content – and update to any world model – reflecting a weighted average of those representations. Or, as with binocular rivalry, RM might rapidly sample from competing "realities" and consequently rapidly update internal world models. In fact, these non-MAP policies appear in constructing first-order estimates (e.g., multisensory causal inference; (Rohe & Noppeney, 2015; Wozny, Beierholm, & Shams, 2010)), making them plausible for RM, too.

Why might RM use one of these alternative decision policies, and not stick with MAP? As mentioned above, MAP is only optimal under symmetric error costs (Berger, 2014; Bishop, 2006), meaning the main challenge is choosing the one correct first-order candidate as "real now." But real organisms rarely live in tidy, symmetric cost conditions. Sometimes tagging an only-slightly-wrong first-order representation as "real" might be much less costly than hesitating; other times, a false positive could be catastrophic. Further, empirical evidence links faulty RM to difficulty distinguishing uninformative noise from environmental change (Katthagen et al., 2022), suggesting RM abilities may be related to ability to separately monitor world versus representational variability. Under uneven costs and imperfect volatility source monitoring, the brain may be incentivized to keep more than one option for "reality" actively in play, especially when any "real now" tag is not 100% certain. And with strong stability priors, RM could be using any of these decision policies internally – MAP, sampling, or averaging – and still we would experience a single, slowly updating conscious percept.

If not to support model-based planning under longer sensory horizons, what else would RM-based consciousness be good for? Perhaps the answer lies in characterizing RM's decision policy. For example, MAP-like systems learn quickly but can also be rigid and prone to confirmation bias (Lange et al., 2021); systems that sample or keep multiple options in play learn more flexibly and smoothly shift beliefs, but at the cost of speed. Thus, RM's decision policy carries important implications for its function and interaction with other cognitive capacities, and for understanding the evolutionary pressures that shaped its emergence.

**Acknowledgments.** I thank Rachel Denison for helpful comments.

**Financial support.** This work was supported in part by a fellowship from the Canadian Institute for Advanced Research Program in Brain, Mind, & Consciousness.

**Competing interests.** I declare none.

## References

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221.

Berger, J. (2014). *Statistical decision theory and Bayesian analysis*. Springer.

Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.

Gershman, S. J., Vul, E., & Tenenbaum, J. B. (2012). Multistability and perceptual inference. *Neural Computation*, *24*(1), 1–24.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.

Katthagen, T., Fromm, S., Wieland, L., & Schlagenhauf, F. (2022). Models of dynamic belief updating in psychosis-A review across different computational approaches. *Frontiers in Psychiatry*, *13*, 814111.

Lange, R. D., Chattoraj, A., Beck, J. M., Yates, J. L., & Haefner, R. M. (2021). A confirmation bias in perceptual decision-making due to hierarchical approximate inference. *PLoS Computational Biology*, *17*(11), e1009517.

Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(37), 12366–12378.

Peters, M. A. K. (2022). Towards characterizing the canonical computations generating phenomenal experience. *Neuroscience and Biobehavioral Reviews*, *142*(104903), 104903.

Peters, M. A. K., & Asrari, H. A. (2025). How brains build higher order representations of uncertainty. In arXiv [q-bio.NC]. arXiv. https://arxiv.org/abs/2506.19057

Rohe, T., & Noppeney, U. (2015). Cortical hierarchies perform Bayesian causal inference in multisensory perception. *PLoS Biology*, *13*(2), e1002073.

Wozny, D. R., Beierholm, U., & Shams, L. (2010). Probability Matching as a Computational Strategy Used in Perception. *PLoS Computational Biology*, *6*(8), e1000871–e1000871.

# Postdiction and the speed of consciousness

Ian Phillips*

William H. Miller III Department of Philosophy and Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, USA
ianbphillips@jhu.edu
https://www.ianbphillips.com/

*Corresponding author.

doi:10.1017/S0140525X25102859, e187

**Abstract**

A central plank of Fleming and Michel's thought-provoking paper is that postdiction reveals a lower bound for the speed of consciousness, showing that perceptual awareness is slow and motivating their perceptual reality monitoring theory of consciousness. The plank cannot bear the weight: Postdiction neither demonstrates a lower bound for the speed of consciousness nor shows that awareness is slow.

In postdiction, perceptual experience of an initial, target stimulus is seemingly altered or eliminated by a second stimulus, despite presentation after target onset. In metacontrast masking, a mask presented 70ms after a target apparently renders the target invisible (Breitmeyer & Öğmen, 2006). In other cases, the second stimulus does not eliminate but modulates target awareness: collinear lines may appear misaligned due to a rectangle presented 200ms after line onset (Sugita, Hidaka, & Teramoto, 2018); and a stream of Verniers may appear as fused, contingent on the presentation of an anti-Vernier within 450ms (Drissi-Daoudi, Doerig, & Herzog, 2019). Fleming and Michel also discuss attentional effects, such as "retro-perception" in which, as they see it, an attentional post-cue presented within 400ms, "postdictively determines [a] target's conscious or unconscious fate" (Sergent et al., 2013; Thibault et al., 2016). Here, though, Thibault et al.'s mixture-model modeling

approach rests on faulty assumptions (Schurgin, Wixted, & Brady, 2020) and retro-cueing is better understood in terms of a decaying memory trace consolidated by a valid cue (Cohen, Keefe, & Brady, 2023).

According to Fleming and Michel, postdictive effects show that "conscious perception does not generally occur before 350ms after stimulus onset". Given such sluggishness, they infer that its purpose cannot be online action guidance. Are they right? Does postdiction show that consciousness is slow? In arguing their case, Fleming and Michel make two implicit assumptions. First, that we can ascribe a single speed to the conscious perception of a target (once context, eccentricity, and attention are fixed). Second, that there is a single, determinate moment when signals reach consciousness (a Cartesian "finish line"; Dennett & Kinsbourne, 1992). Both assumptions should be rejected. There is extensive evidence that target – and especially motion – detection unfolds faster than feature discrimination, and that different features are discriminable at different speeds. Moreover, just as signal detection theorists have long discarded the idea of a threshold above which signals are perceived and below which not, we should equally give up the idea that there is some privileged moment at which perceptual processing delivers its first and final report to consciousness.

Dispensing with these assumptions, there is no inconsistency between postdiction and a model on which conscious perception is rapid but initially partial and unsettled – prone to subsequent enrichment and revision. Such a model secures the best of both worlds: "a continuous and dynamic process of perception and planning" in which information is continuously integrated over time and yet results are made rapidly available for perceptual action guidance.

Does this account make an objectionable "overflow gambit," allowing for the rapid conscious perception of targets only at the price of denying their cognitive inaccessibility? I do not think so.

First, it is unclear that overflow is objectionable. Fleming and Michel claim that overflow is contradicted by data from Cowan and Greenspahn (1995). Cowan and Greenspahn presented subjects with displays comprising either two successive targets on the left (50% trials), or a target on the left and then a target on the right – generating a percept of continuous motion. Subjects were cued to respond rapidly either when the target reached the endpoint or when it passed the midpoint. According to Fleming and Michel, "The overflow gambit predicts that subjects should respond faster in the end point condition [since] subjects should first perceive the target at the end point, and then perceive the apparent motion only after this." Reaction times, however, were the same across conditions. But the overflow theorist is not committed to the falsified prediction. An alternative possibility is discussed by Cowan and Greenspahn (1995: 548). On this, information about motion becomes available at the very same time as information about the location of the second target. This leads to rapid awareness of the second target and also to a pure (i.e., location independent) motion percept (Nakayama, 1985). In the endpoint condition, subjects respond as soon as they see the second target; in the midpoint condition, they respond on the basis of the pure motion percept, with no significant difference in reaction times. Once further motion processing has taken place, this partial experience is updated to include continuous motion from left to right.

Second, Fleming and Michel do not consider the possibility that subjects might be rapidly conscious of unmodulated targets *and briefly cognitively able to access them*. Yet, there is evidence of exactly this. Lachter and colleagues found that a subject's ability to discriminate target-mask from mask-mask pairs in a metacontrast paradigm was dependent on how rapidly they responded (Lachter, Durgin, & Washington, 2000; Lachter & Durgin, 1999). When making speeded responses, subjects performed much better than when responding slowly, suggesting that information was briefly cognitively available before being lost. In general, then, a model in which postdiction revises our awareness does not mean that early editions are not fleetingly accessible.

Fleming and Michel will doubtless claim that such speeded responses are driven by unconscious perception. But this cannot be assumed at this stage in the dialectic. We are told that "a wealth of evidence supports the view that online action control can be achieved unconsciously." However, this view should not be regarded as settled science. Significant controversy rightly continues to attend alleged unconscious action guidance in visual form agnosia (e.g., Schenk, 2012; Rossetti, Pisella, & McIntosh, 2017), dissociations between action and perception in visual illusions (e.g., Kopiske et al., 2016), performance/awareness dissociations in blindsight (e.g., Overgaard, 2011; Phillips, 2021), and other influential cases. The upshot is that the hypothesis that the perceptual information used to guide rapid actions is conscious remains very much on the table – and with it, the idea that one central function of consciousness is online action guidance.

**Financial support.** None.

**Competing interests.** None.

## References

Breitmeyer, B. G., & Öğmen, H. (2006). *Visual masking: Time slices through conscious and unconscious vision* (2nd ed.) Oxford University Press.

Cohen, M. A., Keefe, J., & Brady, T. F. (2023). Perceptual awareness occurs along a graded continuum: No evidence of all-or-none failures in continuous reproduction tasks. *Psychological Science*, *34*(9), 1033–1047.

Cowan, N., & Greenspahn, E. (1995). Timed reactions to an object in apparent motion: Evidence on Cartesian and non-Cartesian perceptual hypotheses. *Perception & Psychophysics*, *57*(4), 546–554.

Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain sciences*, *15*(2), 183–201.

Drissi-Daoudi, L., Doerig, A., & Herzog, M. H. (2019). Feature integration within discrete time windows. *Nature Communications*, *10*(4901), 1–8.

Kopiske, K. K., Bruno, N., Hesse, C., Schenk, T., & Franz, V. H. (2016). The functional subdivision of the visual brain: Is there a real illusion effect on action? A multi-lab replication study. *Cortex*, *79*, 130–152.

Lachter, J., & Durgin, F. H. (1999). Metacontrast masking functions: A question of speed? *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 936–947.

Lachter, J., Durgin, F. H., & Washington, T. (2000). Disappearing percepts: Evidence for retention failure in metacontrast masking. *Visual Cognition*, *7*, 269–279.

Nakayama, K. (1985). Biological image motion processing: A review. *Vision Research*, *25*, 625–660.

Overgaard, M. (2011). Visual experience and blindsight: a methodological review. *Experimental Brain Research*, *209*(4), 473–479.

Phillips, I. (2021). Blindsight is qualitatively degraded conscious vision. *Psychological Review*, *128*(3), 558–584.

Rossetti, Y., Pisella, L., & McIntosh, R. D. (2017). Rise and fall of the two visual systems theory. *Annals of physical and rehabilitation medicine*, *60*(3), 130–140.

Schenk, T. (2012). No dissociation between perception and action in patient DF when haptic feedback is withdrawn. *Journal of Neuroscience*, *32*(6), 2013–2017.

Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a unified theory of visual memory strength. *Nature Human Behaviour*, *4*(11), 1156–1172.


https://doi.org/10.1017/S0140525X25000068

Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Current Biology : CB*, *23*(2), 150–155.
Sugita, Y., Hidaka, S., & Teramoto, W. (2018). Visual percepts modify iconic memory in humans. *Scientific Reports*, *8*(13396), 1–7.
Thibault, L., Van den Berg, R., Cavanagh, P., & Sergent, C. (2016). Retrospective attention gates discrete conscious access to past sensory stimuli. *PLoS ONE*, *11*(2), e0148504.

# Speed of consciousness results: Implications for conscious agency and free will

Adina L. Roskies*

Department of Philosophy, University of California Santa Barbara, Goleta, US
aroskies@ucsb.edu

*Corresponding author.

doi:10.1017/S0140525X25102963, e188

**Abstract**

Speed of consciousness results may contribute to the popular notion that our consciousness is inefficacious. I argue that Fleming and Michel's claim that consciousness is slow, coupled with their idea that consciousness marks a predictive world model, instead supports a view of agency in which we are consciously engaged with the world as it unfolds, both in perception and response.

Fleming and Michel (F&M) argue that many of our responses to the world can be mediated by unconscious processes, and they provide compelling evidence of fast unconscious processes able to mediate action in humans and throughout the evolutionary tree. If so much can be done unconsciously, what is conscious sensory processing for? And how does F&M's discussion bear upon recent challenges to the efficacy of consciousness in directing human action?

According to F&M, consciousness is important for model-based planning, a need engendered by extended sensory horizons. Model-based planning involves (usually unconscious) simulations on a world model, and when a system encompasses both online perception and simulation, these two informational streams need to be distinguished. Thus, they suggest that consciousness is related to perceptual reality monitoring. As models of perceptual reality must be sensitive to estimates of reliability and should modulate evidence gathering, F&M argue that what we are conscious of is a constructed world-model that passes some threshold for reliability as indicated by subjective confidence.

I am generally in agreement with their account of the function of consciousness as identifying a high-credence model to support flexible, deliberative planning and decision-making. This provides a partial answer to one of the most vexing questions in both philosophy of mind and evolutionary theory. It is partial because nothing about their story (or any extant story) answers the deepest puzzle about the function of consciousness, which is why it has phenomenal character. In their view, consciousness is meant to mark the reliable and fixed version of a world model, but information can be marked as special in many ways. Their story leaves unaddressed why (and how) its marking should result in phenomenality.

Despite finding their functional story plausible, the way in which they tell this story may lead to some misconceptions regarding how and when consciousness functions in ordinary agency. What follows is more a clarification than a criticism of F&M's view.

F&M suggest that consciousness is slow: on the basis of postdictive effects on conscious reports, they estimate that it takes approximately 350–400 ms for a visual stimulus to reach conscious awareness. The implication is that our awareness lags significantly behind the worldly present and that, for the most part, our reactions to the world occur unconsciously. This claim needs to be scrutinized in order to avoid misapprehensions about what the temporal profile of consciousness means for human agency and for free will.

It is undeniably true that neural processing takes time, and thus, there are limits to what we can respond to, unconsciously or consciously. Even so, the timing data for conscious experience seems surprisingly slow. F&M canvas the literature on postdictive visual illusions, trying to fix on a number to characterize the time of consciousness, which they identify as approximately 350–400 ms. Their discussion implies that there is a correct single number to be found, but another possibility exists: that different aspects of perception are processed with different time courses. Such a phenomenon is commonplace, as when one becomes aware of having touched something scalding before registering the pain the burn brings. With respect to the Vernier effect F&M describe and other visual illusions, it is possible that the fine details of a visual stimulus take longer to reach consciousness than more general characteristics, and the time-course of retrodictive effects may depend on reliability, magnitude, and so on. If consciousness is related to confidence, as they suggest, then it is likely that different aspects of the world are available to consciousness at different times. Global masking effects that prevent any conscious percept may suggest otherwise, but they are also difficult to interpret because of potential interference with memory and response processing.

The point remains that dynamic aspects of consciousness should not be ignored or underestimated. These dynamic aspects may also be related to the temporal profile of what is experienced. As William James recognized more than a century ago, the experienced present is not an instant, but a temporal slice of some thickness that already encodes information about the future just as it does about the past. This is the "specious present," a construct that fits well with F&M's view that the contents of consciousness are predictive models.

One reason all this matters is that it bears upon deep questions about human agency. For example, many people already erroneously assent to the claim that most of our actions are initiated unconsciously and that "our brains decide before we do," largely based on a common misunderstanding of experimental results involving the "readiness potential" from Benjamin Libet dating from the 1980s (Libet et al, 1983; Libet 1985). This view has negatively influenced discussions of free will for the past forty years, despite the fact that the data do not warrant this common interpretation (see, e.g., Schurger et al, 2021).

The discussion of F&M about the speed of consciousness is likely to feed into the very same misconceptions, for it implies that our conscious selves are always half a second behind reality. If this