Action enhances predicted touch

Emily R. Thomas[1*], Daniel Yon[1,2], Floris P. de Lange[3] & Clare Press[1]

1. Department of Psychological Sciences, Birkbeck, University of London, UK.

2. Department of Psychology, Goldsmiths, University of London, UK.

3. Donders Institute for Brain, Cognition and Behaviour, Radboud University, NL.

* Email: ethoma09@mail.bbk.ac.uk

Author Contributions

All authors contributed to developing the study concept and design. E.R. Thomas performed data collection and analysis under the supervision of C. Press. All authors approved the final version of the manuscript for submission.

Acknowledgments

Leverhulme Trust (RPG-2016-105) and Wellcome Trust (204770/Z/16/Z) grants awarded to C. Press supported this work. F. de Lange was supported by a Vidi Grant (Nederlandse Organisatie voor Wetenschappelijk Onderzoek, 452-13-016) and ERC Starting Grant (Horizon 2020 Framework Programme, 678286).

Abstract

It is widely believed that predicted tactile action outcomes are perceptually attenuated. The present experiments determined whether predictive mechanisms necessarily generate attenuation, or instead can enhance perception – as typically observed in sensory cognition domains outside of action. We manipulated probabilistic expectations in a paradigm often used to demonstrate tactile attenuation. Participants produced actions and subsequently rated the intensity of forces on a static finger. Experiment 1 confirmed previous findings that action outcomes are perceived less intensely than passive stimulation but demonstrated more intense perception when active finger stimulation was removed. Experiments 2 and 3 manipulated prediction explicitly and found that expected touch during action is perceived *more* intensely than unexpected touch. Computational modelling suggested that expectations increase the gain afforded to expected tactile signals. These findings challenge a central tenet of prominent motor control theories and demonstrate that sensorimotor predictions do not exhibit a qualitatively distinct influence on tactile perception.

Statement of Relevance

Perception of expected action outcomes is thought to be attenuated. Such a mechanism may be adaptive because surprising inputs are more useful – e.g., signalling the need to take new courses of action – and is thought to explain why we cannot tickle ourselves and unusual aspects of action and awareness in clinical populations. However, theories outside of action purport that predicted events are perceptually enhanced, allowing us to generate largely accurate representations of our noisy sensory world. We do not know whether action predictions really alter perception differently from other predictions because different manipulations have been performed. Here we perform similar manipulations and demonstrate that action predictions can enhance, rather than attenuate, touch. We thereby demonstrate that action predictions may not have a qualitatively distinct influence on perception, meaning we must re-examine theories of perceptual prediction across domains and clinical theories based upon their assumptions.

Keywords: Prediction, Perception, Attenuation, Motor Processes

When we produce actions we predict their sensory consequences. Prominent motor theories (Blakemore et al., 1998; Dogge et al., 2019; Kilteni & Ehrsson, 2017) propose that we attenuate – or downweight – perception of expected action outcomes. Such downweighting mechanisms are thought to finesse the limited capacity of our sensory systems, prioritising perception of more informative unexpected events that signal the need to perform new actions or update our models of the world (Press et al., 2020b; Wolpert & Flanagan, 2001). For example, if we lift a cup of coffee that is lighter than expected, attenuated processing of expected signals (e.g., touch on our fingertips) will prioritise perception of unexpected events (e.g., accelerating motion of the cup) allowing swift updating of our beliefs about the environment (e.g., the weight of the cup) and support corrective action to avoid spillage. These downweighting mechanisms are invoked to explain why self-produced tactile sensations generate lower secondary somatosensory cortex activity (Blakemore et al., 1998; Kilteni & Ehrsson, 2020), and are perceived less intensely (Bays et al., 2005, 2006; Kilteni et al., 2019), than externally-produced forces. This theory also provides an explanation for why it is difficult to tickle oneself (Blakemore et al., 1998).

However, outside of action it is thought that prediction mechanisms generate a qualitatively opposite influence on perception. In these theories – typically couched in Bayesian frameworks – it is proposed that we combine our expectations (prior) with the input (likelihood) to determine what we perceive (posterior; Kersten et al., 2004). Such a process would *upweight*, rather than downweight, perception of expected events, enhancing the detectability and apparent intensity of events (Brown et al., 2013) – and thereby enabling rapid generation of largely veridical experiences in the face of sensory noise (de Lange et al., 2018). For example, some theories propose that we use predictions to increase the gain on sensory units tuned to expected events and, via competitive local interactions, inhibit sensory populations tuned to unpredicted events (de Lange et al., 2018; Kok et al., 2012; Press & Yon, 2019; Yon et al., 2018). However, it is perhaps unclear why the adaptive arguments presented for downweighting (informativeness) and upweighting (veridicality) predicted perceptual experiences should apply differentially in the domain of action (note that attenuation and downweighting, as well as enhancement and upweighting, will be used interchangeably). Specifically, it appears just as crucial to optimise informativeness and veridicality regardless of how predictions are formed (Press et al., 2020b), and some evidence from the visual domain suggests that predictive influences on perception do not exhibit the qualitative differences assumed in the literature (Yon et al., 2018, 2020; Yon & Press, 2017).

Given the comparability of this recent visual evidence, a notable stark difference between studies purporting to demonstrate upweighting and downweighting is that the former study visual perception whereas the latter study tactile perception. It is therefore widely believed that action predictions shape tactile perception in a qualitatively distinct way – including proposals that differences relate to tactile events being body-related (Dogge et al., 2019) and tightly coupled with the motor system (Kusnir et al., 2019), in a way that many predicted visual or auditory events are not. Similarly, differences may also relate to assumptions that tactile attenuation during action is dependent upon somatosensory-cerebellar connectivity (Blakemore et al., 1998; Kilteni & Ehrsson, 2020), in contrast with hippocampal mediation of prediction in visual processing (Kok & Turk-Browne, 2018).

However, studies examining touch perception during action have not manipulated predictability like those in wider sensory cognition and therefore it would be premature to assume that qualitatively distinct mechanisms influence touch. The defining feature of prediction mechanisms is that they operate according to stimulus probabilities (de Lange et al., 2018). As such, prediction mechanisms outside of action contexts are typically measured by presenting events with high and low conditional probabilities, allowing comparison of perception of ‘expected’ (e.g., 80% likely based upon a preceding cue) and ‘unexpected’ events (20% likely; Cheadle et al., 2015; Kok et al., 2012; Richter & de Lange, 2019). In contrast, studies demonstrating tactile attenuation during action compare the perception of events in the presence or absence of action, or when events are coincident versus delayed with respect to action (Bays et al., 2005; Blakemore et al., 1998; Kilteni et al., 2019; Wolpe et al., 2018). In these experiments it is assumed that the sensory events which coincide with action are more predicted, explaining why perception of them is attenuated. However, it is unclear whether these effects indeed reflect the operation of predictive mechanisms when stimulus probabilities have not been manipulated and various non-predictive mechanisms influence perception during action (Press et al., 2020a; Press & Cook, 2015). For example, when we move we suppress *all* tactile input to a moving effector (Williams & Chapman, 2000).

Relatedly, ‘active inference’ predictive processing accounts and even classic working memory models would expect reduced perception of all sensory events in the presence of action (Press et al., 2020a), regardless of the extent to which they were predicted on its basis. Thus, to test whether action predictions really influence touch perception via qualitatively distinct mechanisms, the present studies adopted a force judgement paradigm used widely in action domains to examine tactile attenuation in combination with a probabilistic predictive manipulation typical in broader sensory cognition domains.

General Method

Participants

Thirty distinct participants were tested in Experiment 1 (16 female, mean age = 25.53 years [SD = 5.25]), Experiment 2 (20 female, mean age = 22.80 years [SD = 3.18]) and Experiment 3 (22 female, mean age = 24.3 years [SD = 4.34]). Eight participants in Experiment 1, six participants in Experiment 2, and nine participants in Experiment 3, were replacements for those where acceptable psychometric functions could not be modelled to their responses (flat functions), where they were unable to follow instructions concerning movement performance (>20% recorded movement errors), or where there was technical malfunction. These criteria were established a priori to participant testing and replacements resulted in a total sample of 30 participants in each experiment. One participant’s PSE score from Experiment 2 was winsorized to meet the normality assumptions of parametric tests (from $z = 3.34$ to $z = 3$, Tukey, 1962). Participants were recruited from Birkbeck, University of London, and paid a small honorarium for their participation. All participants reported no current neurological or psychiatric illness and provided written informed consent prior to participation. The experiments were performed with local ethical committee approval (Birkbeck, University of London) and in accordance with the ethical standards laid down in the 1964 Declaration of Helsinki. The sample size was determined a priori on the basis of pilot testing to estimate effect size – to have at least 80% power of detecting medium effect sizes ($d = 0.5$) – and parametric assumptions were met. Data collection stopped after attaining the pre-determined sample size of 30 participants adhering to the above criteria in each experiment.

Experiment 1

A ‘Contact’ condition in Experiment 1 aimed to determine whether we could replicate typical action attenuation effects within our set-up. Participants therefore moved an active right index finger to make contact with a button, generating a simultaneous mechanical force to their left index finger below (see Procedure and Fig. 1A). We thereby examined whether left-hand stimulation is reported as less forceful during right-hand action than when the right-hand remains still (‘passive’). We were also interested in the nature of effects in a ‘No Contact’ condition, where a similar right-finger downward motion triggered the same left-hand stimulation but did not make contact with a button. Stimulation to the left hand was instead triggered via motion-tracker detection of right-hand intransitive (i.e., not directed towards an object) motion. Given that intransitive actions frequently produce sensory effects, active-finger contact should not be required to form sensorimotor predictions *per se,* and this generation of an active-finger sensory event simultaneous with target stimulation especially complicates interpretation. Specifically, studies examining predictive attenuation have examined perception of events on passive effectors due to potential confounds of ‘generalised gating’ when measuring perception on active effectors. Generalised gating is thought to reduce perception of *any* events delivered to moving effectors (Williams & Chapman, 2000), without exhibiting specificity to predicted consequences because it is hypothesised to occur at the earliest relay in the spinal cord (Seki & Fetz, 2012). One possible reason that active-finger contact complicates interpretation of attenuation in ‘Contact’ set-ups, is that concurrent gated sensory events on active effectors could bias responses about stimulation on passive effectors, for instance, due to response biases (Firestone & Scholl, 2016). If this is the case, we may observe different effects in Contact and No Contact conditions when this non-predictive influence is removed in the latter.

Procedure

The experiment was conducted in MATLAB using the Cogent toolbox. Participants held their left-hand palm upwards (Fig. 1A) with their index finger positioned against a solenoid (diameter of metal rod = 4 mm; diameter of solenoid = 15 mm; TACT-CONTR-TR2, Heijo Research Electronics) sitting on the apex of the fingertip. Their right-hand rested on a right armrest, positioned such that the index finger distal phalange was directly above the left-hand distal phalange, but rotated 90 degrees anticlockwise relative to their left hand (Fig 1A). A small button box or infrared motion tracker (Leap Motion Controller using the Matleap MATLAB interface) was placed on a shelf supporting the solenoid (depending on block type, Contact or No Contact respectively), positioned directly above it.

At the start of each trial, participants were cued onscreen to move their right index finger ('move'; Active trials – 50%) or remain stationary ('do not move'; Passive trials – 50%). In both Contact and No Contact blocks participants were required to hold their right-hand palm facedown, with their arm parallel with the computer monitor and the coronal body midline. In all conditions, participants' left-hand was positioned laterally from the body midline and in line with the shoulder. The experimenter provided demonstration of appropriate action execution for both conditions before the onset of the experiment. These steps ensured that participants' starting position remained the same regardless of trial type (Active vs Passive), and that actions were executed in approximately the same way across Contact and No Contact blocks (bar the contact with the button; see Supplementary Materials). Participants' hands were visually occluded during the experiment and white noise was played through headphones (53 Db; piloting confirmed that this level resulted in inaudible solenoid movement) throughout testing.

On Active trials, they rotated their index finger downwards at the metacarpophalangeal joint. The target stimulus was delivered to the left index finger for 30 ms immediately after motion was detected by a button press in Contact blocks, or as soon as the active finger movement achieved approximately equivalent movement distance in No Contact blocks (distance of at least 20 mm from the stationary finger starting position at trial onset, see Supplementary Materials). Piloting confirmed that stimulation was in apparent synchrony with movement termination in both Contact and No Contact blocks. After 1000 ms, a reference stimulus was presented for 30 ms. The target stimulus presented one of seven logarithmically-spaced suprathreshold forces, and the reference stimulus always presented the fourth (middle) force (see Supplementary Materials for more details about the generation of tactile events). After a 300 – 500 ms delay, participants were asked which tap was more forceful, responding with a left foot pedal for the first stimulus and a right foot pedal for the second stimulus. The next trial started after 1000 ms. In Passive trials, the target stimulus was delivered 500 ms after the cue to remain still. Due to the fact that our

Contact vs No Contact manipulation only affected active trials, passive trials were identical in both Contact and No Contact blocks and served as comparisons within blocks.

There were 560 trials in total; 140 for each of the Active and Passive conditions, in both the Contact and No Contact blocks. The order of blocks was counterbalanced across participants and trial type (Active vs Passive) order was randomized across blocks. Participants completed eight practice trials before the main test blocks.

Modelling Psychometric Functions

Participant responses were modelled by cumulative Gaussians to estimate psychometric functions, using the Palamedes Toolbox (Prins & Kingdom, 2018) in MATLAB. This procedure was performed separately for active and passive trials during the test phase in Experiment 1. The mean of the modelled Gaussian was taken as the Point of Subjective Equivalence (PSE), describing the point at which participants judge the target and reference events to have equal force. Lower values are indicative of more intense target percepts.

Results

PSE values were analysed in a 2x2 within-participants ANOVA, revealing no main effect of Contact ($F(1, 29) = 3.11$, $p = .089$ $\eta p^2 = .10$) or Movement ($F(1, 29) = 1.24$, $p = .274$, $\eta p^2 = .04$). However, there was a significant interaction between Contact and Movement ($F(1, 29) = 15.39$, $p < .001$, $\eta p^2 = .35$), driven by lower force judgements (higher PSEs) in Active ($M = 4.73$, $SD = 1.22$) compared to Passive trials ($M = 4.35$, $SD = .80$) in the Contact condition ($t(29) = 2.07$, $p = .047$, $d = .38$), but higher force judgements (lower PSEs) in Active ($M = 3.82$, $SD = 1.11$) compared to Passive trials ($M = 4.45$, $SD = .97$) in the No Contact condition ($t(29) = -3.80$, $p = .001$, $d = .69$, see Fig. 1C).

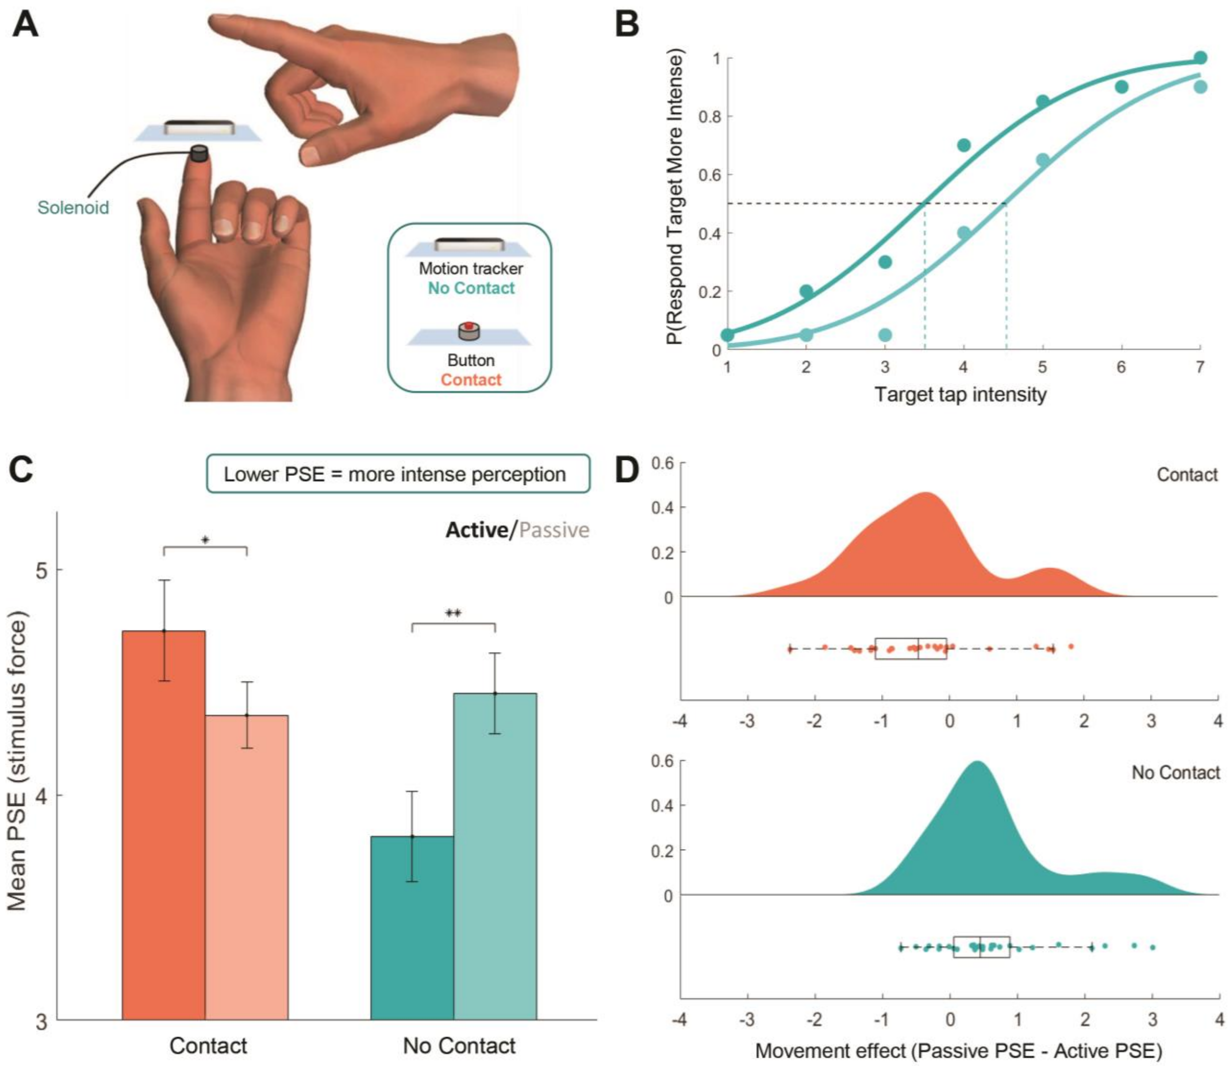

Figure 1. Experiment 1. (A) On each trial, participants made downward movements with their right index finger either over a motion tracker (No Contact condition), or towards a button with which they made contact (Contact condition). Each movement elicited a tactile punctate event to the left index finger positioned directly below. (B) PSEs were calculated for each participant (data represents an example participant in the No Contact condition for Active [dark blue] and Passive [light blue] trials). (C) Mean PSEs (± SEM) were higher in Active than Passive trials in the Contact conditions, but lower in Active than Passive trials in the No Contact condition. Lower PSEs indicate more intense target percepts (* $p < .05$, ** $p = .001$). (D) PSE effect of movement (Passive – Active) for the Contact (top) and No Contact (bottom) condition, plotted with raincloud plots displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive effects of movement indicate more intensely perceived active target events relative to passive events, but negative values indicate the reverse – less intensely perceived active events.

Experiment 2

Experiment 1 replicated previous findings (Bays et al., 2005; Kilteni et al., 2019) that tactile events on a stationary left finger are perceived less intensely during active right-hand movement, but only when the specifics of the paradigm were replicated such that the active finger makes contact with a button. When there is no contact, events are perceived *more* intensely during movement. One possible explanation of the difference between conditions is that perception of gated stimulation on the active effector contributes to the Active-Passive difference in the Contact condition.

Having established that we can observe typical attenuation during action – but that attenuation becomes enhancement in a No Contact condition – Experiment 2 isolated the particular functional influence of prediction mechanisms by manipulating conditional probabilities between actions and outcomes. This is particularly important for establishing the role of action predictions in determining perception, because – as outlined in the introduction – an Active-Passive comparison does not isolate predictive influences of action even in a No Contact condition, e.g., because it confounds the number of tasks. Therefore, we compared perception of tactile events when they were expected or unexpected based on learned action-outcome probabilities established in a preceding training session, but always in the presence of action. While downweighting accounts predict that expected tactile events will be rated less intensely (forceful) than unexpected events, upweighting theories predict that expected events will be rated more intensely.

Procedure changes relative to Experiment 1

Participants now performed one of two movements that predicted one of two tactile effects (Fig. 2A). Participants were positioned with their left index and middle finger making contact with independent solenoids (Fig. 2A). At the beginning of each trial an arbitrary cue (either a square or circle) instructed participants to move their right index finger either upwards or downwards from the metacarpophalangeal joint, tracked by an infrared motion sensor. This action triggered delivery of the target stimulus to either solenoid in the same way as in the No Contact condition of Experiment 1 (see Supplementary Materials for more details). During training, participants' right-hand index-finger action (e.g., downwards movement) was 100% predictive of the location of left-hand tactile events (e.g., index finger). In a test session 24 hours later, the action-outcome relationship was degraded to measure perception of expected and unexpected events – the expected finger was stimulated on 66.7% of trials, and the unexpected was stimulated on the remaining 33.3% of trials.

Presenting two action types and two stimulation types also allowed us to compare perception of expected and unexpected events, while controlling for repetition effects. It should be noted in explicating the logic of the procedure here that any action predictions should determine *where* stimulation will be received, rather than its intensity. However, in Bayesian models it is assumed that enhanced detection and intensity of expected events relates to the precision of the estimate (Brown et al., 2013, e.g., a force precisely estimated to have occurred on a certain region of tactile space should be feel more intense because of the precise estimate of spatial information, rather than an estimate of the force *per se*). Since these models assume that predictions enhance the precision of resultant estimates, they would also predict enhancements in perceived force (and indeed other sensory attributes, like brightness or loudness).

There were 420 trials in each session. Trial order was randomised and the action-stimulus mapping was counterbalanced across participants. The cue-action mapping was reversed halfway through each session to account for effects resultant from possible learning of cue-outcome associations instead of action-outcome associations. The training and test sessions were carried out at the same time on consecutive days. Participants completed 12 practice trials before the main session trials.

Results

Psychometric functions were modelled to participants' responses similarly to in Experiment 1, but now separately for expected and unexpected events. PSE values were lower on expected trials (M = 3.72, SD = .96) than unexpected trials (M = 3.93, SD = .80; $t(29) = -2.13$, p = .041, d = .39, see Fig. 2D), demonstrating more forceful perception of expected than unexpected action outcomes.

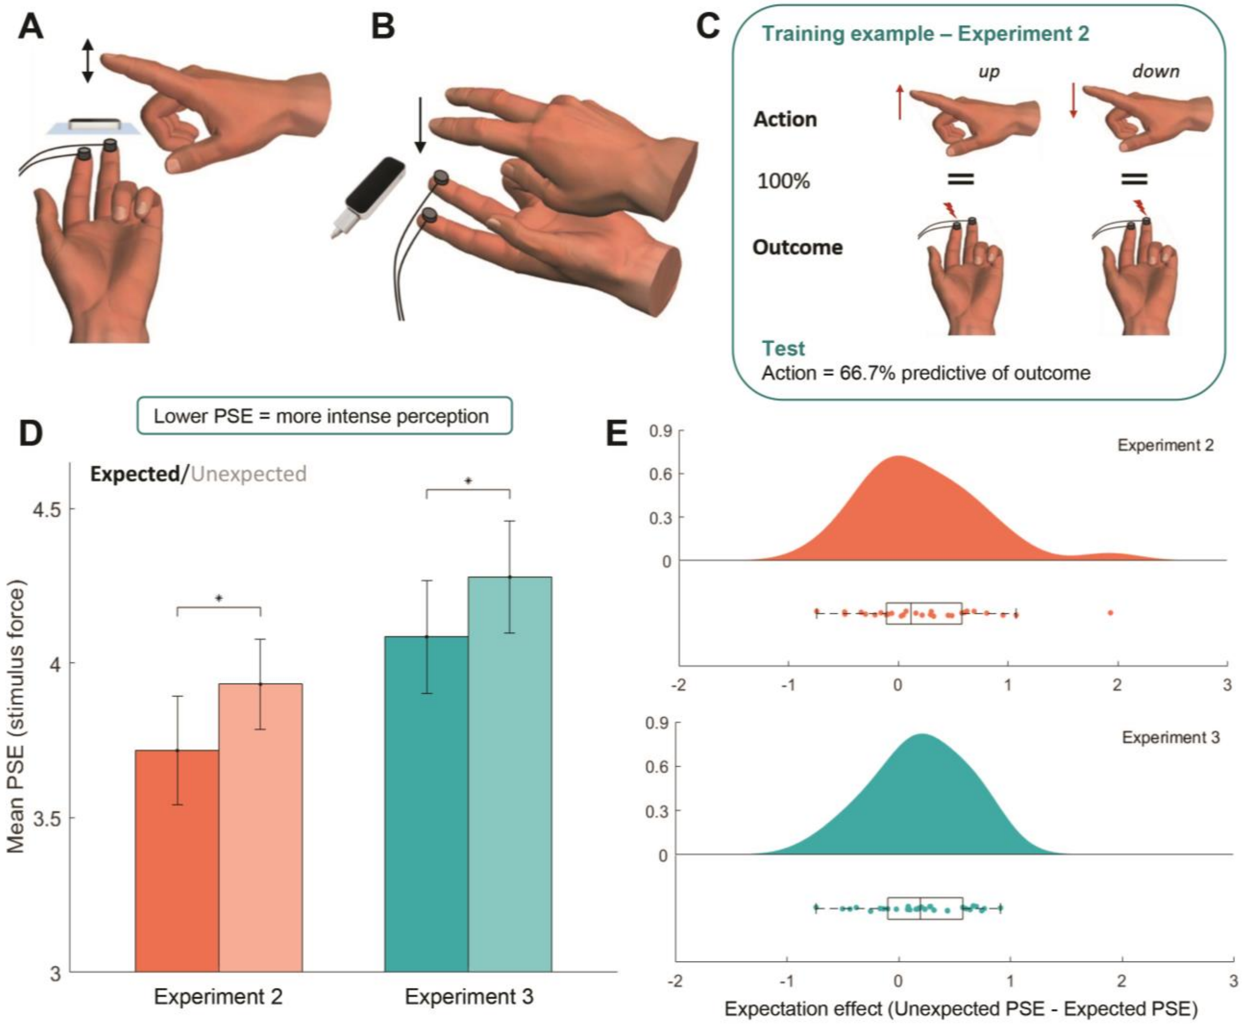

Figure 2. (A) Experiment 2 set-up. On each trial, participants made a downwards or upwards movement with their right index finger over a motion tracker, which elicited tactile punctate events to the left index or middle finger. (B) Experiment 3 set-up. Participants made only downwards movements, now with either their right index or middle finger. (C) Example training procedure for Experiment 2. Movements were perfectly predictive of tactile events during the training session, and 66.7% predictive in the test session. This procedure was also adopted in Experiment 3 but with different action types. Flash symbols illustrate location of stimulation. (D) Mean PSEs (± SEM) were lower for expected than unexpected trials in both Experiment 2 and Experiment 3. Lower PSEs indicate a more intensely perceived target stimulus (* $p < .05$). (E) PSE expectation effect (Unexpected – Expected) plotted with raincloud plots displaying probability density estimates (upper) and box and scatter plots (lower), for Experiment 2 (top) and Experiment 3 (bottom). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive expectation effect values indicate more intensely perceived expected events relative to unexpected events.

Experiment 3

Experiment 3 was designed to provide a conceptual replication of Experiment 2, using a set-up more closely aligned with typical action paradigms (e.g., Experiment 1) – whereby one always makes a movement towards another effector. We additionally controlled for possibilities that the expectation effect in Experiment 2 resulted from cue-outcome learning by removing the cue stimulus and requiring free selection of action. The explicit reference stimulus was also removed and comparisons were made against an implicit reference, eliminating the possibility that effects are determined by forming predictions about the reference stimulus.

Procedure changes relative to Experiment 2

Independent solenoids were now attached to the left index and middle fingers via adhesive tape (diameter of metal rod = 4 mm; diameter of solenoid = 18 mm; TactAmp 4.2 Dancer Design). The foot pedals were positioned at either 45 (for stronger) or 90 (for weaker) degree angles relative to their right foot to record responses, to account for any spatial biases resulting from positioning foot pedals as ‘left’ and ‘right’. At the start of each trial, participants selected to make a downwards movement with their right index or middle finger. These action types were selected to ensure that effects in Experiment 2 were not specific to those action types and to determine whether similar effects could be observed with actions that are always made towards another effector. Participants’ hands, and therefore index and middle fingers, were spatially aligned with each other (Fig. 2B). Actions were freely selected and the frequency of index and middle finger movements was monitored to ensure approximately equal numbers of both action types. Participants’ actions (e.g., right index downwards movement) were still perfectly predictive of the location of tactile events (e.g., left index finger) during training, and this contingency was again degraded to 66.7% in the following test session. Participants were asked whether they perceived the test force to be more or less forceful than the average force intensity. An example of the average force was presented to each finger once at the end of short breaks every 21 trials (NB: the average force was identical to the reference force intensity).

The experiment consisted of two training blocks followed by a test block, all occurring in the same session of testing. In the first training block participants responded yes/no to the question ‘Tap on index or middle finger?’, and in the second training block they were asked about the force, similarly to in Experiment 2 and in subsequent test blocks. Half of the participants experienced a mapping whereby moving the right-hand index finger resulted in left-hand index stimulation and middle finger movement resulted in middle stimulation. The other half experienced a mapping whereby index finger movement resulted in middle stimulation, and middle finger movement in index stimulation. There were 210 trials in each session.

Results

Like in Experiment 2, PSE values were lower in expected (M = 4.08, SD = 1.00) than unexpected (M = 4.28, SD = .99) trials ($t(29)$ = -2.56, p = .016, d = .47, see Fig. 2C), again demonstrating more forceful perception of expected than unexpected action outcomes. Additional post-hoc analyses revealed that the specific kinematics of action were similar in expected and unexpected trials, and that the PSE expectation effect was comparable at the start and end of 21-trial mini-blocks (see Supplementary Materials).

Computational modelling

The present findings are consistent with predictive upweighting theories of perception, which propose that observers combine sensory evidence with prior knowledge – biasing perception towards what we expect (Kersten et al., 2004). This may be achieved mechanistically by altering the weights on sensory channels, increasing the gain of expected relative to unexpected signals (de Lange et al., 2018). However, expectation effects may instead reflect biasing in response-generation circuits – such that action biases people to *respond* that expected events are more intense, rather than altering perception itself (Firestone & Scholl, 2016).

Drift Diffusion Modelling procedure

Perceptual and response biases can be dissociated in computational models that conceptualise perceptual decisions as a process of evidence accumulation. Perceptual biases are modelled as growing across time – every time response units sample from perceptual units they will be sampling from a biased representation, therefore increasing the magnitude of biasing effects across a larger number of samples

(Yon et al., 2020). In contrast, response biases are modelled as operating regardless of current incoming evidence and to be present from the outset of a trial. According to this logic, we can model the decision process with drift diffusion modelling (DDM; Ratcliff & McKoon, 2008) to identify the nature of the biasing process. We can thus establish whether action expectations shift the starting point of evidence accumulation towards a response boundary (‘start biasing’, *z* parameter; Fig. 3A), or instead bias the rate of evidence accumulation (*db* parameter, ‘drift biasing’, Fig. 3B).

We fit drift diffusion models to participant choice and reaction time data from Experiment 3 using the hDDM package implemented in Python (Wiecki et al., 2013; NB: reaction times were not collected in Experiments 1 and 2). In the hDDM, 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: 1) a *null* model where no parameters were permitted to vary between expected and unexpected trials; 2) a *start bias* model where the start point of evidence accumulation (*z*) could vary between expectation conditions; 3) a *drift bias* model where a constant added to evidence accumulation (*db*) could vary according to expectation; 4) a *start + drift bias* model where both parameters could vary according to expectation.

All models were estimated with MCMC sampling, and parameters were estimated with 30,000 samples (‘burn-in’=7,500). Model convergence was assessed by inspecting chain posteriors and simulating reaction time distributions for each participant. Models were compared using deviance information criteria (DIC) as an approximation of Bayesian model evidence, a common method used to determine model fit. Lower DIC values relative to a baseline, or null, model are indicative of a better model fit.

A posterior predictive check was conducted using the hDDM package to establish how well each model was able to reproduce the patterns in our data. The posterior model parameters for the *start bias*, *drift bias*, and *start + drift bias* models were used to simulate a distribution of 500 reaction times and choices for each trial for each participant. From this simulated data we calculated the probability that a ‘stronger than average’ response was given at each intensity level, separately for expected and unexpected trials.

This allowed us to model simulated psychometric functions for expected and unexpected trials, exactly as we had done for empirical decisions. Performing this procedure for each model yielded separate simulated expectation effects (Unexpected PSE – Expected PSE) for each participant under the *start bias*, *drift bias,* and *start* + *drift bias* models.

Results

Fitting the DDM to the behavioural data found that the model allowing both start and drift biases to vary according to expectation provided the best fit (DIC relative to null = -234.8) relative to both the start bias (DIC relative to null = -191.06) and drift bias (DIC relative to null = -8.62) models. This finding may suggest that observed biases are a product of both start and drift rate biasing. However, although the DIC measure does include a penalty for model complexity, it is thought to be biased towards models with higher complexity (Wiecki et al., 2013) and it indeed favoured the most complex model here.

We conducted a posterior predictive check to evaluate how well simulated data from each of the models could reproduce key patterns in our data. Correlations were calculated to quantify how well simulated expectation effects reproduced empirical expectation effects, which revealed significant relationships for all three models (Start bias model: r_{30}= .39, p = .034; Drift bias model: r_{30}= .43, p = .017; Start + Drift bias model: r_{30}= .53, p = .003, see Fig. 3C).

Given that we were interested in whether any of the PSE expectation effect is generated by sensory biasing – rather than possible additional contributions of response biasing – we examined whether drift biasing accounted for any further variance in expectation effects than start biasing alone, by conducting a stepwise linear regression to predict the empirical expectation effect (Unexpected PSE – Expected PSE). In the first step, we included the simulated expectation effect from the start bias model to predict the empirical expectation effect. The simulated start bias data was able to predict the empirical expectation effect (R^2 = .15, F(1,28) = 4.96, p = .034). In the second step, we included the simulated expectation effect from the drift bias model as an additional predictor of the empirical expectation effect, importantly providing a significant improvement to the model fit ($Fchange$(1,27) = 6.72, p = .015; R^2 = .32, F(2,27) = 6.34, p = .006). This analysis reveals that a model implementing a drift biasing mechanism better predicts empirical effects of expectation on perceptual decisions, by explaining unique variance in participant decisions that cannot be explained by start biasing.

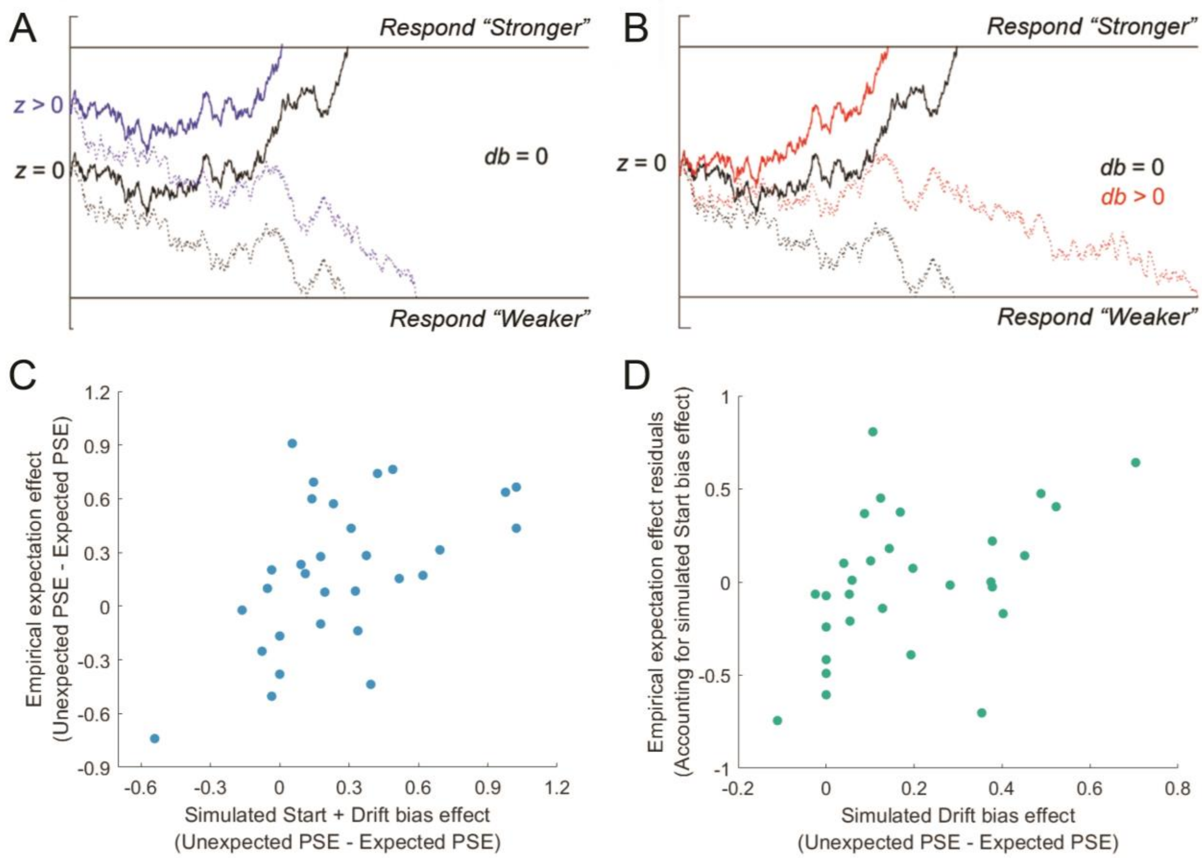

Figure 3. Illustration of how the DDM could explain expectation biases and results of computational modelling. (A) For an unbiased decision process (black lines), sensory evidence integrates towards the upper response boundary when stimuli are stronger than average (solid lines) and towards the lower response boundary when weaker than average (dotted lines). Baseline shifts in decision circuits could shift the start point of the accumulation process nearer to the upper boundary for expected 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 (influencing parameter *db*; red lines – Drift bias model). (C) Simulated Start + Drift bias (winning DIC model) expectation effect plotted against the empirical expectation effect, showing a significant positive correlation. (D) Simulated Drift bias expectation effect plotted against the empirical expectation effect accounting for simulated Start bias effects (plotted as the residuals from a model where the simulated Start bias effect predicts the empirical effect), again showing a significant positive correlation. Importantly, our regression analysis revealed that drift biases accounted for significant additional variance once accounting for start biases. All expectation effects were calculated by subtracting Expected PSEs from the Unexpected PSEs.

Discussion

Extant models disagree about how predictions should shape perception of action outcomes. We examined whether sensorimotor prediction influences touch perception via qualitatively distinct mechanisms from other types of prediction by adapting the force judgement paradigm typically used in the action literature (Bays et al., 2005, 2006; Kilteni et al., 2019), and applying predictive manipulations from broader sensory cognition (de Lange et al., 2018; Richter & de Lange, 2019). Experiment 1 replicated typical findings that self-produced forces are rated as less intense than externally-generated ones, but this effect reversed when there was no active-finger contact with a button. Experiments 2 and 3 manipulated the predictability of tactile action-outcomes and found that expected events were perceived *more*, not less, intensely than unexpected events. Computational modelling suggested that expectations alter the way sensory evidence is integrated – increasing the gain afforded to expected tactile signals.

These findings are consistent with predictive upweighting accounts from outside of action domains, which propose that prior expectations are combined with sensory evidence to generate veridical perceptual interpretations of our noisy environment – thereby rendering expected events more intense. The present findings indicate that these upweighting mechanisms operate similarly in touch. It is therefore essential to consider how the present findings can be resolved with data cited in support of downweighting theories. As well as the data already outlined in humans, there are a range of related findings in other species – e.g., attenuating internally-generated electric fields in Mormyrid fish is thought to improve detection of prey-like stimuli (Enikolopov et al., 2018) and virtual reality trained mice show suppressed auditory responses to self-produced tones generated by treadmill running (Schneider et al., 2018). However, studies have not demonstrated whether underlying mechanisms operate according to stimulus probabilities. There are a number of non-predictive mechanisms which could explain attenuation, and based upon the current findings we propose that many effects are instead generated by identity-general gating mechanisms, and others possibly by mechanisms shaping perception according to event repetition – given that repetition is frequently confounded with expectation (e.g., Kilteni et al., 2019, see Feuerriegel et al., 2020). The assumption that action predictions specifically attenuate perception is central to a number of clinical models, including accounts of sensory differences in healthy ageing (Wolpe et al., 2016), motor severity in Parkinson's disease (Wolpe et al., 2018), and hallucinations in schizophrenia (Corlett et al., 2019). However, if predictions shape perception similarly regardless of domain then these theories may need revisiting.

Importantly however, the present data should not be taken to reflect that predictive attenuation cannot occur, especially given the importance of generating perceptual experiences that are informative across domains (Press et al., 2020b). Nevertheless, they suggest that predictive mechanisms during action operate differently from current assumptions. It has been widely claimed that attenuation likely results from subtracting the prediction from the input (Wolpert & Flanagan, 2001), and such a mechanism would be hard to reconcile with the present findings demonstrating upweighted sensory gain of predicted action outcomes. These data more likely suggest that purported predictive mechanisms must have the capability of generating both up- and downweighting, but under different circumstances. One possible resolution to the current debate has been recently outlined by some of us assuming opposing processes with differing roles (Press et al., 2020b), and resolving it must prove a focus of future work.

Data Availability

All data and documentation will be deposited in the Birkbeck Research Data Depository (BiRD - https://researchdata.bbk.ac.uk/) and be openly accessible by an associated DOI.

References

Bays, P. M., Flanagan, J. R., & Wolpert, D. M. (2006). Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biology*, *4*(2), 281–284. https://doi.org/10.1371/journal.pbio.0040028

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*(12), 1125–1128. https://doi.org/10.1016/j.cub.2005.05.023

Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, *1*(7), 635–640. https://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*(4), 411–427. https://doi.org/10.1007/s10339-013-0571-3

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*(3), 308–313. https://doi.org/10.1038/nn1194

Cheadle, S., Egner, T., Wyart, V., Wu, C., & Summerfield, C. (2015). Feature expectation heightens visual sensitivity during fine orientation discrimination. *Journal of Vision*, *15*(14), 14. https://doi.org/10.1167/15.14.14

Corlett, P. R., Horga, G., Fletcher, P. C., Alderson-Day, B., Schmack, K., & Powers, A. R. (2019). Hallucinations and Strong Priors. *Trends in Cognitive Sciences*, *23*(2), 114–127. https://doi.org/10.1016/j.tics.2018.12.001

de Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception? *Trends in Cognitive Sciences*, *22*(9), 764–779. https://doi.org/10.1016/j.tics.2018.06.002

Dogge, M., Custers, R., & Aarts, H. (2019). Moving Forward: On the Limits of Motor-Based Forward

Models. *Trends in Cognitive Sciences*, *23*(9), 743–753. https://doi.org/10.1016/J.TICS.2019.06.008

Enikolopov, A. G., Abbott, L. F., & Sawtell, N. B. (2018). Internally Generated Predictions Enhance Neural and Behavioral Detection of Sensory Stimuli in an Electric Fish. *Neuron*, *99*(1), 135–146. https://doi.org/10.1016/j.neuron.2018.06.006

Feuerriegel, D., Vogels, R., & Kovács, G. (2020). Evaluating the Evidence for Expectation Suppression in the Visual System. *PsyArXiv*. https://doi.org/10.31234/OSF.IO/Y7639

Firestone, C., & Scholl, B. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioural and Brain Sciences*, *39*, e229. https://doi.org/https://doi.org/10.1017/S0140525X15000965

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference. *Annu. Rev. Psychol.*, *55*, 271–304. https://doi.org/https://doi.org/10.1146/annurev.psych.55.090902.142005

Kilteni, K., & Ehrsson, H. (2020). Functional connectivity between the cerebellum and somatosensory areas implements the attenuation of self-generated touch. *Journal of Neuroscience*, *40*(4), 894–906. https://doi.org/10.1523/JNEUROSCI.1732-19.2019

Kilteni, K., & Ehrsson, H. H. (2017). Body ownership determines the attenuation of self-generated tactile sensations. *Proceedings of the National Academy of Sciences*, *114*(31), 8426–8431. https://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). https://doi.org/10.7554/eLife.42888

Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*(2), 265–270. https://doi.org/10.1016/j.neuron.2012.04.034

Kok, P., & Turk-Browne, N. B. (2018). Associative prediction of visual shape in the hippocampus. *Journal of Neuroscience*, *38*(31), 6888–6899. https://doi.org/10.1523/JNEUROSCI.0163-18.2018

Kusnir, F., Pesin, S., Moscona, G., & Landau, A. N. (2019). When temporal certainty doesn't help. *Journal of Cognitive Neuroscience*, *32*(2), 315–325. https://doi.org/10.1162/jocn_a_01482

Press, C., & Cook, R. (2015). Beyond action-specific simulation: domain-general motor contributions to perception. *Trends in Cognitive Sciences*, *19*(4), 176–178. https://doi.org/10.1016/j.tics.2015.01.006

Press, C., Kok, P., & Yon, D. (2020a). Learning to Perceive and Perceiving to Learn. *Trends in Cognitive Sciences*, *24*(4), 260–261. https://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*(1), 13–24. https://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*(15), R751–R753. https://doi.org/10.1016/j.cub.2019.06.054

Prins, N., & Kingdom, F. A. A. (2018). Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes toolbox. *Frontiers in Psychology*, *9*(1250). https://doi.org/10.3389/fpsyg.2018.01250

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Richter, D., & de Lange, F. P. (2019). Statistical learning attenuates visual activity only for attended stimuli. *ELife*, *8*, e47869. https://doi.org/10.1101/653782

Schneider, D. M., Sundararajan, J., & Mooney, R. (2018). A cortical filter that learns to suppress the acoustic consequences of movement. *Nature*, *561*, 391–395. https://doi.org/10.1038/s41586-018-0520-5

Seki, K., & Fetz, E. E. (2012). Gating of Sensory Input at Spinal and Cortical Levels during Preparation and Execution of Voluntary Movement. *Journal of Neuroscience*, *32*(3), 890–902. https://doi.org/10.1523/JNEUROSCI.4958-11.2012

Shergill, S. S., White, T. P., Joyce, D. W., Bays, P. M., Wolpert, D. M., & Frith, C. D. (2013). Modulation of somatosensory processing by action. *NeuroImage*, *70*, 356–362. https://doi.org/10.1016/J.NEUROIMAGE.2012.12.043

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. https://doi.org/10.1038/nrn3838

Tukey, J. W. (1962). The Future of Data Analysis. *The Annals of Mathematical Statistics*, *33*(1), 1–67.

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. https://doi.org/10.3389/fninf.2013.00014

Williams, S. R., & Chapman, E. C. (2000). Time-course and magnitude of movement-related gating of tactile detection in humans. II. Effects of Stimulus Intensity. *Journal of Neurophysiology*, *84*(2), 863–875. https://doi.org/https://doi.org/10.1152/jn.2002.88.4.1968

Wolpe, N., Ingram, J. N., Tsvetanov, K. A., Geerligs, L., Kievit, R. A., Henson, R. N., Wolpert, D. M., Rowe, J. B., Tyler, L. K., Brayne, C., Bullmore, E., Calder, A., Cusack, R., Dalgleish, T., Duncan, J., Matthews, F. E., Marslen-Wilson, W., Shafto, M. A., Campbell, K., … Rowe, J. B. (2016). Ageing increases reliance on sensorimotor prediction through structural and functional differences in frontostriatal circuits. *Nature Communications*, *7*(1), 13034. https://doi.org/10.1038/ncomms13034

Wolpe, N., Zhang, J., Nombela, C., Ingram, J. N., Wolpert, D. M., & Rowe, J. B. (2018). Sensory attenuation in Parkinson’s disease is related to disease severity and dopamine dose. *Scientific Reports*, *8*(1), 15643. https://doi.org/10.1038/s41598-018-33678-3

Wolpert, D. M., & Flanagan, J. R. (2001). Motor prediction. *Current Biology*, *11*(18), R729–R732. https://doi.org/10.1016/S0960-9822(01)00432-8

Yon, D., Gilbert, S. J., De Lange, F. P., & Press, C. (2018). Action sharpens sensory representations of expected outcomes. *Nature Communications*, *9*(1), 4288. https://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*(6), 1073–1083. https://doi.org/10.1037/xhp0000385

Yon, D., Zainzinger, V., Lange, F. de, Eimer, M., & Press, C. (2020). Action biases perceptual decisions toward expected outcomes. *Journal of Experimental Psychology: General*, Advance online publication. https://doi.org/10.1037/xge0000826

Action enhances predicted touch: Supplementary Information

Supplementary Methods

Tactile stimuli. The tactile forces were mechanically delivered to participants' left index and middle fingers in these experiments via solenoids (index finger only in Experiment 1). The tactile forces were presented via a soundcard. In order to achieve tactile forces that could be perceived as a single tap we presented a non-sinusoidal half square wave of low frequency (16.67 Hz) for 30 ms for each individual force. Stimulus intensity was manipulated via the amplitude of the square wave and seven logarithmically spaced amplitudes were selected to provide a range of detectable intensities. The fourth intensity in the range of target forces was used as the reference force.

Experiment 1 design. In Contact blocks, participants' right hand was positioned 5 cm above their left hand, and in No Contact blocks it was moved to 12 cm above to allow movements to be made without touching the lower shelf holding the motion tracker. It is worth noting that the difference in palm separation generated a difference between our Contact and No Contact conditions additionally to contact. However, this allowed us to replicate the typical setup (e.g., Bays et al., 2005, 2006; Kilteni et al., 2019) in the Contact condition while allowing movement to be registered with the infrared tracking and ensuring no contact with external surfaces in the No Contact condition. Importantly, conclusions relate primarily to the simple effects (between active and passive trial types) within the Contact and No Contact conditions, so this additional difference should not alter the conclusions.

Actions recorded by the motion tracker. In all conditions of all experiments, except for the Contact condition of Experiment 1, actions were recorded using a motion tracker (Leap Motion Controller using the Matleap MATLAB interface) that measured and tracked the coordinates of fingers and hand positions throughout trials. Tactile stimulation was triggered once the task-relevant right-hand finger had traversed at least 20 mm from its initial starting position in the y-axis. In Experiment 3, we stored these coordinates and calculated movement distance and duration relative to the starting coordinates. Distance traversed (Expected: M = -25.20 (mm), SD = 1.89, Unexpected: M = -25.02 (mm), SD = 1.75, $t(29)$ = -1.38, p = .178, d = -.25) and duration (Expected: M = .63 (s), SD = .19, Unexpected: M = .63 (s), SD = .20, $t(29)$ = -

.23, p = .817, d = -.04) were similar in both conditions – an unsurprising similarity given that participants were unaware of the outcome at the time of action execution. It is likely that actions were performed similarly in Experiments 1 and 2, given that movements were observed by the experimenter and corrected via instruction.

Supplementary Results

Experiment 3 additional analysis. We conducted an additional analysis of the PSE effects in Experiment 3 to determine whether any degradation in memory for the reference across each mini-block (21 trials) interacted with the reported expectation effect. We calculated PSEs separately for responses in the first (trials 1-10) and second half (trials 11-21) of blocks, separately for expected and unexpected trials. This calculation generated four PSE measures per participant – Expected first half (trials 1-10), Expected second half (trials 11-21), Unexpected first half, and Unexpected second half. These values were analysed in a 2x2 within-participants ANOVA, revealing a main effect of Expectation ($F(1, 29) = 6.85$, $p = .014$ $\eta p^2 = .19$), equivalent to our main finding in Experiment 3. However, there was no main effect of Block half ($F(1, 29) = 1.07$, $p = .310$, $\eta p^2 = .035$) nor a significant interaction between Expectation and Block half ($F(1, 29) = .002$, $p = .966$, $\eta p^2 < .001$). These results therefore suggest it is unlikely that any degradation in memory for the reference is interacting with the reported expectation effects.

Cross-Experiment comparisons. It is notable that PSEs in Experiment 1 and 3 are numerically higher than 4, whereas those in Experiment 2 are numerically lower than 4. In principle, the particular deviation of a given PSE value from 4 could be taken as reflective of whether there is enhancement or attenuation of the target event. Caution should be exercised with this interpretation due to the number of factors that would influence perceptual decisions (e.g., biases to select the first of two options; Firestone & Scholl, 2016), which is why our conclusions relate to differences between conditions within each experiment. Nevertheless, to examine this trend we computed the average PSE for each participant in each experiment and performed a one-way ANOVA with the factor of Experiment (1, 2, and 3). The results demonstrated no significant main effect of Experiment ($F = 2.96$, $p = .057$, $\eta p^2 = .064$), and Bonferroni corrected pairwise comparisons revealed no significant differences between Experiments 1 ($M = 4.34$, SD

= .67) and 2 (*M* = 3.82, *SD* = .84, p = .06), 1 and 3 (*M* = 4.18, *SD* = .97, *p* = 1.00), or 2 and 3 (*p* = .306). These trends may therefore reflect pure noise due to different samples, or perhaps also some contributions of non-predictive influences of action on perception (Press & Cook, 2015).

Supplementary References

Bays, P. M., Flanagan, J. R., & Wolpert, D. M. (2006). Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biology*, *4*(2), 281–284. https://doi.org/10.1371/journal.pbio.0040028

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*(12), 1125–1128. https://doi.org/10.1016/j.cub.2005.05.023

Kilteni, K., Houborg, C., & Ehrsson, H. H. (2019). Rapid learning and unlearning of predicted sensory delays in self-generated touch. *ELife*, *8*(e42888). https://doi.org/10.7554/eLife.42888

Press, C., & Cook, R. (2015). Beyond action-specific simulation: domain-general motor contributions to perception. *Trends in Cognitive Sciences*, *19*(4), 176–178. https://doi.org/10.1016/j.tics.2015.01.006