pattern. PAA may be a novel, fourth category, or it may overlap with one or more of the three established categories.

Anticipation of intentional motor activity is supported by neurophysiological evidence indicating that the neural anticipation of our conscious awareness of having a will to move occurs at least 500 ms (Libet et al., 1983; Haggard and Eimer, 1999) and as much as 10 s (Soon et al., 2008; Bode et al., 2011) before the first conscious report of the will to move. Anticipation of stimulus detection is supported by the fact that EEG alpha activity during the pre-stimulus period for trials presenting stimuli that will be detected differs from alpha activity during pre-stimulus periods preceding stimuli that will not be detected (Ergenoglu et al., 2004; Mathewson et al., 2009; Panzeri et al., 2010). The explanation here is that specific phases and/or amplitudes of neural oscillations facilitate or suppress detection of the upcoming stimulus. Also, anticipation during sleep of a complex firing pattern to be used in the future, a phenomenon dubbed "preplay," has been observed in mouse hippocampal neurons during sleep before entering a novel maze (Dragoi and Tonegawa, 2011). The firing pattern recorded during sleep has greater-than-expected similarity to the patterns recorded when the mouse eventually navigates the maze, an effect explained by the researchers with the idea that the hippocampus recycles generalizable firing patterns from its recent history to create the complex firing patterns accompanying spatial exploration. For these three categories of anticipatory effects, reasonable explanations using the usual cause-preceding-effect assumption are sufficient to explain the results. The usual causal temporal assumptions do not suffice, however, when attempting to understand PAA, because it apparently represents a statistically reliable *retrocausal* effect. In the next section we examine potential mechanism for PAA.

PART 2. POTENTIAL MECHANISMS FOR PAA AND IMPLICATIONS OF PAA
POTENTIAL MECHANISMS FOR PAA
Delayed conscious experience
One seemingly reasonable explanation for PAA is that our conscious mind is wrong about when events occur. That is, our conscious experience of events is delayed by seconds relative to some external/physical time of which we are not conscious. Meanwhile, unconscious neural processes are much less delayed relative to this external time. The explanation goes as follows: one important role of the unconscious is to assess the environment and mobilize physiological resources when it senses challenging external events. Once resources are mobilized and the body is readied, the conscious mind is presented with an ordered version of events that is necessarily temporally delayed so the conscious mind does not initiate counter-productive actions that might interfere with the preparation of physiological resources. Because challenging external events can occur at any time, the conscious mind is always receiving delayed and filtered information about sensory and motor events. Virtually all behavior is unconscious, and

conscious awareness rides on top of this activity like an unfolding and delayed story.

This *delayed conscious experience* hypothesis predicts that we should find brain areas with activity that predicts upcoming consciously perceived events seconds before they occur. As mentioned previously, outside the PAA literature, Bode et al. (2011) and Soon et al. (2008) have reported that up to 10 s before the conscious experience of a decision to produce a motor event, brain activity predicts conscious decisions. These data support the delayed conscious experience hypothesis, but do they reflect PAA?

Although PAA may seem like a sensory counterpart to the predictive coding observed in the motor system, it differs in terms of the order of events and by not involving inferred events or making a decision. In PAA experiments, the physiological and stimulus events are in the wrong order to be explained causally, and they are time-stamped by a computer (not subjectively reported by research participants). Regardless of the absolute times when these events occur, a physiological reaction occurs before the stimulus to which it seems to be linked. Thus PAA neither supports nor refutes the delayed conscious experience hypothesis, but this hypothesis is not a viable explanation for PAA.

Quantum biology
A potentially more viable way to understand PAA effects is that they might reflect an epiphenomenon associated with quantum processing in biological systems. Aharonov et al. (1964, 1988) suggested that one way to explain quantum effects is via interactions between future and past events. This idea has recently been supported by advances in quantum measurement, so-called "weak measurements," which demonstrate that observations in the future do indeed affect observations in the past (Aharonov et al., 1964, 1988; Hosten and Kwiat, 2008; Dixon et al., 2009). Further support of a similar "retrocausal" phenomenon in physics is provided by experimental verification of delayed-choice entanglement (Ma et al., 2012). Finally, because quantum effects have been shown to manifest in biological systems at physiological temperatures, e.g., in photosynthetic reactions (e.g., Scholes, 2011; Dawlaty et al., 2012; Olaya-Castro et al., 2012), it is no longer inconceivable that retrocausal quantum effects can occur in the human nervous system.

However, one problem with a quantum biological explanation for PAA is that retrocausal effects *on the order of seconds* would have to be explainable via quantum processes, and we know of no evidence so far that these effects can occur at that time scale[5]. Nevertheless, exploration into biological quantum effects is in its infancy, and most biological models have yet to entertain the consequences of retrocausation. Thus, the idea that PAA may be related to quantum effects is speculative and currently difficult to test. However, the quantum biology hypothesis demonstrates the value of anomalous phenomena in driving science forward by motivating scientists to search for novel explanations based on emerging scientific concepts. For further discussion of the philosophical and quantum mechanical arguments for time symmetry and retrocausation the reader is referred to an article on

---

backward causation in the Stanford Encyclopedia of Philosophy[6] and to Bierman (2010).

IMPLICATIONS OF PAA FOR PHYSIOLOGY AND CONSCIOUSNESS RESEARCH
Physiology
The most mundane implication of PAA for physiologists is that the time-honored convention of establishing a baseline for physiological post-event measures by subtracting mean pre-event activity may obscure important physiological effects in two ways (Bierman and Radin, 1997; Mossbridge et al., 2012). First, by assuming that pre-event activity is equivalent between event types (without testing this assumption), PAA may be hiding in plain sight. Indeed, several re-examinations of pre-event activity reported in psychophysiology studies conducted for other purposes suggest that the PAA effect actually is present but overlooked (Bierman, 2000; Mossbridge et al., 2012). Of course, when researchers are performing a conventional psychophysiology experiment it is unlikely that they will feel the need to closely examine expectation bias or use truly random techniques to specifically rule out order bias. Thus PAA results found in data from such studies could potentially be due to these mundane explanations.

A second way that baselining data to pre-event physiological activity could obscure important physiological effects is by falsely increasing or decreasing post-event physiological differences. This can occur because pre-event activity is rarely equivalent between event types due to PAA, so subtracting these differing values can produce misleading post-event data as a result.

A more intriguing implication of PAA for those attempting to understand human physiology is that there seems to be a correlation between pre- and post-event responses, such that the magnitude of a post-event response seems to correlate with the magnitude of the corresponding pre-event response. Note that a simple correlation between pre- and post- responses is not what is being discussed here; certainly across individuals there is a strong correlation between the physiological state before and after any event, partially due to the fact that there are stronger inter-individual physiological variations than intra-individual variations. Instead, what has been observed is a correlation between the magnitude of the *change* in a physiological measure before an event and the magnitude of a *change* of that physiological measure after the event. We call this "temporal mirroring." To investigate the idea of temporal mirroring in PAA, Radin (2004) quantitatively tested temporal mirroring using independent ratings of emotionality for different stimuli and found a significant correlation between pre-stimulus response size and stimulus emotionality. In another study, men had large blood-oxygen level dependent (BOLD) post-responses to erotic images randomly distributed along with violent and neutral images, and in men the only significant PAA effect occurred for erotic but not violent images (Bierman and Scholte, 2002). Meanwhile, women in the same study had large BOLD post-responses to violent but not erotic images, which also matched their PAA effects. Others have noted similar gender effects for which post-event differences in responding mirrored pre-event

---

[5]Elementary forms of coherence vanish quickly, but it is possible that coherence in living systems might be "pumped" and sustained for a relatively longer time.

[6]http://plato.stanford.edu/entries/causation-backwards/

differences (McCraty et al., 2004b; Radin and Lobach, 2007; Radin and Borges, 2009; Mossbridge et al., 2011).

Another parameter dependence suggesting temporal mirroring is the relationship between PAA effects for single-modality events (auditory or visual stimuli presented alone) vs. events that are more ecologically valid in that they incorporate multiple modalities (e.g., emotional auditory and visual stimuli presented simultaneously). Well known in the multimodal literature is the idea that responses to stimuli presented in multiple modalities are more robust than responses to each modality alone (e.g., Meredith and Stein, 1986; Meredith et al., 1987). According to the "Hypothesis of Functional Equivalence," any pre-stimulus PAA activity, if it exists, should be drawn on with same readiness as a post-stimulus response to a stimulus (Carpenter, 2012), supporting the functional value of temporal mirroring. Along these lines, in at least one report (Radin, 2004), PAA effects for events presented using multiple modalities are quantitatively larger than PAA effects for single-modality events, though several methodological differences between experiments preclude strong conclusions being drawn from these data.

Thus, it appears that the pre- and post-event physiological responses may mirror one another in size across participants (though pre-event responses are generally smaller than their post-event counterparts), implying that physiological processes are predicting either the importance of the future event to the organism or perhaps the future physiological response itself. These two interpretations may seem similar, but they have different implications for understanding the physiological mechanisms underlying PAA. The idea that PAA predicts the importance of the future event to the organism suggests that even if no physiological post-event response occurs due to some manipulation of the organism's physiology or due to a probable event not occurring, PAA would occur before a highly probable important event. In contrast, if PAA predicts an organism's future physiological response, then no PAA effect would occur if no physiological post-event response occurs. These differing interpretations of the pre-to-post-event mirroring phenomenon have important implications for applications that attempt to harness and amplify PAA, and these implications will be discussed below (see Part 3: Potential Applications of PAA Sensing Tools).

One final implication of PAA for our understanding of physiological systems is that PAA responses apparently decay with time prior to an event. If the size of a PAA response did not decay with time, one would never expect to find PAA, as arbitrarily timed future events that are important would be "temporally blurred" with arbitrarily timed future events that are not important. This is clearly not the case, as inter-trial intervals as short as 10 s have produced significant PAA effects. This does not necessarily mean that PAA effects completely vanish beyond 10 s, but it does indicate that there is some decay of PAA with time preceding the event. Thus, the physiological mechanisms underlying PAA are temporally localized in relation to each event[7].

---

[7]The temporal issue might be related to the boundary conditions of each trial. In other words, the backward temporal information flow could end with the presentation of the stimulus. This could potentially be tested by delaying feedback over multiple trials, some of which are a long sequence of emotional targets, and some of which are a long sequence of calm targets.

Consciousness

A major implication of PAA for our understanding of consciousness is that there must be a necessity for PAA to remain non-conscious most of the time. That is, for most people at most times there is a clear difference between the forward temporal flow of information and experience of which we are aware and a seemingly symmetric flow of information within the non-conscious portions of our experience, as evidenced by the existence of PAA. Why should this be the case? If some part of our nervous system can obtain information about events seconds in the future, would we not have evolved to make this information conscious?

One answer to this question is that the information is not conscious because most of the time it is not useful, like the majority of information that is processed unconsciously. Under this assumption, the mirroring of future physiological states by our unconscious physiological processes is just a side effect of how physiological systems (and in a more general sense the unfolding of events in time) work. The idea is that there has to be a physiological post-event response to produce the PAA prediction of that response, so there is then no point to being consciously aware of PAA effects as the event will occur in short order and there may be nothing we can do to stop it.

A seemingly paradoxical PAA experiment called a "bilking experiment" is one in which a participant's PAA response is used to avoid a future emotional event that presumably caused the PAA response to occur in the first place. If such a PAA response can be shown to exist when there is no accompanying future emotional event, this would invalidate the idea that PAA requires a post-stimulus response and would support the idea that PAA predicts probable vs. actual events. One of us (Tressoldi et al., 2013) has preliminary data that support this idea. However, such bilking experiments are in their infancy, making it difficult to draw conclusions.

Another answer to "Why aren't we conscious of PAA?" is that the conscious mind is not especially skilled at making quick decisions. Unconscious processing is increasingly being recognized as a powerful resource that provides the results of its calculations to conscious awareness for further use and elaboration (e.g., Kahneman, 2011). Converging evidence suggests that unconscious processing can result in learning and decision making that betters, or at least matches, those resulting from conscious processing (e.g., De Houwer et al., 1997; Dijksterhuis et al., 2006; Strick et al., 2011; Voss et al., 2012; Atas et al., 2013; Hassin, 2013). Thus, it might be evolutionarily advantageous for unconscious processing to assess upcoming events, filter them, mobilize resources, and only then inform conscious awareness (see *Potential Mechanisms for PAA* and *Delayed Conscious Experience*, above).

PART 3. POTENTIAL APPLICATIONS OF PAA SENSING TOOLS

Assuming we can understand PAA well enough to amplify and characterize it for a given event of interest, the potential uses of PAA-sensing tools (PAASTs) largely depend on the delay between the PAA signal and the event of interest. Potential applications also depend on whether the PAA is caused by the high a priori probability of an emotional event or is the result of an actual and unavoidable emotional event.

Applications that could benefit from a few seconds advance notice may include: slamming on the brakes of a vehicle to prevent a crash, taking cover in advance of an explosion, or in general orchestrating quick movements to effect a fortuitous result within a few seconds. Applications requiring 10s of seconds might include: course corrections for vehicles, preparing to move out of a location, locating a hiding spot and then hiding there, preparing for a medical emergency, communicating information verbally, or in general, orchestrating more complex chains of action.

POTENTIAL ROADBLOCKS TO DESIGNING RELIABLE PAA SENSING TOOLS

How does one amplify and characterize PAA for an event of interest so that one can determine the practical temporal window of usability for a PAAST? The amplification and characterization process itself could present several difficulties, which we speculate about here (see Recommendations for Designing Reliable PAA Sensing Tools, below, for potential solutions to these problems).
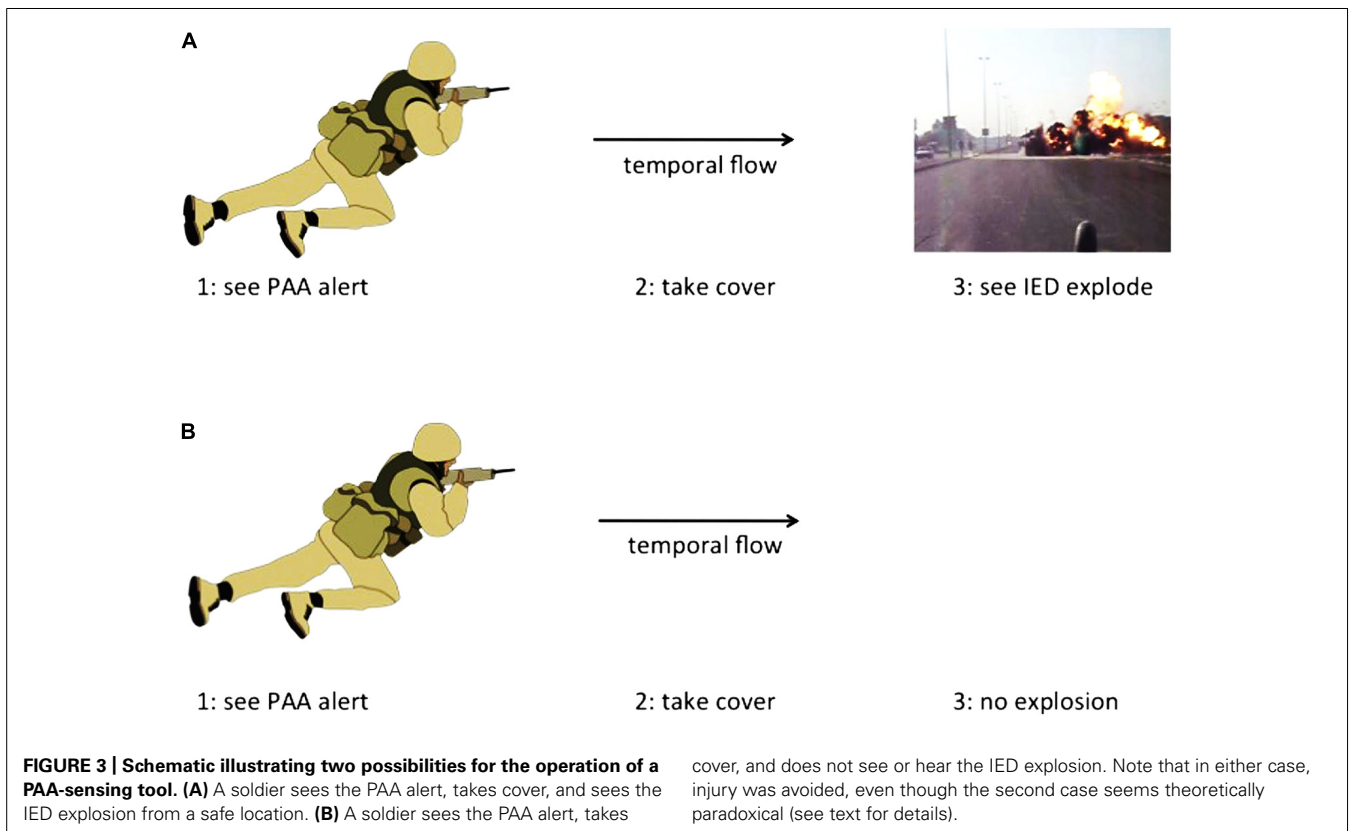
One critical problem to overcome during the characterization process is that when a PAAST is used outside of the laboratory, many events occur beyond the event of interest. For instance, one potentially life-saving application of PAASTs could be to predict the detonation of an improvised explosive device (IED). A PAAST device, when working perfectly, would emit an alert 10–20 s before the detonation of an IED, providing time to avoid the device or take cover. When working well, the PAAST device should not be triggered by emotional events such as an individual soldier's phone call from a friend or a near miss in a firefight. The specific

physiological PAA "signature" of a soldier to the IED detonation must be isolated.

When the signal is characterized and is being amplified, time course is likely critical. Any "temporal blurring" that might occur due to an overlap in physiological responses either before or after the IED detonation must be minimized. For example, if a soldier goes into a simulation where he will be faced with 20 IED detonations in a short time period, a temporal blurring of his or her physiology may occur because responses to later detonations will include responses to previous detonations.

A third potential roadblock is that in the process of characterizing and amplifying PAA to a simulated IED detonation, the event itself may become uninteresting, producing small post-event physiological responses. If the post-event physiological response is small, this would likely reduce the size of the matching PAA response (see Physiology, above). Thus a soldier must be kept engaged and the simulated IED detonation must maintain its emotional and cognitive impact value.

Finally, it is important to address a perceived paradox that seems like a potential stumbling block, but may not be. The paradox goes like this: if PAA is a reflection of the future physiological state of an individual, and the actions taken once a PAA alert is given change the physiological state of that person, the PAA alert may not work in the first place. Imagine this scenario (**Figure 3A**): a PAAST gives a soldier an IED alert. The soldier takes cover away from a nearby trash heap but hears the sound of the subsequent explosion of the IED. The soldier has a typical post-detonation emotional response, which is what produced the PAA alert. Now imagine

FIGURE 3 | Schematic illustrating two possibilities for the operation of a PAA-sensing tool. (A) A soldier sees the PAA alert, takes cover, and sees the IED explosion from a safe location. **(B)** A soldier sees the PAA alert, takes cover, and does not see or hear the IED explosion. Note that in either case, injury was avoided, even though the second case seems theoretically paradoxical (see text for details).

this paradoxical scenario (**Figure 3B**): a PAAST gives a soldier an IED alert. The soldier hides in a tank and puts on headphones. The soldier does not have a typical post-detonation emotional response, so there could not have been a PAA alert. However, there was a PAA alert. How could this work? If the second scenario can occur, it means one of two things: (1) that PAA does not reflect the future physiological state of the individual, or (2) that a non-typical post-detonation response is still an emotional response to saving one's own life and so can still produce PAA. The PAAST still saved the soldier's life, regardless of how it worked (see above for Potential Mechanisms for PAA and Implications of PAA). There may be a theoretical paradox here[8], but perhaps not a practical one.

RECOMMENDATIONS FOR DESIGNING RELIABLE PAA-SENSING TOOLS

Recommendations for designing reliable PAASTs are speculative, as they are based on results from studies varying broadly in their events of interests, methodology, and participants. However, here we make an attempt to outline best practices, as they would currently be defined, for designing a reliable PAAST.

The first step in designing any PAAST would seem to be finding the time course of the decay of PAA for the event of interest. Again, taking the example of IED detonation, once this delay is known, the optimal inter-detonation interval to produce the most reliable PAA effect will become apparent, and the usefulness of that delay can be assessed for the application. Thus, one critical initial experiment would be to use variations in the inter-detonation interval to determine the extent to which the PAA effect can be amplified by reducing temporal blurring with respect to the PAA response (see *Implications of PAA for Physiology and Consciousness Research* and *Physiology*, above).

After finding the critical delay for the IED detonation event across a group of soldiers, the next step would be to characterize the specific physiological "tell" or signature of each soldier who will be using the PAAST. What are the respiration, temperature, blood volume, skin conductance, EEG, and heart rate signatures for each soldier? Obtaining this information would require exposing each soldier to multiple simulated experiences with IED detonations. Based on the little knowledge we have about PAA, these simulation protocols should share several characteristics. To ensure the robust emotional responses that produce reliable PAA, each simulated IED detonation should include at least auditory and visual information, with somatosensory and olfactory cues where possible. To reduce the influence of responses to previous events (temporal blurring) and to ensure continued strong responses after the upcoming event that seem to produce strong PAA (temporal mirroring), delays between simulated detonations should be relatively long (on the order of minutes or hours) and randomly timed. To ensure the generalizability of the PAA signature across multiple situations, detonations should occur in differing scenarios. Because the device must distinguish PAA to IED detonations from emotional responses to other events, events causing emotional responses (but that are not IED detonations) should be intermixed within the series of IED detonations. Finally, to ensure the continued autonomic responses to the IED detonation, any

decrement in the post-stimulus response should be monitored and the simulation stopped at this point and resumed again when the post-stimulus response is robust.

During the entire simulation time, the soldier should be moving and behaving in a life-like environment while physiological data are being recorded continuously. Once enough simulated IED detonations are recorded (likely 30–60), the data from multiple physiological systems preceding each detonation can be analyzed on a trial-by-trial basis using non-linear machine-learning methods to find the characteristic PAA for that soldier. The reason we suggest automated learning is that non-linear, complex relationships between physiological variables and their time courses could exist that allow a fuller characterization of PAA. Along this line of thinking, one of us (Mossbridge) has preliminary data showing that such algorithms may be able to use EEG activity to determine the correct response to an upcoming random event (pressing the left or right mouse key as a correct response to an unpredictable event) on a trial-by-trial basis with above 75% accuracy in ¾ of untrained individuals[9]. We note that without the use of machine learning, finding the combination of electrodes and time points that would produce this predictive effect would have been difficult.

Characterizing the PAA for each soldier could be time consuming, and one possible way to reduce this investment is to screen soldiers to find those who have particularly reliable PAA effects, then characterize PAA to simulated IED explosions in only those soldiers. These PAA-sensitive soldiers could then use the resulting personalized PAAST as "a canary in the coal mine" for their whole team. Another possible time-saving method might be to run multiple soldiers on the IED simulation, then combine their data to characterize a generalized PAA to the simulated IED detonations that can be applied to soldiers who were not tested in the simulated environment. The reliability of this method, of course, would depend on the physiological similarity between the soldiers from whom the data were obtained and those in the field. One possible way to bridge this difference would be to use a generic PAAST on multiple soldiers in combat, so that, for instance, if three of four soldiers have their PAAST activated, all four soldiers would take cover. Assuming an algorithm combining the soldiers' data would not alert the soldiers when only one soldier has a PAA response, this kind of generic physiological profile/multiple-user approach could potentially also reduce the chance of false alarms and increase the likelihood of avoiding real danger.

SUMMARY AND CONCLUSIONS

In summary we have made the following points in this article.

- PAA, the predictive physiological anticipation of a truly randomly selected and thus unpredictable future event, has been under investigation for more than three decades, and a recent conservative meta-analysis suggests that the phenomenon is real.
- Neither QRP, expectation bias, nor physiological artifacts seem to be able to explain PAA.

---

[8]If PAA predicts probable rather than actual futures, there would be no paradox.

- The mechanisms underlying PAA are not yet clear, but two viable yet difficult-to-test hypotheses are that quantum processes are involved in human physiology or that they reflect fundamental time symmetries inherent in the physical world.

- The evidence indicates that there is a temporal mirroring between pre- and post-event physiological events, so that the nature of the post-event physiological response is a reflection of the characteristics of the PAA for that event.

- Temporal blurring, in which closely overlapped emotional events may confuse or minimize both post-event responses and PAA before the event, may be a critical factor in isolating and amplifying PAA.

- The principles of temporal mirroring and temporal blurring both guide the recommendations for designing reliable PAASTs.

- Future research with multiple stimulus modalities, long inter-trial intervals, multiple individuals simultaneously exposed to the same stimulus, and machine-learning techniques will advance our understanding of the nature of PAA and allow a better harnessing of the delay before future events unfold.

REFERENCES

Aharonov, Y., Albert, D. Z., and Vaidman, L. (1988). How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100. *Phys. Rev. Lett.* 60, 1351–1354. doi: 10.1103/PhysRevLett.60.1351

Aharonov, Y., Bergmann, P. G., and Lebowitz, J. L. (1964). Time symmetry in the quantum process of measurement. *Phys. Rev.* 134, B1410–B1416. doi: 10.1103/PhysRev.134.B1410

Atas, A., Faivre, N., Timmermans, B., Cleeremans, A., and Kouider, S. (2013). Nonconscious learning from crowded sequences. *Psychol. Sci.* 25, 113–119. doi: 10.1177/0956797613499591

Bem, D. (2011). Feeling the future: experimental evidence for anomalous retroactive influences on cognition and affect. *J. Pers. Soc. Psychol.* 100, 407–425. doi: 10.1037/a0021524

Bierman, D. (2000). Anomalous baseline effects in mainstream emotion research using psychophysiological variables. *J. Parapsychol.* 64, 239–240.

Bierman, D. (2010). Consciousness induced restoration of time-symmetry (CIRTS), a psychophysical theoretical perspective. *J. Parapsychol.* 74, 273–300.

Bierman, D., and Radin, D. (1997). Anomalous anticipatory response on randomized future conditions. *Percept. Mot. Skills* 84, 689–690. doi: 10.2466/pms.1997.84.2.689

Bierman, D., and Scholte, H. (2002). A fMRI brain imaging study of presentiment. *J. ISLIS* 20, 380–389.

Bode, S., He, A. H., Soon, C. S., Trampel, R., Turner, R., and Haynes, J. D. (2011). Tracking the unconscious generation of free decisions using ultra-high field fMRI. *PLoS ONE* 6:e21612. doi: 10.1371/journal.pone.0021612

Bradley, R., Gillin, M., McCraty, R., and Atkinson, M. (2011). Non-local intuition in entrepreneurs and non-entrepreneurs: results of two experiments using electrophysiological measures. *Int. J. Entrep. Small Bus.* 12, 343–372. doi: 10.1504/IJESB.2011.039012

Carpenter, J. (2012). *First Sight: ESP and Parapsychology in Everyday Life.* Lanham, MD: Rowman & Littlefield Publishers.

Dalkvist, J., Mossbridge, J., and Westerlund, J. (2013). How to handle expectation bias in presentiment experiments. *Paper Presented at the Proceedings of the Parapsychological Association, 56th Meeting*, Paris.

Dalkvist, J., Westerlund, J., and Bierman, D. (2002). "A computational expectation bias as revealed by simulations of presentiment experiments," in *Proceedings of the 45th Annual Convention of the Parapsychological Association*, Paris, 62–79.

Dawlaty, J., Ishizaki, A., De, A., and Fleming, G. (2012). Microscopic quantum coherence in a photosynthetic-light-harveting antenna. *Philos. Trans. R. Soc. A* 370, 3672–3691. doi: 10.1098/rsta.2011.0207

De Houwer, J., Hendrickx, H., and Baeyens, F. (1997). Evaluative learning with "subliminally" presented stimuli. *Conscious. Cogn.* 61, 87–107. doi: 10.1006/ccog.1996.0281

Dijksterhuis, A., Bos, M. W., Nordgren, L. F., and van Baaren, R. B. (2006). On making the right choice: the deliberation-without-attention effect. *Science* 311, 1005–1007. doi: 10.1126/science.1121629

Dixon, P. B., Starling, D. J., Jordan, A. N., and Howell, J. C. (2009). Ultrasensitive beam deflection measurement via interferometric weak value amplification. *Phys. Rev. Lett.* 102, 173601. doi: 10.1103/PhysRevLett.102.173601

Don, N., Donough, B., and Warren, C. (1998). Event-related potentials (ERP) indicators of unconscious PSI: a replication using subjects unselected for PSI. *J. Parapsychol.* 62, 127–145.

Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* 469, 397–401. doi: 10.1038/nature09633

Ergenoglu, T., Demiralp, T., Bayraktaroglu, Z., Ergen, M., Beydagi, H., and Uresin, Y. (2004). Alpha rhythm of the EEG modulates visual detection performance in humans. *Cogn. Brain Res.* 20, 8. doi: 10.1016/j.cogbrainres.2004.03.009

Haggard, P., and Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Exp. Brain Res.* 126, 128–133. doi: 10.1007/s002210050722

Hartwell, J. (1978). Contingent negative variation as an index of precognitive information. *Eur. J. Parapsychol.* 2, 83–103.

Hassin, R. (2013). Yes it can: on the functional abilities of the human unconscious. *Perspect. Psychol. Sci.* 8, 195–207. doi: 10.1177/1745691612460684

Hosten, O., and Kwiat, P. (2008). Observation of the spin hall effect of light via weak measurements. *Science* 319, 787–790. doi: 10.1126/science.1152697

Kahneman, D. (2011). *Thinking, Fast and Slow.* New York: Farrar, Straus and Giroux.

Libet, B., Wright, E. W. Jr., and Gleason, C. A. (1983). Preparation- or intention-to-act, in relation to pre-event potentials recorded at the vertex. *Electroencephalogr. Clin. Neurophysiol.* 56, 367–372. doi: 10.1016/0013-4694(83)90262-6

Ma, X., Zotter, S., Kofler, J., Ursin, R., Jennewein, T., Brukner, C., et al. (2012). Experimental delayed-choice entanglement swapping. *Nat. Phys.* 8, 479–484. doi: 10.1038/nphys2294

Maier, M. A., Büchner, V. L., Kuhbandner, C., Pflitsch, M., Fernández-Capo, M., and Gámiz-Sanfeliu, M. (in press). Feeling the future again: retroactive avoidance of negative stimuli. Available at: http://ssrn.com/abstract=2388097

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., and Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *J. Neurosci.* 29, 8. doi: 10.1523/JNEUROSCI.3963-08.2009

May, E. C., Paulinyi, T., and Vassy, Z. (2005). Anomalous anticipatory skin conductance response to acoustic stimuli: experimental results and speculation about a mechanism. *J. Altern. Complement. Med.* 11, 8. doi: 10.1089/acm.2005.11.695

McCraty, R., Atkinson, M., and Bradley, R. T. (2004a). Electrophysiological evidence of intuition: Part 1. The surprising role of the heart. *J. Altern. Complement. Med.* 10, 11.

McCraty, R., Atkinson, M., and Bradley, R. T. (2004b). Electrophysiological evidence of intuition: Part 2. A system-wide process? *J. Altern. Complement. Med.* 10, 12.

McDonough, B., Don, N., and Watson, C. (2002). Differential event-related potentials to targets and decoys in a guessing task. *J. Sci. Explor.* 16, 187–206.

Meredith, M. A., Nemitz, J. W., and Stein, B. E. (1987). Determinants of multisensory integration in superior colliculus neurons. I. Temporal factors. *J. Neurosci.* 7, 3215–3229.

Meredith, M. A., and Stein, B. E. (1986). Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration. *J. Neurophysiol.* 56, 640–662.

Meyer, D. E., and Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *J. Exp. Psychol.* 90, 227–234. doi: 10.1037/h0031564

Mossbridge, J., Grabowecky, M., and Suzuki, S. (2011). Physiological markers of future outcomes: three experiments on subconscious psi perception during concurrent performance of a guessing task. *Paper Presented at the 54th Conference of the Parapsychology Association*, Paris.

Mossbridge, J., Tressoldi, P., and Utts, J. (2012). Predictive physiological anticipation preceding seemingly unpredictable stimuli: a meta-analysis. *Front. Psychol.* 3:390. doi: 10.3389/fpsyg.2012.00390

Olaya-Castro, A., Nazir, A., and Fleming, G. (2012). Quantum-coherent energy transfer: implications for biology and new energy technologies. *Philos. Trans. R. Soc. A* 370, 3613–3617. doi: 10.1098/rsta.2012.0192

Panzeri, S., Brunel, N., Logothetis, N. K., and Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends Neurosci.* 33, 10. doi: 10.1016/j.tins.2009.12.001

Radin, D. (1997). Unconscious perception of future emotions: an experiment in presentiment. *J. Sci. Explor.* 11, 163–180.

Radin, D. (2004). Electrodermal presentiments of future emotions. *J. Sci. Explor.* 18, 253–273.

Radin, D., and Borges, A. (2009). Intuition through time: what does the seer see? *Explore (NY)*, 5, 200–211. doi: 10.1016/j.explore.2009.04.002

Radin, D., and Lobach, E. (2007). Toward understanding the placebo effect: investigating a possible retrocausal factor. *J. Altern. Complement. Med.* 13, 733–739. doi: 10.1089/acm.2006.6243

Radin, D., Taylor, R., and Braud, W. (1995). Remote mental influence of electrodermal activity: a pilot replication. *Eur. J. Parapsychol.* 11, 19–34.

Radin, D., Vieten, C., Michel, L., and Delorme, A. (2011). Electrocortical activity prior to unpredictable stimuli in meditators and nonmeditators. *Explore* 7, 286–299. doi: 10.1016/j.explore.2011.06.004

Ritchie, S., Wiseman, R., and French, C. (2012). Failing the future: three unsuccessful attempts to replicate Bem's "Retroactive facilitation of recall" effect. *PLoS ONE* 7:e33423. doi: 10.1371/journal.pone.0033423

Sartori, L., Massaccesi, S., Martinell, M., and Tressoldi, P. (2004). Physiological correlates of ESP: heart rate differences between targets and nontargets. *J. Parapsychol.* 68, 351–360.

Scholes, G. (2011). Coherence in photosynthesis. *Nat. Phys.* 7, 448–449. doi: 10.1038/nphys2013

Simmons, J. P., Nelson, L. D., and Simonsohn, U. (2011). False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychol. Sci.* 22, 1359–1366. doi: 10.1177/0956797611417632

Soon, C. S., Brass, M., Heinze, H. J., and Haynes, J. D. (2008). Unconscious determinants of free decisions in the human brain. *Nat. Neurosci.* 11, 543–545. doi: 10.1038/nn.2112

Spottiswoode, S., and May, E. (2003). Skin conductance prestimulus response: analyses, artifacts and a pilot study. *J. Sci. Explor.* 17, 617–641.

Strick, M., Dijksterhuis, A., Bos, M., Sjoerdsma, A., and van Baaren, R. B. (2011). A meta-analysis on unconscious thought effects. *Soc. Cogn.* 29, 738–762. doi: 10.1521/soco.2011.29.6.738

Tressoldi, P. (2011). Extraordinary claims require extraordinary evidence: the case of non-local perception, a classifcal and Bayesian review of evidences. *Front. Psychol.* 2:117. doi: 10.3389/fpsyg.2011.00117

Tressoldi, P. E., Martinelli, M., and Semenzato, L. (2013). Does pupil dilation predict real or future probable events? Available at: http://ssrn.com/abstract=2371577

Tressoldi, P., Marinelli, M., Massaccessi, S., and Sartori, L. (2005). Heart rate differences between targets and non-targets in intuitivate tasks. *Hum. Physiol.* 31, 6446–6650. doi: 10.1007/s10747-005-0108-y

Tressoldi, P., Martinelli, M., Semenzato, L., and Cappato, S. (2011). Let your eyes predict: prediction accuracy of pupillary responses to random alerting and neutral sounds. *Sage Open* 1, 1–7. doi: 10.1177/2158244011420451

Tressoldi, P., Martinelli, M., Zaccaria, E., and Massaccessi, S. (2009). Implicit intuition: how heart rate can contribute to predict future events. *J. Soc. Psych. Res.* 73.

Voss, J., Lucas, H., and Paller, K. (2012). More than a feeling: pervasive influences of memory without awareness of retrieval. *Cogn. Neurosci.* 3, 193–207. doi: 10.1080/17588928.2012.674935

Wackermann, J. (2002). On cumulative effects and averaging artefacts in randomised S-R experimental designs. *Paper Presented at the 45th Annual Convention of the Parapsychological Association*, Paris, 293–305.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.