

Action Biases Perceptual Decisions Toward Expected Outcomes

Daniel Yon

Birkbeck, University of London and Goldsmiths, University of London

Vanessa Zainzinger

Birkbeck, University of London

Floris P. de Lange
Radboud University

Martin Eimer and Clare Press
Birkbeck, University of London


We predict how our actions will influence the world around us. Prevailing models in the action control literature propose that we use these predictions to suppress or “cancel” perception of expected action outcomes, to highlight more informative surprising events. However, contrasting normative Bayesian models in sensory cognition suggest that we are more, not less, likely to perceive what we expect—given that what we expect is more likely to occur. Here we adjudicated between these models by investigating how expectations influence perceptual decisions about action outcomes in a signal detection paradigm. Across three experiments, participants performed one of two manual actions that were sometimes accompanied by brief presentation of expected or unexpected visual outcomes. Contrary to dominant cancellation models but consistent with Bayesian accounts, we found that observers were biased to report the presence of expected action outcomes. There were no effects of expectation on sensitivity. Computational modeling revealed that the action-induced bias reflected a sensory bias in how evidence was accumulated rather than a baseline shift in decision circuits. Expectation effects remained in Experiments 2 and 3 when orthogonal cues indicated which finger was more likely to be probed (i.e. task-relevant). These biases toward perceiving expected action outcomes are suggestive of a mechanism that would enable generation of largely veridical representations of our actions and their consequences in an inherently uncertain sensory world.

Keywords: action, perception, prediction, sensorimotor, modeling

Effectively acting on the world around us requires predicting the consequences of our actions (James, 1890). We select actions based on their predicted outcomes and use these predictions to generate rapid corrective movements when we experience deviant sensory input (Hommel, Müsseler, Aschersleben, & Prinz, 2001; Wolpert, Ghahramani, & Jordan, 1995). Influential “Cancellation” models in the action control literature propose that we also use these predictions to suppress perception of expected sensory inputs across sensory modalities (Bays & Wolpert, 2007; Blakemore, Wolpert, & Frith, 1998; Fiehler, Brenner, & Spering, 2019; Kiltner & Ehrsson, 2017; Kiltner, Houborg, & Ehrsson, 2019; see also Müsseler & Hommel, 1997; Figure 1a). Such a mechanism would

allow us to ignore predictable sensations and therefore remain maximally sensitive to more behaviorally relevant unexpected events. Such cancellation models provide an appealing explanation for why it is difficult to tickle oneself (Weiskrantz, Elliott, & Darlington, 1971). The idea has also drawn wide support from studies showing that sensory events predictably resulting from action are perceived as less intense than similar events presented in the absence of action (Bays, Wolpert, & Flanagan, 2005; Sato, 2008). Indeed, interest in cancellation mechanisms has been galvanized by studies suggesting an intimate link between such effects and the feelings of control that accompany our movements (the “sense of agency”), with dysfunctions of cancellation associ-

This article was published Online First December 7, 2020.

 Daniel Yon, Department of Psychological Sciences, Birkbeck, University of London, and Department of Psychology, Goldsmiths, University of London; Vanessa Zainzinger, Department of Psychological Sciences, Birkbeck, University of London; Floris P. de Lange, Donders Institute for Brain, Cognition, and Behavior, Radboud University; Martin Eimer and Clare Press, Department of Psychological Sciences, Birkbeck, University of London.

All authors contributed to the design of the study. Vanessa Zainzinger and Daniel Yon collected the data, which were analyzed and modeled by Daniel Yon in conjunction with Clare Press. Daniel Yon wrote the manuscript, and all authors were involved in revisions. Clare Press supervised this work. This work was funded by grants from the Leverhulme Trust

(RPG-2016–105) and Wellcome Trust (204770/Z/16/Z). We are grateful to Richard Cooper for advice concerning the modeling.

This article has been published under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. Copyright for this article is retained by the author(s). Author(s) grant(s) the American Psychological Association the exclusive right to publish the article and identify itself as the original publisher.

Correspondence concerning this article should be addressed to Daniel Yon, Department of Psychology, Goldsmiths, University of London, Lewisham Way, London, SE14 6NW, United Kingdom. E-mail: d.yon@gold.ac.uk

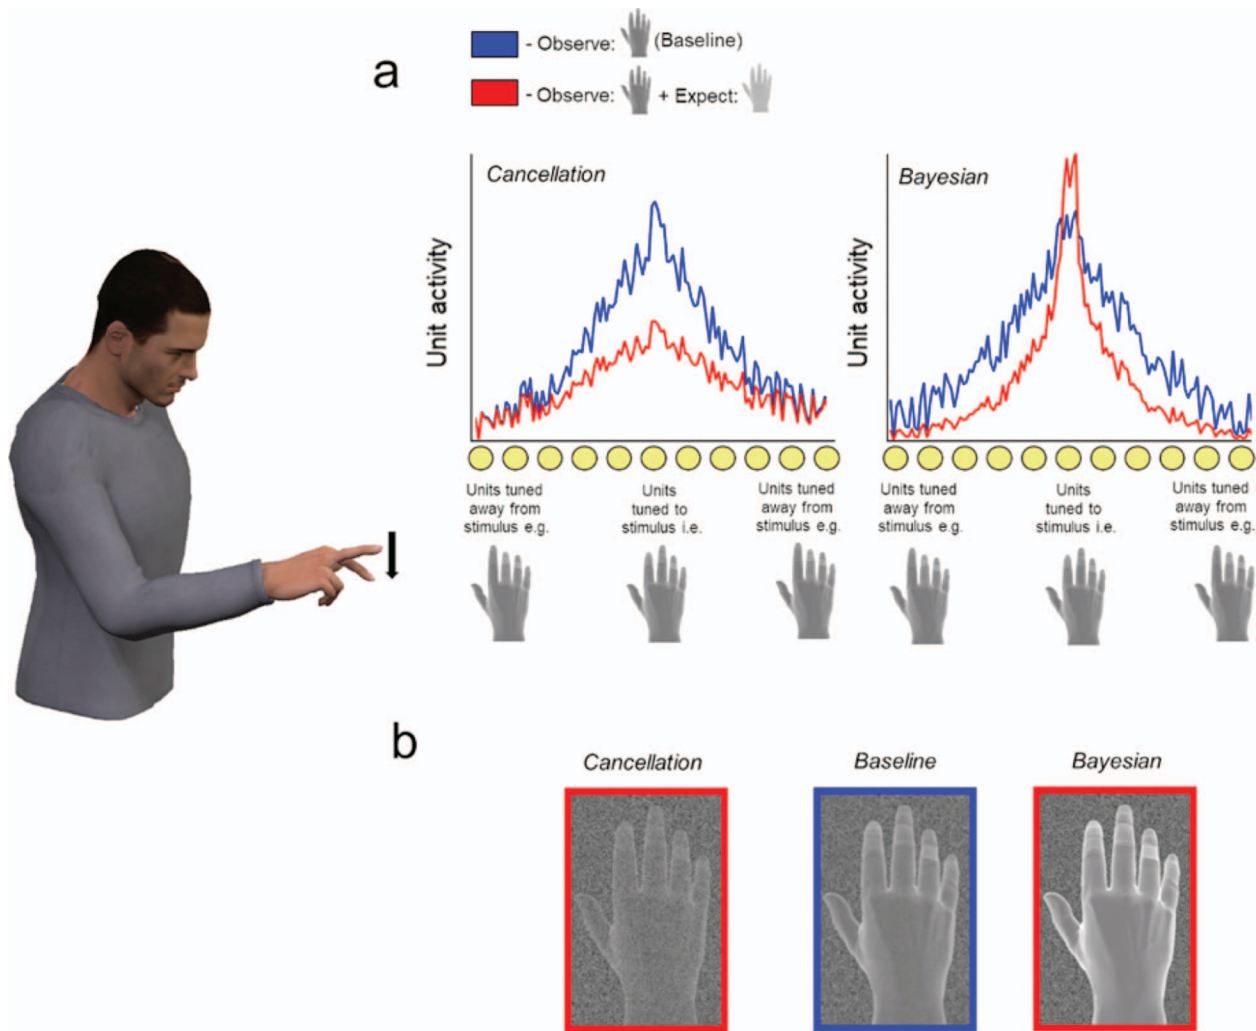


Figure 1. A schematic illustration of how predictive signals influence activation of sensory units under cancellation and Bayesian models, alongside their putative influences on perception. Cancellation models developed in the action literature (a, left) hypothesize that when we move (e.g., depress our index finger), we generate a predictive signal that suppresses activity in sensory units tuned to expected perceptual outcomes (e.g., visual units tuned to the sight of a hand with a depressed index finger). Weakening activity in these units reduces the signal-to-noise ratio of the sensory population—leading to less intense percepts and biasing observers away from perceiving these outcomes (b, left). In contrast, Bayesian models developed in the wider sensory cognition literature (a, right) suggest that predictive signals alter the weights on sensory channels such that the volume is “turned up” (e.g., the gain is increased) on expected relative to unexpected signals. Such weighting leads to a higher signal-to-noise ratio when expectations are valid, leading to more intense percepts and biases toward perceiving expected outcomes (b, right). See the online article for the color version of this figure.

ated with pathologies of agency in a variety of psychiatric conditions (Frith, Blakemore, & Wolpert, 2000).

However, the core principle guiding Cancellation models—that perception of predicted inputs is suppressed—contrasts with prominent Bayesian models in the wider sensory cognition literature (Figure 1b). These models suggest that we are more, not less, likely to perceive what we expect (de Lange, Heilbron, & Kok, 2018; Press & Yon, 2019). They emphasize how in an inherently ambiguous sensory world, it is adaptive for organisms to combine sampled sensory evidence with prior knowledge about what is likely to occur. Mechanistically this can be achieved by altering the weights on sensory

channels, by increasing the “gain” of expected relative to unexpected signals (de Lange et al., 2018; Summerfield & de Lange, 2014). Increasing the gain afforded to expected sensory signals in this fashion—effectively “turning up the volume” of events that conform to our prior predictions—would bias perceptual processing and predispose observers to perceive events that they expect to occur (e.g., Hudson, Nicholson, Ellis, & Bach, 2016; Hudson, Nicholson, Simpson, Ellis, & Bach, 2016; Wyart, Nobre, & Summerfield, 2012). For example, a range of curious illusory phenomena—such as the tendency of observers to perceive concave faces as convex (Gregory, 1997)—could arise via such mechanisms. Importantly, while these

kinds of biases may lead to occasional misperceptions, they may nonetheless reflect an adaptive and efficient way of generating veridical percepts, since expected events are, by definition, more likely to occur. While these models have been developed outside of action contexts, a mechanism that biases perception in line with expectations could be just as adaptive during action. For example, if we are trying to flick the light switch in a darkened room, we will generate more veridical estimates of our ongoing actions if we are biased to perceive expected events (e.g., the sight of a moving hand).

Cancellation and Bayesian accounts of how predictions shape perception have been difficult to compare directly because experimental approaches differ between disciplines (Press, Kok, & Yon, 2020b). Support for Bayesian models within the normative sensory cognition literature typically examines an organism's ability to detect a low intensity stimulus (Stein & Peelen, 2015; Wyart et al., 2012). In contrast, action studies reporting cancellation have typically asked participants to judge the intensity of action outcomes (Bays et al., 2005; Blakemore et al., 1998; Kilteni et al., 2019; Sato, 2008; Weiss, Herwig, & Schütz-Bosbach, 2011; see also Yon & Press, 2017, 2018). To render the paradigms more comparable, the present series of perceptual experiments used a signal detection paradigm within the domain of action (see also Cardoso-Leite, Mamassian, Schütz-Bosbach, & Waszak, 2010; Schwarz, Pfister, Kluge, Weller, & Kunde, 2018). They also presented visual action outcomes, given that both Cancellation and Bayesian theories hypothesize comparable operation of mechanisms across sensory modalities (Brown, Adams, Parees, Edwards, & Friston, 2013; Wolpert, Doya, & Kawato, 2003) and that visual events have been used more commonly in the normative sensory cognition literature (see General Discussion).

Participants produced manual actions (abducting their index or middle finger) and detected visual action outcomes, which could be congruent or incongruent with their own movement. This congruency manipulation exploits the fact that congruent action outcomes will be more expected than incongruent ones, based either on inherited evolutionary expectations or our extensive experience of controlling our actions (Hommel et al., 2001). Under Cancellation models, suppressing activity in units tuned to expected stimuli should make it harder for predictable action outcomes to reach detection threshold (see Figure 1), making observers either less sensitive to these events or biased to report that they did not occur. In contrast, under normative Bayesian models selectively increasing the relative weight on expected sensory channels should either bias observers to report that congruent events occurred or make them more sensitive to congruent outcomes. However, given recent findings in the broader sensory cognition literature, we anticipated that sensitizing sensory channels tuned to expected events would increase hits and also false alarms—in signal detection terms, biasing observers rather than increasing their sensitivity (Wyart et al., 2012). We subsequently used computational modeling to pinpoint which aspect of perceptual decision making is influenced by expectations.

Experiment 1: How do Actions Influence Detection of Expected Outcomes?

Method

Participants. Twenty-four healthy participants (19 female, 5 male, mean age = 24.9 years, $SD = 5.38$) took part in Experiment 1. All participants in all experiments reported normal or corrected-

to-normal vision and no history of psychiatric or neurological illness. A sample size of 24 participants per experiment was selected such that each would have at least 80% power to detect a medium-sized effect of action-outcome congruency on perceptual decisions (Cohen's $d_z = .6$, $N = 24$, $\alpha = .05$ provides 80.3% power - G*Power 3.1.9.2). All experiments were approved by the local ethics committee at Birkbeck, University of London.

Procedure. The experiment took place in a dimly lit testing cubicle. Participants sat ~55 cm from the monitor (153 × 32 cm, 60 Hz) used for stimulus presentation, with their hands placed above two keypads. The participant's right hand was rotated 90°, such that their knuckles were aligned with the body midline. Each trial began with the presentation of a greyscale avatar hand (Poser 10, Smith Micro Software). This image remained on screen until participants executed either an index or middle finger tapping action—depressing the relevant key. Movements were freely selected, but the experiment ran until participants executed at least 100 of each type. On 50% of trials, participant's actions triggered a synchronous movement of the on-screen hand (signal present) that was displayed for 17 ms. On the remaining 50% of trials, the hand remained still (signal absent). On signal present trials, half of observed movements were congruent with the participant's own action (e.g., execute index tap, observe index tap), and half were incongruent (e.g., execute index tap, observe middle tap). Regardless of signal presence, the index and middle region of the avatar hand was backward-masked by an oval texture comprised of avatar fingers (Cutting, Moore, & Morrison, 1988) for 100 ms. This image was in turn followed by a visual white noise mask presented for 300–600 ms. Participants were subsequently asked about the movement of one of the two fingers (e.g., did the INDEX finger move?). They registered their decision with a button press with their left thumb. On half of the trials, participants were probed about the congruent finger of the avatar hand (i.e. the index finger if they moved their index finger), and on the remaining half, the incongruent finger was probed. Participants in Experiment 1 also made a confidence judgment about their decision (“high confidence” or “low confidence”) to collect pilot data for an additional experiment.

Participants completed at least 200 trials.¹ Trial types were randomized across the experiment, and breaks were taken every 40 trials. Before the main experiment, participants completed a short practice block that familiarized them with the main task (16 trials). Participants subsequently completed a longer practice block without producing actions where they completed a 1 up–1 down adaptive staircase to adjust the difficulty of the perceptual discrimination such that it was approximately matched for all participants. This staircase targeted the amount of observed finger movement (minimum 1° rotation around the metacarpophalangeal joint, maximum 16°) that was required for detection on ~50% of trials. The staircase terminated after 12 reversals, and the average of the last six turning points was taken as an estimate of the participant's threshold. The main experiment began after completing the adaptive staircase, and test stimuli (when present) were shown at this threshold value on all subsequent trials.

¹ Participants in Experiment 1 were successful in following the instruction to execute roughly equal numbers of each action and therefore completed on average 218 trials ($SEM = 2.5$), comprising on average 49.8% congruent trials ($SEM = .0014$).

Results and Discussion

All tests in all experiments used an alpha level of .05, and for nonsignificant results we calculated Bayes Factors to quantify evidence for the absence of an effect (i.e. the null hypothesis; Dienes, 2014). Separate signal detection theoretic measures of sensitivity (d') and bias (c) were calculated using hit rates and false alarm rates on congruent and incongruent trials. d' reflects the extent to which participants are more likely to report the presence of a stimulus when it is present than when it is absent ($d' = z[\text{hit rate}] - z[\text{false alarm rate}]$), while c reflects the extent to which participants are more likely to respond “present” or “absent” regardless of objective stimulus presence ($c = -.05[z(\text{hit rate}) + z(\text{false alarm rate})]$). For some participants in some conditions, response counts were empty (e.g., no misses), which can preclude calculation of d' and c . In line with previous recommendations (Hautus, 1995), this issue was overcome by adjusting counts of hits, misses, false alarms, and correct rejections by $+.5$ in all experiments, and this adjustment was applied to all participants to avoid introducing biases into group-level analyses (Snodgrass & Corwin, 1988).

Sensitivity and bias on congruent and incongruent trials were compared using t tests. These analyses revealed that participants were more liberal in reporting the presence of congruent action outcomes (lower $c - t_{23} = 2.35, p = .028, dz = .480$; see Figure 2) and were also more sensitive when judgments probed the congruent finger (higher $d' - t_{23} = 2.29, p = .031, dz = .467$; see Table 1). These findings are predicted by the Bayesian account and inconsistent with the Cancellation account.

Experiment 2: Dissociating Effects of Expectation and Attention on Detection Performance

Experiment 1 found that participants were more sensitive to (higher d') and biased to report the presence of (lower c) congruent action outcomes. However, one possibility is that these results reflect effects of “attention” rather than “expectation” per se. That is, in both laboratory tasks and natural settings, top-down expectations (i.e. what is likely to occur) are often confounded with top-down attention (i.e. what is relevant for task performance; Summerfield & Egner, 2016). While in our task movements of congruent fingers are just as probable as incongruent ones—making both types of event equally task-relevant—actors may have learned outside the laboratory to allocate top-down attention to congruent fingers as these are typically more relevant for controlling our actions (note of course that our logic also assumes that congruent movements are more expected due to learning outside of the experimental setting; see Introduction). This interpretation is particularly likely, given that when expectation and attention have been orthogonalized in the sensory cognition literature, the former has been found to generate biasing effects and the latter sensitivity effects (Wyart et al., 2012). Experiment 2 examined this possibility by orthogonally manipulating action-outcome congruency and task relevance with a new sample. This was achieved by introducing two separate cues to the task: a number cue indicating which action participants should perform (thereby manipulating expectations about outcomes) and an orthogonal arrow cue indicating which outcomes will be relevant for perceptual decisions (indicating what participants should attend to).

Method

Participants. A new sample of 24 participants took part in Experiment 2 (17 female, 7 male, mean age = 24.4 years, $SD = 4.23$). Data from one additional participant was lost due to a technical malfunction.

Procedure. The procedure of Experiment 2 was identical to Experiment 1 with the following changes. Trials began with the presentation of the neutral observed hand overlaid with an arrow cue. On 50% of trials, the arrow cue was valid, pointing to the finger that was subsequently probed. On 25% of trials the cue was invalid, pointing to the finger that was not subsequently probed. On the remaining 25% of trials, the cue was neutral, pointing to both fingers and indicating that they would be probed with equal probability. After 700 ms, an imperative cue (“1” or “2”) was presented above the arrow, indicating which action participants were required to perform (index or middle tap, respectively). After participants executed the correct action, the same stimulus sequence was triggered as in Experiment 1. Participants made the same detection judgments, but confidence judgments were not collected in this experiment. The experiment comprised 320 trials. Trial types were randomized across the experiment, and breaks were taken every 20 trials.

Results and Discussion

Signal detection theoretic measures of sensitivity and bias were calculated for each combination of action-outcome congruency (congruent, incongruent) and cue validity (valid, neutral, invalid), and effects were evaluated using ANOVAs with the same factorial structure. Both analyses revealed an effect of cue validity, such that participants were more sensitive to (higher $d' - F_{2,46} = 4.246, p = .020, \eta_p^2 = .156$; Table 1) and more liberal in reporting (lower $c - F_{2,46} = 14.57, p < .001, \eta_p^2 = .388$; Figure 2) validly cued events—validating that participants used these cues to guide their attention (Hawkins et al., 1990). T tests decomposing these cuing effects found that judgments were more sensitive when cues were valid than invalid ($t_{23} = 2.85, p = .009, dz = .581$), though differences between sensitivity on valid and neutral cuing trials ($t_{23} = 1.98, p = .059, dz = .404, BF_{01} = 1.36$) or neutral and invalid trials ($t_{23} = .937, p = .359, dz = .191, BF_{01} = 3.01$) were nonsignificant. Comparable tests found observers were more liberal on validly cued trials than invalid ones ($t_{23} = 4.11, p < .001, dz = .838$) and more liberal on neutrally cued than invalidly cued trials ($t_{23} = 4.81, p < .001, dz = .981$), but differences between valid and neutral cuing trials were not significant ($t_{23} = 1.28, p = .214, dz = .261, BF_{01} = 2.21$). Moreover, effects of cue validity on judgment sensitivity did not interact with action-outcome congruency ($F_{2,46} = .184, p = .833, \eta_p^2 = .008, BF_{01} = 7.19$).

Importantly, these analyses also found that participants were biased to report the presence of congruent action outcomes ($F_{1,23} = 8.47, p = .008, \eta_p^2 = .269$)², and the magnitude of this bias did not interact with the focus of attention ($F_{2,46} = 1.29, p = .286, \eta_p^2 = .053, BF_{01} = 5.95$), suggesting that observers are biased to perceive predictable action outcomes irrespective of task relevance. However, there was no significant effect of congruency on sensitivity ($F_{1,23} = 1.70, p = .205, \eta_p^2 = .069, BF_{01} = 1.52$).

² In Experiment 2, one participant was an outlier, showing an especially large action-induced bias (z score = 3.76). Re-running the same analyses without data from this participant revealed the same effect of action congruency on c values ($p = .003$), the same correlation between congruency effects and modeled drift biases ($r = -.924, p < .001$) and did not change any other statistical patterns observed.




Figure 2. Action execution and detection task, with signal detection c results from Experiments 1–3. a. Participants performed actions, which were paired with synchronous congruent or incongruent movements of an avatar hand that they were required to detect. In Experiments 2 and 3, attentional arrow cues also informed participants about which finger of the avatar hand was likely to be probed. b. We calculated the signal detection theoretic measure c to index biases induced by action. These values were lower (i.e. responses were more liberal) on congruent (saturated) relative to incongruent (desaturated) trials, irrespective of attentional focus. This effect demonstrates that perceptual decisions were biased toward expected action outcomes. Error bars show 95% within-participant confidence intervals of the mean difference between conditions. See the online article for the color version of this figure.

Therefore, when attention is orthogonally manipulated, participants are still biased (lower c) to report the presence of congruent action outcomes. The effect of congruency on sensitivity was not detectable in Experiment 2, consistent with the possibility that this effect in Experiment 1 was determined by attentional processes (Wyart et al., 2012).

Experiment 3: Controlling Imperative-Outcome Mapping

Experiment 3 investigated whether the congruency biasing effect was driven by expectations engendered by action or a mapping between imperative cues and stimulus types (1 = index/left; 2 =

Table 1
Mean (SD) Sensitivity (d') Values Across All Conditions in Experiments 1–3

Condition	Congruent			Incongruent		
	Valid	Neutral	Invalid	Valid	Neutral	Invalid
Experiment 1	1.42 (.567)			1.28 (.588)		
Experiment 2	0.73 (.532)	0.68 (.539)	0.63 (.455)	0.69 (.536)	0.59 (.537)	0.57 (.521)
Experiment 3	1.14 (.743)	1.05 (.684)	0.90 (.590)	1.14 (.680)	0.993 (.591)	0.923 (.555)

middle or right—a common mapping used in musical training). To rule out possible mapping between cues and outcomes, we compared one group of participants on an identical procedure as Experiment 2 to another group who received arbitrary shapes as imperatives rather than numbers (NB: no cue-outcome mapping can be learnt within the experiment as there is no within-experiment contingency).

Method

Participants. A new sample of 48 participants took part in Experiment 2 (33 female, 15 male, mean age = 24.7 years, $SD = 4.08$), with 24 in each of the two groups. This provides the same power to detect the effect individually within each of the two groups as in Experiments 1 and 2, as well as enabling examination of the interaction according to cue type.

Procedure. One group of 24 participants completed a procedure identical to Experiment 2, where numbers (“1” or “2”) cued participants to perform index or middle finger actions. A separate group of 24 participants completed a near identical procedure, except circles and squares indicated to participants that they should execute index and middle finger movements. Half of these participants received a circle-index/square-middle mapping, and half received a circle-middle/square-index mapping.

Results and Discussion

Measures of sensitivity and bias were calculated separately for each combination of experimental conditions and analyzed with separate ANOVAs. Therefore, the only change with respect to Experiment 2 was the addition of a between-participants factor of cue type.³ Greenhouse Geisser corrections were employed where appropriate. Participants were again more liberal in reporting the presence of congruent than incongruent action outcomes ($F_{1,46} = 14.59, p < .001, \eta_p^2 = .241$; Figure 2) and more liberal in reporting the presence of validly cued events ($F_{2,92} = 5.062, p = .008, \eta_p^2 = .009$). T tests decomposing the attentional cuing effect on c values found that observers were more liberal on validly cued trials than invalid ones ($t_{47} = 2.73, p = .009, dz = .394$) and more liberal on neutrally cued than invalidly cued trials ($t_{47} = 2.47, p = .017, dz = .356$), but differences between valid and neutral cuing trials were not significant ($t_{47} = 1.19, p = .242, dz = .171, BF_{01} = 3.30$).

The effect of action-outcome congruency did not interact with the focus of attention ($p = .080, \eta_p^2 = .025, BF_{01} = 3.81$). Crucially, cue type (shapes vs. numbers) also did not interact with the factor action-outcome congruency ($F_{1,46} = 1.192, p = .281, \eta_p^2 = .025, BF_{01} = 2.61$) or generate a three-way interaction with congruency and attentional focus ($F_{2,92} = .234, p = .792, \eta_p^2 =$

$.005, BF_{01} = 7.69$). Indeed, separate analyses of the number ($F_{1,23} = 7.50, p = .012, \eta_p^2 = .246$) and shape cuing conditions ($F_{1,23} = 8.02, p = .009, \eta_p^2 = .259$) revealed a significant effect of congruency in both groups—underscoring that participants were more likely to report the presence of congruent action outcomes irrespective of cue type.

Equivalent analyses of d' found that judgment sensitivity again improved with valid cues ($F_{2,92} = 9.87, p < .001, \eta_p^2 = .257$; Table 1). T tests decomposing this cuing effect found judgments were more sensitive when cues were valid compared to invalid ($t_{47} = 4.49, p < .001, dz = .648$), valid compared to neutral ($t_{47} = 2.42, p = .019, dz = .349$), and neutral compared to invalid ($t_{47} = 2.07, p = .044, dz = .298$). However, d' was unaffected by action-outcome congruency ($F_{1,46} = .092, p = .763, \eta_p^2 = .002, BF_{01} = 7.63$), and effects of relevance did not interact with congruency ($F_{2,92} = .638, p = .531, \eta_p^2 = .014, BF_{01} = 9.80$). Neither the main effect of validity ($F_{2,92} = .194, p = .824, \eta_p^2 = .004, BF_{01} = 11.7$) nor any interaction with congruency ($F_{2,92} = .542, p = .583, \eta_p^2 = .012, BF_{01} = 6.21$) was affected by cue type. Experiment 3 therefore replicated the findings from Experiment 2 while controlling for the particular imperative mapping, suggesting that it was the action-based expectations that generated the biasing effects.

Combined Analysis

No significant congruency effects on d' were found in Experiments 2 and 3, in contrast to Experiment 1 where observers were more sensitive to the presence of congruent action outcomes. As noted, one plausible explanation for this discrepancy is that effects on sensitivity are driven by attentional mechanisms (see Wyart et al., 2012), and sensitivity effects were correspondingly abolished when attention was orthogonally cued in Experiments 2 and 3. However, Bayes Factors suggested evidence for a null result was convincing in Experiment 3 ($BF_{01} > 3$) but indecisive in Experiment 2 ($BF_{01} < 3$). Combining data from Experiments 2 and 3 for greater precision, we found convincing evidence for the absence of a congruency effect on d' ($t_{71} = 1.02, p = .312, dz = .120, BF_{01} = 4.69$, suggesting that sensitivity effects were absent across the two experiments).

³ In Experiment 3, one participant was an outlier, showing an especially large action-induced bias (z score = 4.12). Rerunning the same analyses without data from this participant revealed the same effect of action congruency on c values ($p < .001$), the same correlation between congruency effects and modeled drift biases ($r = -.807, p < .001$) and did not change any other statistical patterns observed.

Computational Modeling: Which Aspect of Perceptual Decision Making Is Biased by Action?

We used computational modeling of participant choices and reaction time (RT) distributions to pinpoint which aspect of perceptual decision making is biased by action. Drift diffusion models (DDMs) of perceptual decision making have enjoyed growing prominence in the cognitive sciences (Ratcliff, Smith, Brown, & McKoon, 2016; see Figure 3). These models assume that when making perceptual choices (e.g., was a stimulus present or not?), observers have an internal representation of “sensory evidence” that is sampled by decision circuits. Decision circuits continuously sample from the representations of sensory evidence, and when the accumulated decision variable meets a response boundary (e.g., ‘respond present’), the appropriate response is triggered. The representation of sensory evidence is therefore separable from the representation of decisions about that evidence.

There are two ways that the DDM could accommodate the action-induced bias found in Experiments 1–3. First, expectations during action could shift the starting point of the evidence accumulation process toward the “respond present” decision bound (varying parameter z of the DDM; see Figure 3a). Such effects are often thought to reflect biases in the decision circuits. For example, predictive cues can induce preparatory motor activity before a stimulus is presented (de Lange, Rahnev, Donner, & Lau, 2013), and these starting point effects could operate even if participants are insensitive to the perceptual information (e.g., they closed their eyes). However, a second alternative is that action directly biases sensory representations, as though agents have selectively increased the gain (or “precision”) afforded to expected sensory signals (Friston, 2018). This kind of bias would manifest as an asymmetric bias in the rate of evidence accumulation (Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; “drift biasing,” affecting parameter db of the DDM; see Figure 3b) because, as the sensory evidence is sampled more, it provides progressively more evidence in favor of the expected event.

Method

To investigate whether either of these possibilities could account for the action-induced bias we observed in our experiments, we fit hierarchical DDMs to participant choice and RT data using the hDDM package implemented in Python (Wiecki, Sofer, & Frank, 2013). In the hierarchical DDM, model parameters for each participant are treated as random effects drawn from group-level distributions, and Bayesian Markov Chain Monte Carlo (MCMC) sampling is used to estimate group and participant level parameters simultaneously.

We specified four different models for data in each experiment: (a) a *null* model where no parameters were permitted to vary between congruent and incongruent trials; (b) a *start bias* model where the start point of evidence accumulation (z) could vary on congruent and incongruent trials; (c) a *drift bias* model where a constant bias in evidence accumulation (db) could vary between congruent and incongruent trials; (d) a *start + drift bias* model where both parameters could vary. To improve model fits, all models in Experiments 2 and 3 also allowed global drift rate (but not drift bias—i.e. biases to drift asymmetrically toward present or absent decisions) to vary as a function of cue validity to account

for the effect of attentional cues on d' (Ratcliff et al., 2016). Varying drift rate (rather than drift bias) captures effects that reflect more reliable evidence accumulation to the appropriate response boundary (i.e. “respond present” when stimuli are present and “respond absent” when absent) and therefore accounts for sensitivity effects seen as a function of cue validity, rather than a biased accumulation toward one response boundary over another. In no model did we allow drift rate to vary between congruent and incongruent trials, as this would amount to assuming that observers are generally more sensitive to all kinds of sensory evidence on congruent trials (e.g., higher d'), rather than sensitizing particular sensory channels that encode expected outcomes.

All models were estimated with MCMC sampling with 30,000 samples (“burn-in” = 7500). Model convergence was assessed by inspecting chain posteriors and simulating RT distributions for each participant. Models were compared using deviance information criteria (DIC) as an approximation of Bayesian model evidence. Estimated parameters in each model were compared using the Bayesian significance test implemented in hDDM, which computes the posterior probability that group-level parameters differ across conditions.

Results and Discussion

Fitting the DDM to the behavioral data found—in all experiments—that the drift biasing model provided a better fit than the start biasing model and that in Experiments 2–3 a model only implementing a drift bias outperformed a model implementing both biases (see Figure 3c). Analyzing modeled parameters revealed higher drift biases on congruent relative to incongruent trials (posterior probabilities that congruent $db >$ incongruent db : Experiment 1 = 0.819, Experiment 2 = 0.959, Experiment 3 = 0.899—higher values indicate greater differences between conditions; Wiecki et al., 2013). To confirm that these differences in drift bias explained the effect of expectations on perceptual decisions, we calculated and correlated the difference between drift bias parameters (congruent $db -$ incongruent db) and the magnitude of the behavioral bias (congruent $c -$ incongruent c) for each participant. This analysis revealed strong relationships within all three samples (Experiment 1: $r_{24} = -.877, p < .001$; Experiment 2: $r_{24} = -.973, p < .001$; Experiment 3: $r_{48} = -.855, p < .001$, see Figure 3d).

This modeling thereby suggests that action-induced biases were best accounted for by a sensory drift biasing mechanism—where observers are biased to accumulate sensory evidence in line with their expectations—rather than a change in later decision circuits.

General Discussion

Cognitive scientists have proposed models of perceptual prediction that disagree about how our expectations should shape what we perceive (Press et al., 2020b). Cancellation models influential in action control have suggested that agents are less likely to perceive the predictable consequences of their actions (Bays & Wolpert, 2007), in contrast to Bayesian models from the wider sensory cognition literature that suggests observers weight perception toward prior knowledge (see de Lange, Heilbron, & Kok, 2018; Press & Yon, 2019). The present experiments suggest that evidence accumulation is biased in line with expectations during

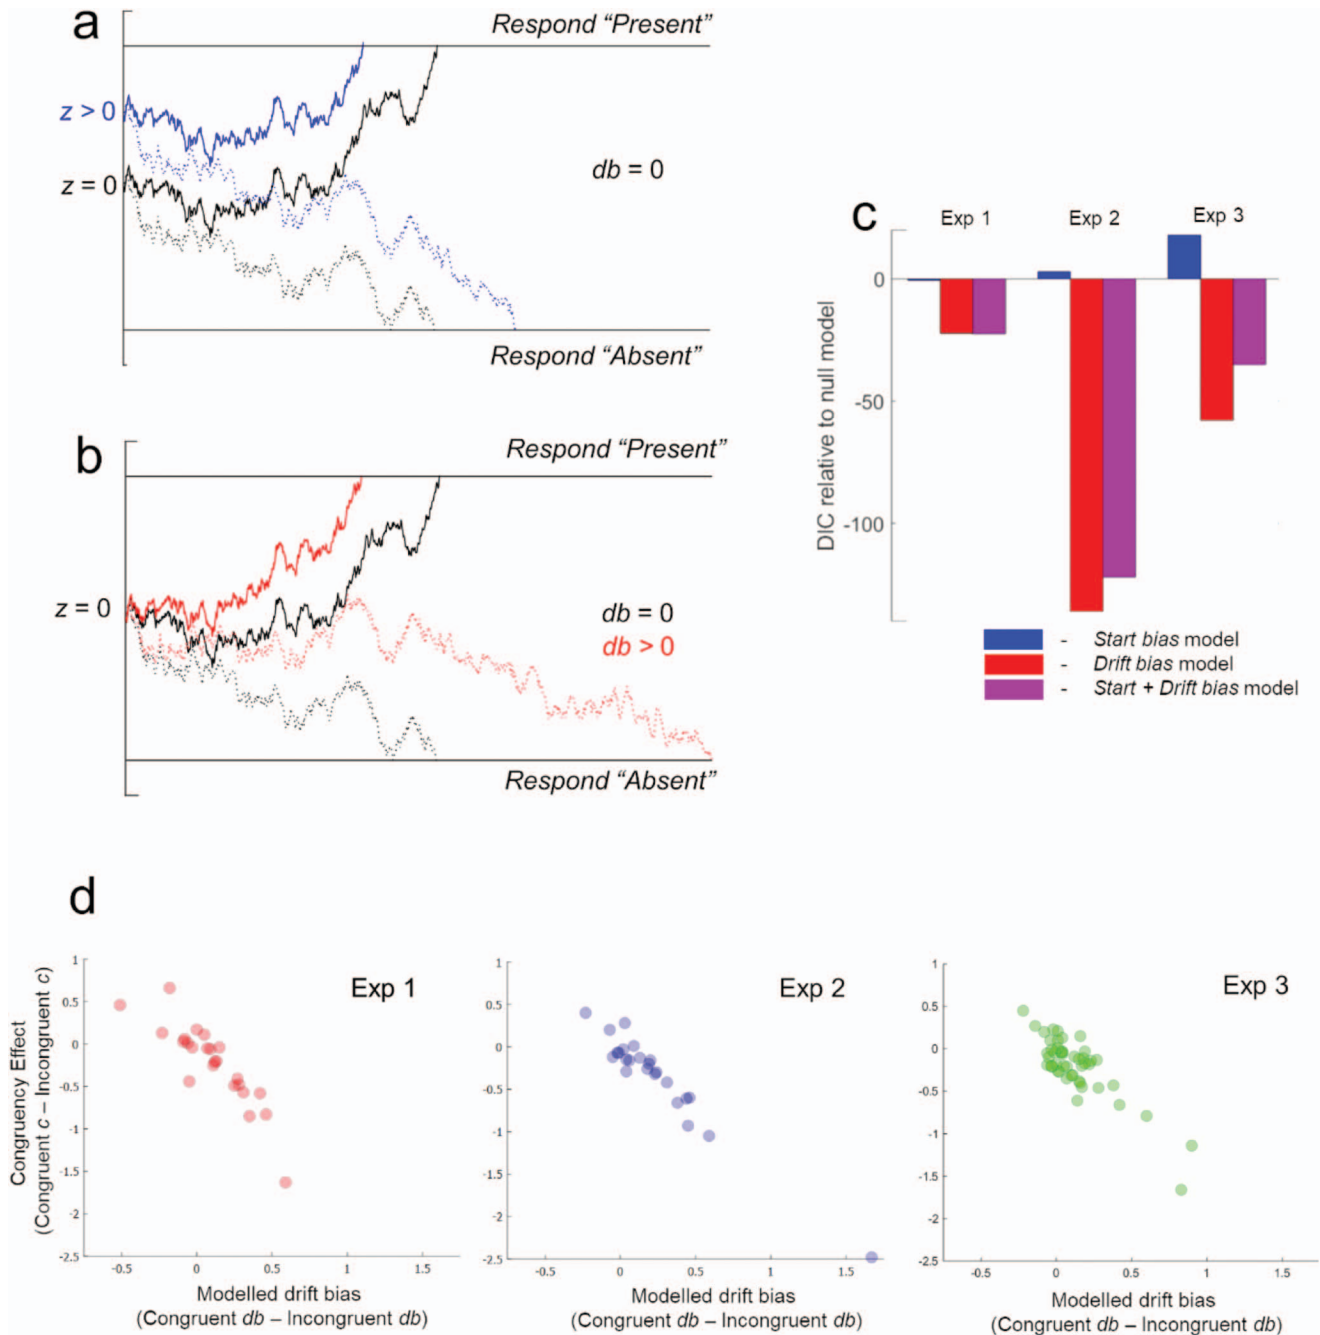


Figure 3. Illustration of how the DDM could explain action-induced biases, and results of computational modeling. a. For an unbiased decision process (black lines) sensory evidence integrates toward the upper response boundary when stimuli are present (solid lines) and toward the lower response boundary when stimuli are absent (dotted lines). Baseline shifts in decision circuits could shift the start point of the accumulation process nearer to the upper boundary for congruent events (influencing the parameter z ; blue lines – *Start bias model*). b. Alternatively, selectively altering the weights on sensory channels could bias evidence accumulation in line with expectations (influence parameter db ; red lines – *Drift bias model*). c. Across all experiments, the *Drift bias* model provided a better fit than the *Start bias* model (lower deviance information criteria [DIC] indicates better model fit). d. Moreover, in each experiment, there was a strong correlation between the drift bias values modeled to each participant and the empirical action-induced bias on c values. See the online article for the color version of this figure.

action such that observers are more likely to perceive the outcomes they expect. This pattern concords with Bayesian accounts of expectation developed in the wider sensory cognition literature, which assume observers increase the weight they give to expected inputs when making perceptual judgments. Biasing perceptual decisions in this fashion is an effective way to rapidly generate more veridical percepts from sensory signals corrupted by irreducible internal and external noise. While this process increases the likelihood that expected signals are detected (i.e. more hits), it also makes observers prone to hallucinate events when confronted with signal-like noise (i.e. more false alarms; Wyart et al., 2012) and therefore is thought to generate biasing rather than sensitivity effects—as observed here. Indeed, we found these action-induced biases in perception were dissociable from top-down attention, with orthogonal task-relevance cues altering judgment sensitivity. This finding is in line with previous studies of expectation and attention outside of action contexts (Wyart et al., 2012) and perhaps therefore indicative of domain-general mechanisms.

These patterns are consistent with findings of how previous decisions about sensations influence subsequent decisions (Talluri, Urai, Tsetsos, Usher, & Donner, 2018) and with a recent fMRI study that found that expected action outcomes are more readily decoded from early and late visual brain areas (Yon, Gilbert, de Lange, & Press, 2018). These neural effects are predicted by Bayesian models, but it has alternatively been suggested that effects of expectation in sensory brain areas could reflect feedback from decision-related areas in higher-level cortex that have no causal effect on perception (Bang & Rahnev, 2017; Choe, Blake, & Lee, 2014). The current findings importantly indicate effects of action expectation on perception and a corresponding drift biasing mechanism that is consistent with an early sensory biasing account.

The predictive relationship exploited in these experiments—for example, that observed index finger movements are an expected consequence of moving one's index finger—is strong and stable. It reflects an expectation that is likely acquired through our extensive experience of controlling our actions (Hommel et al., 2001). However, we would predict that in principle the same underlying mechanisms operate when we acquire new predictions. In line with this assumption, our effects are akin to those found—outside of action settings—when participants learn within an experiment that color cues predict the orientation of gratings (Wyart et al., 2012). The hypothesis that similar mechanisms operate when we acquire new predictions may appear inconsistent with previous conflicting findings when participants are presented with novel action-outcome mappings within an experiment. In an elegant training study, Cardoso-Leite et al. (2010) found evidence that participants are less sensitive to grating orientations “predicted” by actions that had been paired with the gratings during training, with no influence of actions on bias measures. However, sensitivity was especially high in this study, such that a bias toward perceiving the expected, accompanied by ceiling effects on the hit rate, would appear as a sensitivity reduction. Perhaps more importantly, this single study used a small sample, and the findings may not replicate (Schwarz et al., 2018). It may therefore be suggested that there was insufficient opportunity in this paradigm to acquire the action-outcome mappings reliably. However, in principle, given sufficient opportunity for prediction acquisition, we would hypothesize prediction mechanisms to operate in a qualitatively similar

fashion when predictions are both “old” and “new” (Dogge, Custers, Gayet, Hoijtink, & Aarts, 2019).

This pattern of results is difficult to reconcile with cancellation models and their central claim that observers are less likely to perceive the multisensory predictable consequences of their movements (Bays & Wolpert, 2007; Blakemore et al., 1998). For example, key support for cancellation models has come from studies that show predictable signals generated by action are perceived to be less intense than similar (i.e. unpredicted) events presented in the absence of action (Bays et al., 2005). These studies are in fact often difficult to interpret because there are several differences between the predicted and unpredicted conditions. Most notably, many of the studies compare perception of “predicted” self-generated events with perception of “unpredicted” sensory events generated by external sources while participants themselves remain passive—or where the sensory and motor events overlap less due to temporal misalignment (Blakemore, Frith, & Wolpert, 1999). This comparison perhaps confounds the operation of expectation mechanisms with that of other processes (see Press, Kok, & Yon, 2020a). For example, if conceptualizing action as an additional task, classic working memory models would hypothesize reduced sensory processing when events are presented in combination with action (Baddeley, 1996). It is therefore difficult to isolate effects of prediction mechanisms from those introduced by the dual- versus single-task design. The measures employed typically also cannot isolate perceptual effects from decisional biases (see Firestone & Scholl, 2016) or expectation from attention-based processes (see Summerfield & Egner, 2016).

One explanation for the difference between the present and previous studies could be that we have examined visual action effects, whereas much evidence to support cancellation comes from tactile paradigms (e.g., Juravle, McGlone, & Spence, 2013; Juravle, Binsted, & Spence, 2017; Kiltani & Ehrsson, 2017). The models of which we are aware hypothesize similar operation of predictive Bayesian or Cancellation mechanisms across domains (Brown et al., 2013; Wolpert et al., 2003), and there may be no reason to hypothesize distinct adaptive arguments for the sense of touch. However, touch may be influenced differently because of operation of additional generalized “suppression” mechanisms. Such suppression mechanisms are thought to attenuate tactile sensations during action regardless of whether they are predicted effects of action or not and are thought to be mediated by spinal mechanisms (Seki & Fetz, 2012), and comparable mechanisms may similarly attenuate sensory processing in a nonpredictive fashion across modalities in humans and other animals (Crapse & Sommer, 2008). Importantly, recent experiments in touch suggest that when confounds related to sensory suppression are removed, action predictions may influence perception in a qualitatively similar fashion irrespective of sensory modality (Thomas, Yon, de Lange, & Press, 2020).

Nevertheless, there remain studies that are less prone to these alternative explanations and report “cancelled” percepts for expected, relative to unexpected, action outcomes (e.g., Roussel, Hughes, & Waszak, 2013, 2014). It is therefore important that future research unpacks the features that have driven previous reports of “cancellation”—including the possibility that additional mechanisms are recruited as sensory processing unfolds (see Press et al., 2020b). Specifically, some of us have recently proposed that

the influence of expectation on perception may not be as simple as suggested in either Bayesian or Cancellation theories. Under this proposal, observers are initially biased toward perceiving expected sensory events, but any especially surprising inputs—likely relevant for model updating—are reactively upweighted (Press et al., 2020b). Such a reactive upweighting process only for a subset of “unexpected” events is consistent with evidence from the learning and inference literature and is in line with the proposed adaptive function of a cancellation mechanism—that is, highlighting events that are informative to the organism. This account may explain some discrepancies in the literature—for example, finding perceptual upweighting of rapid, punctate consequences of movement, in contrast with evidence for cancellation-like phenomena when action outcomes unfold dynamically (e.g., Lally, Frendo, & Diedrichsen, 2011). Future studies should therefore establish whether expectation effects are modulated by the time at which perception is probed and the extent of the surprise elicited by an “unexpected” stimulus—which is low with the barely detectable sensations used in the present signal detection task (Press et al., 2020b).

The concept of cancellation has had a wide-ranging influence on research investigating the sense of agency and its aberration in psychiatric disease. For example, experiences of passivity in schizophrenia—where patients feel moved by an external force—have often been attributed to a failure of predictive cancellation, which causes self-produced action outcomes to appear unusually intense (Frith et al., 2000). However, our results suggest that sensorimotor predictions can increase the weight we give to expected sensory signals, which may suggest that these mechanisms contribute to the experience of agency in a different way. Specifically, biases toward perceiving expected outcomes may help observers to overcome sensory noise, giving us higher fidelity representations of our ongoing movements and their consequences. An inability to incorporate this kind of top-down knowledge into perceptual estimates could leave agents with higher levels of uncertainty about their actions, leaving them vulnerable to developing the unusual beliefs that characterize psychosis (Fletcher & Frith, 2009).

Important contributions to work on sensory prediction during action have also come from models of predictive processing and active inference developed in computational neuroscience (Friston, 2005). These models suggest that all aspects of perception, action, and cognition arise as agents minimize the mismatch between their models of the extracranial world and incoming sensory evidence. Notably, this can be achieved either by using the evidence to update the models (i.e. perception) or by using action to change the evidence (i.e. to alter the world so it conforms to the model). In principle, the architecture of these models can account for the tendency of agents to up- or down-weight perception of predictable action outcomes through changes in the gain or “precision” afforded to sensory evidence (Friston, 2008; Yon, de Lange, & Press, 2019). For example, researchers using this framework have suggested that sensory precision on certain channels is attenuated during action—explaining “cancellation” phenomena (Brown et al., 2013; Van Doorn, Paton, Howell, & Hohwy, 2015)—while also suggesting that observers can increase the precision afforded to expected sensory signals—leading to “representational sharpening” and biases toward perceiving what we expect (Friston, 2018). However, while models of precision-weighting during ac-

tion have important advantages over forward-model based accounts (e.g., where it has been difficult to specify computationally how predictions could “cancel” sensory signals—see Brown et al., 2013), it remains the case that the precision on a sensory channel, and perception of a single event, cannot be up- and down-weighted simultaneously. This makes it difficult to use existing models of active inference to make precise predictions about those conditions under which perception of expected action outcomes should be enhanced and those where it should be attenuated. Our present findings may provide some constraints on such models that aid the generation of more precise behavioral predictions in the future, as these results suggest that the precision afforded to visual channels encoding expected outcomes is augmented rather than attenuated—certainly when sensory events coincide temporally with action initiation. We hope that behavioral experiments like ours will play an important role in informing and constraining such models—elaborating a comprehensive account of how and why sensory gain is augmented and attenuated as we act upon the world around us.

In conclusion, these results have shown that observers are biased toward perceiving the expected outcomes of their movements. These findings are difficult to reconcile with dominant cancellation accounts in action control, but concord well with normative models of Bayesian perceptual inference. Namely, increasing the weight we give to expected sensory signals may explain how we develop a largely veridical representation of our actions and their consequences in an inherently uncertain sensory world.

Context

Previous work in the lab has investigated whether specific types of sensorimotor representations—that is, visual-motor mirror representations of action—are the product of domain-general statistical learning processes. After finding broad support for this idea (e.g., Cook, Bird, Catmur, Press, & Heyes, 2014; Press et al., 2012), we began to investigate whether domain-general models can explain the functional influence of sensorimotor predictions on perception as well as the origin of such representations. In a recent neuroimaging study (Yon et al., 2018) we found representations in visual brain areas were reshaped toward expected action outcomes, which is more in line with domain-general ideas in the broader sensory cognition literature than the cancellation models in action. However, while these neuroimaging data are consistent with these models, it has been suggested that expectation-related activity in sensory brain areas plays no causal role in determining what we perceive (Bang & Rahnev, 2017; Choe et al., 2014). Here we directly investigated how expectations influence the perception of action outcomes—finding evidence that we are biased toward perceiving expected action outcomes.

References

- Baddeley, A. (1996). Exploring the central executive. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 49, 5–28. <http://dx.doi.org/10.1080/713755608>
- Bang, J. W., & Rahnev, D. (2017). Stimulus expectation alters decision criterion but not sensory signal in perceptual decision making. *Scientific Reports*, 7, 17072. <http://dx.doi.org/10.1038/s41598-017-16885-2>
- Bays, P. M., & Wolpert, D. M. (2007). Computational principles of sensorimotor control that minimize uncertainty and variability. *The*

- Journal of Physiology*, 578, 387–396. <http://dx.doi.org/10.1113/jphysiol.2006.120121>
- Bays, P. M., Wolpert, D. M., & Flanagan, J. R. (2005). Perception of the consequences of self-action is temporally tuned and event driven. *Current Biology*, 15, 1125–1128. <http://dx.doi.org/10.1016/j.cub.2005.05.023>
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (1999). Spatio-temporal prediction modulates the perception of self-produced stimuli. *Journal of Cognitive Neuroscience*, 11, 551–559. <http://dx.doi.org/10.1162/089892999563607>
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, 1, 635–640. <http://dx.doi.org/10.1038/2870>
- Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14, 411–427. <http://dx.doi.org/10.1007/s10339-013-0571-3>
- Cardoso-Leite, P., Mamassian, P., Schütz-Bosbach, S., & Waszak, F. (2010). A new look at sensory attenuation: Action-effect anticipation affects sensitivity, not response bias. *Psychological Science*, 21, 1740–1745. <http://dx.doi.org/10.1177/0956797610389187>
- Choe, K. W., Blake, R., & Lee, S. H. (2014). Dissociation between neural signatures of stimulus and choice in population activity of human V1 during perceptual decision-making. *The Journal of Neuroscience*, 34, 2725–2743. <http://dx.doi.org/10.1523/JNEUROSCI.1606-13.2014>
- Cook, R., Bird, G., Catmur, C., Press, C., & Heyes, C. (2014). Mirror neurons: From origin to function. *Behavioral and Brain Sciences*, 37, 177–192. <http://dx.doi.org/10.1017/S0140525X13000903>
- Crapse, T. B., & Sommer, M. A. (2008). Corollary discharge across the animal kingdom. *Nature Reviews Neuroscience*, 9, 587–600. <http://dx.doi.org/10.1038/nrn2457>
- Cutting, J. E., Moore, C., & Morrison, R. (1988). Masking the motions of human gait. *Perception & Psychophysics*, 44, 339–347. <http://dx.doi.org/10.3758/BF03210415>
- de Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, 22, 764–779. <http://dx.doi.org/10.1016/j.tics.2018.06.002>
- de Lange, F. P., Rahnev, D. A., Donner, T. H., & Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *The Journal of Neuroscience*, 33, 1400–1410. <http://dx.doi.org/10.1523/JNEUROSCI.1094-12.2013>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5, 781. <http://dx.doi.org/10.3389/fpsyg.2014.00781>
- Dogge, M., Custers, R., Gayet, S., Hoijsink, H., & Aarts, H. (2019). Perception of action-outcomes is shaped by life-long and contextual expectations. *Scientific Reports*, 9, 5225. <http://dx.doi.org/10.1038/s41598-019-41090-8>
- Fiehler, K., Brenner, E., & Spering, M. (2019). Prediction in goal-directed action. *Journal of Vision*, 19(9), 10. <http://dx.doi.org/10.1167/19.9.10>
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 39, e229. <http://dx.doi.org/10.1017/S0140525X15000965>
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews Neuroscience*, 10, 48–58. <http://dx.doi.org/10.1038/nrn2536>
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 360, 815–836. <http://dx.doi.org/10.1098/rstb.2005.1622>
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4, e1000211. <http://dx.doi.org/10.1371/journal.pcbi.1000211>
- Friston, K. (2018). Does predictive coding have a future? *Nature Neuroscience*, 21, 1019–1021. <http://dx.doi.org/10.1038/s41593-018-0200-7>
- Frith, C. D., Blakemore, S. J., & Wolpert, D. M. (2000). Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 355, 1771–1788. <http://dx.doi.org/10.1098/rstb.2000.0734>
- Gregory, R. L. (1997). Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 352, 1121–1127. <http://dx.doi.org/10.1098/rstb.1997.0095>
- Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d' . *Behavior Research Methods, Instruments & Computers*, 27, 46–51. <http://dx.doi.org/10.3758/BF03203619>
- Hawkins, H. L., Hillyard, S. A., Luck, S. J., Mouloua, M., Downing, C. J., & Woodward, D. P. (1990). Visual attention modulates signal detectability. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 802–811. <http://dx.doi.org/10.1037/0096-1523.16.4.802>
- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24, 849–878. <http://dx.doi.org/10.1017/S0140525X01000103>
- Hudson, M., Nicholson, T., Ellis, R., & Bach, P. (2016). I see what you say: Prior knowledge of other’s goals automatically biases the perception of their actions. *Cognition*, 146, 245–250. <http://dx.doi.org/10.1016/j.cognition.2015.09.021>
- Hudson, M., Nicholson, T., Simpson, W. A., Ellis, R., & Bach, P. (2016). One step ahead: The perceived kinematics of others’ actions are biased toward expected goals. *Journal of Experimental Psychology: General*, 145, 1–7. <http://dx.doi.org/10.1037/xge0000126>
- James, W. (1890). *The principles of psychology*. New York, NY: Henry Holt and Company.
- Juravle, G., Binsted, G., & Spence, C. (2017). Tactile suppression in goal-directed movement. *Psychonomic Bulletin & Review*, 24, 1060–1076. <http://dx.doi.org/10.3758/s13423-016-1203-6>
- Juravle, G., McGlone, F. P., & Spence, C. (2013). Context-dependent changes in tactile perception during movement execution. *Frontiers in Psychology*, 4, 913. <http://dx.doi.org/10.3389/fpsyg.2013.00913>
- Kilteni, K., & Ehrsson, H. H. (2017). Body ownership determines the attenuation of self-generated tactile sensations. *Proceedings of the National Academy of Sciences of the United States of America*, 114, 8426–8431. <http://dx.doi.org/10.1073/pnas.1703347114>
- Kilteni, K., Houborg, C., & Ehrsson, H. H. (2019). Rapid learning and unlearning of predicted sensory delays in self-generated touch. *eLife*, 8, e42888. <http://dx.doi.org/10.7554/eLife.42888>
- Lally, N., Frendo, B., & Diedrichsen, J. (2011). Sensory cancellation of self-movement facilitates visual motion detection. *Journal of Vision*, 11(14), 5. <http://dx.doi.org/10.1167/11.14.5>
- Mulder, M. J., Wagenmakers, E. J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience*, 32, 2335–2343. <http://dx.doi.org/10.1523/JNEUROSCI.4156-11.2012>
- Müsseler, J., & Hommel, B. (1997). Blindness to response-compatible stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 861–872. <http://dx.doi.org/10.1037/0096-1523.23.3.861>
- Press, C., Catmur, C., Cook, R., Widmann, H., Heyes, C., & Bird, G. (2012). fMRI evidence of ‘mirror’ responses to geometric shapes. *PLoS ONE*, 7, e51934. <http://dx.doi.org/10.1371/journal.pone.0051934>
- Press, C., Kok, P., & Yon, D. (2020a). Learning to perceive and perceiving to learn. *Trends in Cognitive Sciences*, 24, 260–261. <http://dx.doi.org/10.1016/j.tics.2020.01.002>
- Press, C., Kok, P., & Yon, D. (2020b). The perceptual prediction paradox. *Trends in Cognitive Sciences*, 24, 13–24. <http://dx.doi.org/10.1016/j.tics.2019.11.003>
- Press, C., & Yon, D. (2019). Perceptual prediction: Rapidly making sense of a noisy world. *Current Biology*, 29, R751–R753. <http://dx.doi.org/10.1016/j.cub.2019.06.054>

- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20, 260–281. <http://dx.doi.org/10.1016/j.tics.2016.01.007>
- Roussel, C., Hughes, G., & Waszak, F. (2013). A preactivation account of sensory attenuation. *Neuropsychologia*, 51, 922–929. <http://dx.doi.org/10.1016/j.neuropsychologia.2013.02.005>
- Roussel, C., Hughes, G., & Waszak, F. (2014). Action prediction modulates both neurophysiological and psychophysical indices of sensory attenuation. *Frontiers in Human Neuroscience*, 8, 115. <http://dx.doi.org/10.3389/fnhum.2014.00115>
- Sato, A. (2008). Action observation modulates auditory perception of the consequence of others' actions. *Consciousness and Cognition*, 17, 1219–1227. <http://dx.doi.org/10.1016/j.concog.2008.01.003>
- Schwarz, K. A., Pfister, R., Kluge, M., Weller, L., & Kunde, W. (2018). Do we see it or not? Sensory attenuation in the visual domain. *Journal of Experimental Psychology: General*, 147, 418–430. <http://dx.doi.org/10.1037/xge0000353>
- Seki, K., & Fetz, E. E. (2012). Gating of sensory input at spinal and cortical levels during preparation and execution of voluntary movement. *The Journal of Neuroscience*, 32, 890–902. <http://dx.doi.org/10.1523/JNEUROSCI.4958-11.2012>
- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, 117, 34–50. <http://dx.doi.org/10.1037/0096-3445.117.1.34>
- Stein, T., & Peelen, M. V. (2015). Content-specific expectations enhance stimulus detectability by increasing perceptual sensitivity. *Journal of Experimental Psychology: General*, 144, 1089–1104. <http://dx.doi.org/10.1037/xge0000109>
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, 15, 745–756. <http://dx.doi.org/10.1038/nrn3838>
- Summerfield, C., & Egner, T. (2016). Feature-based attention and feature-based expectation. *Trends in Cognitive Sciences*, 20, 401–404. <http://dx.doi.org/10.1016/j.tics.2016.03.008>
- Talluri, B. C., Urai, A. E., Tsetsos, K., Usher, M., & Donner, T. H. (2018). Confirmation bias through selective overweighting of choice-consistent evidence. *Current Biology*, 28, 3128–3135.e8. <http://dx.doi.org/10.1016/j.cub.2018.07.052>
- Thomas, E. R., Yon, D., de Lange, F. P., & Press, C. (2020). Action enhances predicted touch. *bioRxiv*. <http://dx.doi.org/10.1101/2020.03.26.007559>
- Van Doorn, G., Paton, B., Howell, J., & Hohwy, J. (2015). Attenuated self-tickle sensation even under trajectory perturbation. *Consciousness and Cognition*, 36, 147–153. <http://dx.doi.org/10.1016/j.concog.2015.06.016>
- Weiskrantz, L., Elliott, J., & Darlington, C. (1971). Preliminary observations on tickling oneself. *Nature*, 230, 598–599. <http://dx.doi.org/10.1038/230598a0>
- Weiss, C., Herwig, A., & Schütz-Bosbach, S. (2011). The self in social interactions: Sensory attenuation of auditory action effects is stronger in interactions with others. *PLoS ONE*, 6, e22723. <http://dx.doi.org/10.1371/journal.pone.0022723>
- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in python. *Frontiers in Neuroinformatics*, 7, 14. <http://dx.doi.org/10.3389/fninf.2013.00014>
- Wolpert, D. M., Doya, K., & Kawato, M. (2003). A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 358, 593–602. <http://dx.doi.org/10.1098/rstb.2002.1238>
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269, 1880–1882. <http://dx.doi.org/10.1126/science.7569931>
- Wyart, V., Nobre, A. C., & Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 3593–3598. <http://dx.doi.org/10.1073/pnas.1120118109>
- Yon, D., de Lange, F. P., & Press, C. (2019). The predictive brain as a stubborn scientist. *Trends in Cognitive Sciences*, 23, 6–8. <http://dx.doi.org/10.1016/j.tics.2018.10.003>
- Yon, D., Gilbert, S. J., de Lange, F. P., & Press, C. (2018). Action sharpens sensory representations of expected outcomes. *Nature Communications*, 9, 4288. <http://dx.doi.org/10.1038/s41467-018-06752-7>
- Yon, D., & Press, C. (2017). Predicted action consequences are perceptually facilitated before cancellation. *Journal of Experimental Psychology: Human Perception and Performance*, 43, 1073–1083. <http://dx.doi.org/10.1037/xhp0000385>
- Yon, D., & Press, C. (2018). Sensory predictions during action support perception of imitative reactions across suprasecond delays. *Cognition*, 173, 21–27. <http://dx.doi.org/10.1016/j.cognition.2017.12.008>

Received May 29, 2019

Revision received April 9, 2020

Accepted April 17, 2020 ■