

RESEARCH ARTICLE

Computational Neuroscience

Divisive attenuation based on noisy sensorimotor predictions accounts for excess variability in self-touch

 Nicola Valè,^{1,2}  Ivan Tomić,^{1,3} Zahara Gironés,¹  Daniel M. Wolpert,⁴ Konstantina Kilteni,^{5,6} and  Paul M. Bays¹

¹Department of Psychology, University of Cambridge, Cambridge, United Kingdom; ²Department of Neurosciences, Biomedicine and Movement Sciences, University of Verona, Verona, Italy; ³Department of Psychology, Faculty of Humanities and Social Sciences, University of Zagreb, Zagreb, Croatia; ⁴Zuckerman Mind Brain Behavior Institute and Department of Neuroscience, Columbia University, New York City, New York, United States; ⁵Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The Netherlands; and ⁶Department of Neuroscience, Karolinska Institutet, Stockholm, Sweden

Abstract

When one part of the body exerts force on another part, the resulting tactile sensation is perceived as weaker than when the same force is applied by an external agent. This phenomenon has been studied using a force matching task, in which observers were first exposed to an external force on a passive finger and then instructed to reproduce the sensation by directly pressing on the passive finger with a finger of the other hand: healthy participants consistently exceeded the original force level. However, this exaggeration of the target force was not observed if the observer generated the matching force indirectly, by adjusting a joystick or slider that controlled the force output of a motor. Here, we present the first detailed computational account of the processes leading to the exaggeration of target forces in the force-matching task, incorporating attenuation of sensory signals based on motor predictions. The model elucidates previously unappreciated contributions of multiple sources of noise, including memory noise, in determining matching force output, and it shows that quantifying attenuation as the discrepancy between direct and indirect self-generated forces isolates its predictive component. Our computational account makes the prediction that attenuated sensations will display greater trial-to-trial variability than unattenuated ones because they incorporate additional noise from motor prediction. Quantitative model fitting of new and existing force-matching data confirmed the prediction of excess variability in self-generated forces and provided evidence for a divisive rather than subtractive mechanism of attenuation, while highlighting its predictive nature.

NEW & NOTEWORTHY We formulate a detailed computational account of sensory attenuation in force-matching tasks that disambiguates contributions of perceptual, memory, and prediction noise to isolate a pure measure of attenuation strength. Analysis of data from nearly 500 participants shows that attenuated sensations display increased trial-to-trial variability, consistent with incorporating additional noise inherent to motor prediction. These results support a divisive, rather than subtractive, reduction in the sensation of self-generated forces based on predicted reafference.

forward model; perception; sensory attenuation

INTRODUCTION

The same tactile stimulus is perceived differently when it is the result of your own action compared with when it has an external source. This phenomenon is often referred to as sensory attenuation, as a self-generated stimulus is typically perceived as weaker than an externally generated one. A

predictive mechanism has been proposed for this phenomenon (1–6), inspired by neurophysiological mechanisms of sensory cancellation (7–10). According to the predictive account, when the motor system generates a voluntary movement, a duplicate of the motor command (an efference copy) is provided to a forward model that integrates it with an estimate of the state of the body to predict sensory



Correspondence: P. Bays (pmb20@cam.ac.uk).
Submitted 4 February 2025 / Revised 17 March 2025 / Accepted 9 May 2025



consequences of the movement (11). This prediction is used to attenuate incoming tactile signals (reafference) that are a consequence of the action.

Shergill et al. (2) described a force-matching task, in which participants were asked to reproduce the tactile sensation of a target force, delivered to their passive left index finger by a lever attached to a torque motor (Fig. 1A). The reproduction was performed in two conditions: participants either applied a force to the left index finger by pressing on the lever with their right index finger (direct condition, Fig. 1B) or they used their right index finger to adjust the setting of a response device that controlled the force output of the torque motor (indirect condition, Fig. 1C). In both conditions, the participants were asked to modulate the force on their passive left finger until it matched the tactile sensation of the preceding target force.

The results showed that healthy participants produced matching forces that consistently exceeded the target forces in the direct condition, whereas matching forces were much closer to veridical in the indirect condition. Since the manner of presentation of the target force was identical in both conditions, the authors concluded that the exaggerated matching forces in the direct condition were caused by a relative attenuation of the tactile sensation induced in the passive finger by the participant's action. Moreover, as it was only in the direct condition that the force applied by the active finger would naturally cause a tactile sensation in the passive finger, the authors suggested that this attenuation was based on an internal prediction of the expected sensory consequences of the action. In the indirect condition, the relationship between movement of the active finger and force in the passive finger was mediated by the response device in a novel manner that has previously been shown to impair sensorimotor prediction (12).

These results have been replicated and the predictive account corroborated in a range of studies that have investigated the sensory attenuation phenomenon using different experimental settings and manipulations (3, 4, 13). Nevertheless, a detailed understanding of the computational basis of sensory attenuation has remained elusive. In particular, previous studies have debated whether


attenuation is based on a subtractive mechanism, in which the sensation of an external force exceeds that of a self-generated force by a fixed amount irrespective of the force magnitude, or a divisive mechanism, in which the level of attenuation scales with the amplitude of the predicted force (5, 14). This is significant because a subtractive mechanism could operate with only a coarse expectation that a self-generated force will be felt, whereas a divisive mechanism implies a fine prediction of the force amplitudes resulting from self-action.

Although models of sensory attenuation have been presented in previous work at varying levels of computational detail (1, 15–17), these models have not addressed the perceptual and cognitive demands specific to force reproduction. Here, we present the first detailed model of the computations that lead to the exaggeration of target forces in the force-matching task. We first validated the model framework by testing a prediction about trial-to-trial variability in matching force in new and previously published datasets. Our hypothesis was that the variability of the matching force would be greater in the direct condition compared with the indirect one, even when the mean matching force was equated. We then formally compared model variants with divisive and subtractive mechanisms, finding decisive evidence for a divisive attenuation process. Our findings expand our understanding of the attenuation process and provide further evidence against nonpredictive gating (18–20) as an account of force-matching responses.

MATERIALS AND METHODS

Model

Our model of the processes involved in the force matching task is set out in Fig. 1D. The black path is involved in both direct and indirect conditions, whereas the blue path is unique to direct force generation. The filled circles represent the observable quantities: the target force exerted by the torque motor (F_T) and the matching force exerted by the participant (F_M). These are the only measured variables in the



experiment. Perceptual noise corrupts the participant's tactile perception of both the target force and the matching force (s_T and s_M , respectively). Moreover, after experiencing the target force, the participant has to hold that sensation in memory (m_T) for comparison with the matching force, and this process may introduce additional variability and bias (memory noise).

During the reproduction phase, the participant adjusts their output force F_M until the sensation of the force in the passive finger ($\tilde{s}_M$) is equal to the remembered sensation of the target force m_T , again with the possibility of some inherent variability (comparison noise). In the direct condition, an efference copy of the motor output from the active hand (f_{copy}) provides input to a forward model that predicts the tactile sensation it will produce on the passive finger (s_{pred}). This prediction is combined with the afferent sensory signal s_M to produce an attenuated sensation $\tilde{s}_M$. In the indirect condition, where the causal relationship between active finger movement and passive finger sensation is unnatural, we assume the sensation of the matching force is not attenuated, i.e., $\tilde{s}_M = s_M$.

This formal model identifies two important features of the force-matching task. First, consistent differences observed between target and matching forces cannot be straightforwardly ascribed to the attenuation process, as they could also reflect biases in memory of the target force. Therefore, rather than comparing matching forces to the target force, we can compare matching forces generated in direct and indirect conditions in response to the same target force. Such a comparison has identical memory and decision components, and therefore, it allows us to isolate the attenuation process. Second, the sources of noise on the predictive pathway (blue) are expected to produce trial-by-trial variability in attenuation, which should translate into greater variability in the matching force when the attenuation process is active, i.e., in the direct condition as compared with the indirect condition. Based on these features, we analyzed data from the force-matching task, focusing on differences between direct and indirect conditions in mean amplitude and trial-to-trial variability of the matching force.

Data

In this study, we combined analysis of new data and reanalysis of previously published data. Specifically, we included four datasets from experimental studies on force-matching tasks. We extended a study by Kilteni and Ehrsson (21), adding new data from 93 participants to the existing dataset of 28 participants. The experimental protocol was identical for all 121 participants and is summarized below; see Ref. 21 for

full methodological details. The other three datasets were from previously published studies on this topic (5, 22, 23).

Table 1 outlines the sample characteristics and methodological details of each study. In all cases, target force levels were uniformly distributed across the stated range and tested with equal frequency in a pseudorandomized order of trials. In each trial, participants were given 3 s to reproduce the target force, and the matching force was computed as the mean force exerted in the time interval between 2 and 2.5 s after the go signal.

In Studies 1–3, trials were equally divided into blocks of direct and indirect force matching, with Studies 1 and 2 using a linear potentiometer (or “slider”) and Study 3 using a joystick to control force output in the indirect condition. In the direct condition, participants in Studies 2 and 3 pressed on the top surface of the lever to reproduce the target force, whereas in Studies 1 and 4 they pressed on a force sensor positioned above the lever, while the torque motor continuously transmitted the recorded force to the passive finger, closely mimicking direct contact between the fingers. The force was always transmitted veridically in Study 1, whereas trials in Study 4 were equally divided into three blocks in which the gain between sensed and transmitted force was varied, so the force generated on the passive finger was 2× (double), 1× (veridical), or 0.5× (half) that of the participant's active finger press. As a consequence, the force generated by and tactile feedback received in the active finger varied across gain conditions, while participants produced approximately the same matching force in the passive finger.

We excluded from the analyses four participants from Study 1 and two participants from Study 2 since they had missing data and two participants from Study 2 since the matching forces they produced did not correlate with the target forces, suggesting that they likely did not properly understand the task.

Participants in Studies 2 and 3 were divided into subgroups that were analyzed separately. Participants from Study 2 were part of the Cambridge Center for Aging and Neuroscience (24, 25) cohort of healthy adults spanning the adult age range, and they were separated by age into young adult, middle-aged, and older adult subgroups, following the same criteria as the source study (23). Study 3 comprised a patient subgroup with schizophrenia and an age-matched healthy control subgroup.

Studies 1–3, although using slightly different apparatus and including different populations, investigated the sensory attenuation phenomenon using the same force-matching task, allowing their results to be directly compared. Replicating our analyses across these three independent

Table 1. Summary of datasets

Study	Source	Target Range (n)	Target Levels	Trials Per Participant	Subgroup/Condition	n	Age (Yr)
1	Kilteni and Ehrsson [(21), extended]	1–3.5	6	72		121	27.9 ± 4.9
2	Wolpe et al. (23)	1–2.5	4	64	Young	79	30.7 ± 5.9
					Middle	131	51.9 ± 7.2
					Older	108	75.4 ± 6.2
3	Shergill et al. (22)	0.5–2.75	5	160	Healthy	19	35.9 ± 14.0
					Patients	19	36.4 ± 13.4
4	Bays and Wolpert (5)	1–3	5	50	2× gain	16	18–40
				50	1× gain		
				50	0.5× gain		

datasets strengthens the results and demonstrates the generalization of our findings. In contrast, *Study 4* was analyzed and presented separately from *Studies 1–3*, as it did not have an indirect force-matching condition, instead having a within-participant manipulation of gain that made it specifically suited for evaluating the gating account of attenuation.

Analysis

To assess qualitative predictions of the model framework for direct versus indirect force matching, we first computed, for each participant in each of *Studies 1–3*, the mean ($\bar{F}_M$) and standard deviation (σ_{F_M}) of their matching forces over all trials with each target force level (F_T), separately for each condition. We then performed a linear mixed-effects model (*fitlme* in MATLAB) with subject as a random factor on these measures to generate lines of best fit and estimates of slopes and intercepts. *Study 4* was omitted from this analysis as it had no indirect force-matching condition.

When performing statistical inference on regression parameters, we used the Bayesian approach implemented in MATLAB (<https://klabhub.github.io/bayesFactor/>) with the default Jeffreys–Zellner–Siow prior on effect sizes (26). The reported Bayes factors compare the predictive adequacy of two competing hypotheses (e.g., alternative and null) and quantify the change in belief that the data bring about for the hypotheses under consideration (27). For example, $BF_{10} = 10$ indicates that the data are 10 times more likely to occur under the alternative hypothesis (i.e., there is a difference) than under the null hypothesis (i.e., there is no difference). Evidence for the null hypothesis is indicated by $BF_{10} < 1$, in which case the strength of evidence is indicated by $1/BF_{10}$.

Formal Model Specification

For quantitative fitting, we developed two implementations of the attenuation model illustrated in Fig. 1, incorporating either a subtractive or divisive mechanism.

Subtractive Model

According to the subtractive model, the tactile sensory attenuation is fixed, irrespective of the intensity of the stimulus. That is, given an afferent sensory signal s_D , the participant perceives an attenuated sensation $\tilde{s}_D = s_D - K$, where K is the subtractive attenuation. We assume K is normally distributed with mean $\bar{K}$ and standard deviation σ_K . We further assume that, for a given target force, the matching force in the indirect condition, F_I , is normally distributed with mean $\bar{F}_I$ and SD σ_{F_I} , respectively. This matching force contains all the components of perception for the direct condition, except for any attenuation. Consequently, the matching force in the direct condition, F_D , for the same target force will also be normally distributed with mean and SD,

$$\begin{aligned}\bar{F}_D^{\text{sub}} &= \bar{K} + \bar{F}_I, \\ \sigma_{F_D^{\text{sub}}} &= \sqrt{\sigma_K^2 + \sigma_{F_I}^2}.\end{aligned}$$

In contrast to the divisive model (see *Divisive Model*), the skewness of matching forces predicted by the subtractive model is always zero.

Divisive Model

The divisive model is based on the hypothesis that the attenuation is proportional to the stimulus intensity, i.e., $\tilde{s}_D = s_D/K$. Again, assuming K and F_I are normally distributed, the matching force in the direct condition is the product of two normally distributed variables. This distribution function does not in general have a closed-form expression, but can be approximated by a skew-normal distribution matched for mean, SD, and skewness [(28); see APPENDIX for derivations]:

$$\begin{aligned}\bar{F}_D^{\text{div}} &= \bar{K} \cdot \bar{F}_I, \\ \sigma_{F_D^{\text{div}}} &= \sqrt{\bar{F}_I^2 \sigma_K^2 + (\bar{K}^2 + \sigma_K^2) \sigma_{F_I}^2}, \\ &= \sigma_K \sigma_{F_I} \sqrt{1 + \delta_K^2 + \delta_{F_I}^2}, \\ \text{Skew}[F_D^{\text{div}}] &= \frac{6\mu_x\mu_y\sigma_x^2\sigma_y^2}{\text{Var}[Z]^{3/2}} = \frac{6\delta_x\delta_y}{(1 + \delta_x^2 + \delta_y^2)^{3/2}},\end{aligned}$$

where we have used, $\delta_K := \frac{\bar{K}}{\sigma_K}$ and $\delta_{F_I} := \frac{\bar{F}_I}{\sigma_{F_I}}$.

Hierarchical Bayesian Modeling

We used hierarchical Bayesian estimation to fit the models to the data (29, 30). In this modeling framework, a unique set of parameters is estimated for each individual participant, but all instances of each parameter are considered samples from a population-level distribution. Consequently, the estimation of each individual-level parameter is simultaneously informed by data from all the other participants—this occurs because all individual data contribute to the estimation of the population-level parameters, which, in turn, constrain all the individual-level parameters. Therefore, the hierarchical models offer a twofold advantage compared with traditional fitting approaches: these models share informative power across model hierarchy, which is particularly beneficial in scenarios with limited data (e.g., a low number of trials per participant, as is the case in some of the datasets we examined), and also for regularization of parameter estimates across individuals, enhancing the reliability of these estimates (31). An important feature of the Bayesian approach is that it provides a posterior distribution over potential parameter values, rather than just point estimates of the best-fitting parameters as is the case with classical modeling approaches (e.g., ML). This full posterior distribution not only offers insights into the best point estimate (e.g., MAP) but also provides information about the uncertainty of parameter estimation (e.g., the width of the posterior distribution).

Data Fitting

We fit subtractive and divisive models (see *Formal Model Specification*) separately to each subgroup in each of *Studies 1–4* (Table 1). We simultaneously estimated the posterior distributions of parameters by using the individual trial data from all conditions within each experiment. Detailed model descriptions, along with specified prior distributions, can be found in the APPENDIX. In brief, we fit a separate set of parameters for each input force level in the indirect force reproduction task. Using those measured matching forces, we estimated the attenuation parameters to fit the direct force reproductions.

To obtain the hierarchical fit, we used the Gibbs sampler with the Markov Chain Monte Carlo (MCMC) algorithm implemented in JAGS (32), in combination with the *matjags* package (33) for MATLAB (The MathWorks Inc.). Posterior distributions for all parameters were obtained by running three parallel chains, with a thinning factor of two, and discarding the initial 1,000 samples for burn-in, resulting in a total of 10,000 valid samples per chain. After obtaining the samples for each dataset and model, we performed standard checks for convergence of chains [$\hat{R}$ (34)] and autocorrelation within chains, and found the satisfactory performance of the sampler for both individual- and population-level parameters. To compare how well the two competing models fit the data, we used Deviance Information Criterion (DIC) as a comparison metric. DIC is a generalization of Akaike's Information Criterion (AIC), and it similarly depends on the goodness of fit and model complexity. The model with a smaller DIC is expected to be the one that would best predict and replicate the data.

We summarize the obtained posterior distributions numerically by reporting the median of posterior samples and the 90% highest density interval (HDI). The 90% HDI includes parameter values that have at least some minimal level of posterior credibility as opposed to values outside the HDI (35). Specifically, the probability density is greater for any parameter value within the 90% HDI compared with values outside it, and the total probability of values within the 90% HDI is 90%. Consequently, the width of the HDI serves as a measure of the uncertainty associated with the parameter estimates, i.e., when the HDI is narrow, we can be relatively certain about the parameter estimates.

When analyzing differences in parameter values between different subgroups of participants, we first calculate the group mean difference (MD) by finding the difference between the posterior distributions of two groups. We then report 95% HDI of the difference scores to quantify a range of values that has a 95% probability of containing the true difference in parameter estimates for two groups. Because parameter values with higher density are interpreted as more credible than parameter values with lower density, we set a decision rule to reject a mean difference of zero (i.e., no difference) when the 95% HDI does not contain zero. Intervals that include zero suggest zero as a credible difference value, indicating no difference between groups.

To assess the descriptive adequacy of the model, we generated posterior predictive distributions. The posterior predictive distribution shows what data to expect given the model at hand and the available knowledge, which includes the posterior distribution over parameters informed by the observed data (30). The posterior predictive distribution is a distribution over data and therefore indicates the relative probability of different outcomes (i.e., output forces) after the model has seen the data. By comparing this predicted data to the data already observed, we can evaluate the model's adequacy. We generated posterior predictive distributions of output forces in both direct and indirect reproduction conditions by using randomly sampled parameters from individual-level posterior distributions and the generative process described for each model. We conducted 10,000 simulations of output forces for each participant and condition. We then assess the resemblance between simulated and observed data

through a combination of qualitative and quantitative methods. Qualitatively, we visually examine graphical displays of observed and simulated data.

In addition, we quantitatively assess the fit by calculating the Bayesian P value (not to be confused with the frequentist P value). This value quantifies the probability that the obtained parameters will produce simulated data that deviates more from the predicted (average) data than the observed data does (36). The Bayesian P value of 0.5 indicates that observed and simulated data deviate equally around the model's prediction. If the model's predictions exhibit bias, meaning they are too extreme in either direction, the values will tend to deviate from 0.5, approaching zero or one. In such cases, one should consider rejecting the model and exploring alternative models. Given that the tested models are designed to overfit the data in the indirect reproduction conditions, and hence are expected to produce predictions very similar to the data, we focus on Bayesian P values in the direct reproduction condition.

RESULTS

We first examined whether the predictions of the model were qualitatively consistent with behavior in three existing datasets where participants performed force matching in both direct and indirect conditions. Figure 2A illustrates the typical pattern of matching forces using data from *Study 1* as an example [Kiltner and Ehrsson (21); see Fig. A2 and Fig. A3 for other datasets]. The relationship between target force and mean matching force in each condition was close to linear, as shown by lines of best fit in Fig. 2A (linear regression model $R^2 = 0.94$; for all subgroups of *Studies 1–3*, $R^2 > 0.91$). In line with typical findings, directly generated matching forces (cyan symbols) consistently exceeded target forces (dashed line of equality) at every force level. As a function of target force, the matching force showed increases in both slope (>1 , $BF_{10} = 1.92 \times 10^{18}$) and intercept (>0 , $BF_{10} = 3.35 \times 10^9$) (cyan bars in Fig. 2C), a pattern that was also observed in *Study 2* (23) for every age-based subgroup (young subgroup: slope $BF_{10} = 9.13 \times 10^6$, intercept $BF_{10} = 1.58 \times 10^5$; middle-age subgroup: slope $BF_{10} = 1.22 \times 10^{14}$, intercept $BF_{10} = 1.21 \times 10^{11}$; older group: slope $BF_{10} = 2.64 \times 10^4$, intercept $BF_{10} = 3.74 \times 10^9$). The healthy subgroup in *Study 3* (22) showed a significant increase in intercept only ($BF_{10} = 6.63 \times 10^5$; slope $BF_{10} = 0.24$), while the patient subgroup (which displayed reduced attenuation overall according to the original study) showed an increased intercept ($BF_{10} = 2.14 \times 10^8$) paired with a decrease in slope ($BF_{10} = 7.22$).

Consistent with classical findings, reproduction of target forces was substantially more accurate when the matching force was generated indirectly (red symbols in Fig. 2A). Nonetheless, in *Study 1*, linear regression coefficients indicated deviations from equality in both slope (<1 , $BF_{10} = 20.55$) and intercept (>0 , $BF_{10} = 2 \times 10^{31}$). This pattern was replicated in the other studies and subgroups (red bars in Fig. 2C), with the exception of the young adult subgroup (strong evidence for an increase in intercept, $BF_{10} = 203$; moderate evidence against a change in slope, $BF_{10} = 0.17$) and the middle-aged subgroup (weak evidence for a reduction in slope, $BF_{10} = 1.13$; moderate evidence against a change in intercept, $BF_{10} = 0.23$) of *Study 2*. These results are consistent with predictions of our model




Figure 2. Matching forces in direct and indirect force reproduction. **A:** mean matching force as a function of target force in *study 1*, for direct (cyan) and indirect (red) conditions. Data are shown as symbols and linear fits as lines. **B:** matching force in the direct condition of *study 1* plotted against matching force in the indirect condition, with a linear fit. Symbols indicate mean matching forces for each target force level. **C:** slope (top) and intercept (bottom) parameters for each data set and subgroup, for the regression of direct (cyan) and indirect (red) matching force against target force (as in **A**) and direct matching force against indirect matching force (as in **B**; purple). Error bars in **A** and **B** indicate $\pm$ SE. Shaded areas in **A** and **B** and error bars in **C** indicate 95% CI. In **A** and **C** (cyan and red), a slope greater than one indicates exaggeration of the matching force proportional to the target force magnitude, while the intercept greater than zero indicates exaggeration of the matching force irrespective of the target force magnitude. In **B** and **C** (purple), the slope greater than one means that the matching force in the direct condition was proportionally higher than the matching force in the indirect condition, while intercept greater than zero indicates a systematic offset between direct and indirect matching forces. CI, confidence interval.

of the task (Fig. 1D), in which accumulated bias and variability in memory for the sensation of the target force, possibly supplemented by noise in the comparison process itself, cause matching forces to deviate from target forces even in the absence of predictive attenuation. More specifically, the particular pattern observed across studies for the indirect condition, of a decrease in slope paired with an increased intercept, is consistent with the classical finding of contraction bias in memory (37, 38), whereby memory for the current target force is influenced by the history of target forces on preceding trials, leading to a bias towards the mean of the presented forces. Notably, in *Study 2*, we observed a progressive flattening of indirect matching slopes with age, consistent with the well-established age-related decline in short-term memory performance (39, 40).

Isolating Effects of Predictive Attenuation

The results above are typical of a wide range of previous force-matching studies, in which healthy participants were found to be close to accurate in indirect force reproduction but substantially exceeded target forces in direct reproduction. However, when analyzing slopes and intercepts, there is considerable variation across groups and studies, making it difficult to discern a consistent pattern. We hypothesized that this variation was at least partly due to the memory and comparison components identified in our model of the task, rather than the attenuation of perceived force in the direct condition. As these nonpredictive components are identical in direct and indirect conditions, the effects of predictive attenuation should be better isolated by comparing matching forces in the two conditions for the same target force. Figure 2B plots data from *Study 1* in this way (see Fig. A2 and Fig. A3 for other datasets). This relationship was again close to linear ($R^2 = 0.97$; for all subgroups of *Studies 1–3*, $R^2 =$

0.96), with slopes and intercepts plotted as purple bars in Fig. 2C. Relative to results from the individual conditions (cyan and red bars), comparing direct to indirect matching forces reduced variability across subgroups, revealing consistent evidence for a strong deviation in slope in every subgroup (excepting patients in *Study 3*) and a relatively reduced (though still nonzero) intercept.

This approach also clarified results for the patients in *Study 3* (rightmost bars in Fig. 2C), indicating that the greater accuracy of force matching in patients with schizophrenia reflected a strong reduction in slope compared with age-matched controls, approaching veridicality (slope of 1), but a negligible change in intercept. In *Study 2*, the increasing exaggeration of target forces with age was observed in both slope and intercept, once memory effects were accounted for.

Evidence for Excess Variability in Direct Force Reproduction

Our model of force matching predicts excess variability in the direct compared to indirect condition (Fig. 1D), arising from variability in the predictive processes involved in sensory attenuation (i.e., efference copy noise and prediction noise in Fig. 1D). Figure 3A plots within-participant across-trial variability in matching force as a function of target force level for *Study 1*. Matching force variability was substantially greater in the direct condition (cyan) than in the indirect condition (red) at each target force level compared with the indirect condition. This finding was consistent across all datasets and subgroups as measured by differences in regression parameters between the two conditions (Fig. 3B; all $BF_{10} > 5.6$).

Although matching force variability was consistently larger in direct than in indirect condition for the same target




Figure 3. Trial-to-trial variability in matching force. **A:** average of within-subject standard deviation of matching force as a function of target force in *study 1*, for direct (cyan) and indirect (red) conditions, with linear fits. **B:** slope (top) and intercept (bottom) parameters for each data set and subgroup, for the regression of direct (cyan) and indirect (red) matching force SD against target force (as in **A**). **C:** average of the within-subject standard deviation of matching force as a function of mean matching force in *study 1*, with linear fits. Each symbol corresponds to a single target force level. **D:** parameters of the linear regression of matching force SD against mean matching force (as in **C**) for each data set and subgroup. In **A** and **B**, a slope greater than zero indicates that matching force variability increases with increasing target force, whereas an intercept greater than zero indicates the presence of trial-to-trial variability irrespective of the target force magnitude. In **C** and **D**, a slope greater than zero indicates that matching force variability increases with increasing mean matching force, whereas an intercept greater than zero indicates the presence of trial-to-trial variability that does not depend on the mean matching force magnitude. Error bars in **A** and **C** indicate $\pm$ SE. Shaded areas and error bars in **B** and **D** indicate 95% CI. CI, confidence interval; SE, standard error.

force, the mean matching force was also higher in the direct condition. Variability of force output has been found to scale with the mean amplitude (41, 42), so to examine whether this was a viable explanation for the increased variability, we next regressed matching force SD in the two conditions against mean matching force (Fig. 3C shows results from *Study 1*; Fig. 3D plots regression coefficients for each dataset and subgroup). As predicted, variability increased with mean force amplitude in both direct and indirect conditions (positive slope in all subgroups; all $BF_{10} > 1644$). However, matching force variability was consistently greater in direct than in indirect condition for the same mean force amplitude (cyan symbols above red symbols in Fig. 3C), an effect primarily driven by increased slope in the regression model for every dataset and subgroup (cyan bars vs. red bars in Fig. 3D, top; healthy subgroup $BF_{10} = 1.1$, all other $BF_{10} > 591$).

Subtractive versus Divisive Models of Attenuation

Having confirmed qualitative predictions of the model framework, we next computed formal model fits for each of the datasets and subgroups. This required us to specify how the attenuation process quantitatively changes the sensation of force: we considered two variant models, one in which the attenuation was subtractive and one in which it was divisive. In the subtractive model, the attenuation process subtracts a fixed amount (the attenuation factor) from a perceived force only when it is directly self-generated: as a result, the matching force in the direct condition exceeds that in the indirect condition by a fixed amount. In the divisive model, a fixed fraction of a perceived force is instead removed when self-generated, equivalent to dividing it by the attenuation factor: as a result, the matching force in the direct condition equals the matching

force in the indirect condition multiplied by the attenuation factor. In both variant models, the attenuation factor was fixed with respect to different levels of target and matching force but incorporated trial-to-trial variability. Specifically, the attenuation factor was modeled as a normally distributed random variable with mean and SD as free parameters.

The Divisive Model Provides a Close Fit to Behavioral Data

Formal model comparison demonstrated that the divisive model provided a superior fit to behavioral responses on the force matching task compared to the subtractive model in all subgroups of *Studies 1–3* (*Study 1*, $\Delta\text{DIC} = 1039.4$; *Study 2*, $\Delta\text{DIC} \geq 201$ for all age-groups; *Study 3*, healthy participants: $\Delta\text{DIC} = 169.1$, patients: $\Delta\text{DIC} = 2.5$). Later, we present further results obtained from fitting the divisive model. Results for the subtractive model are available in the APPENDIX (Table A1 and Figs. A6 and A7).

The fit of the divisive model to data is exemplified by the posterior predictive distributions for *Study 1* shown as shaded areas in Fig. 4. These model predictions accurately reproduced the empirical distributions (gray histograms) of matching forces in direct (cyan) and indirect (red) conditions at each force level. The effect of attenuation in the direct condition is observed as an increase compared with the indirect condition in mean, variance, and skewness of the matching force distributions. Skewed distributions (note long tails in direction of larger matching forces in direct condition; Fig. 4) are a prediction of the divisive model, arising from multiplication of two unskewed random variables (attenuation factor and indirect matching force), and contribute to its superior fit over the subtractive model.

To quantify the quality of fit, we calculated Bayesian P values from the posterior predictive distributions in the direct condition. We observed Bayesian P values that were consistently around the ideal value of 0.5 ($0.47 \leq \text{all } P \text{ values} \leq 0.51$), indicating a close match between simulated and observed data. The divisive model provided a similarly close match to the data from *Studies 2* and *3* (Figs. A4 and A5, respectively). Again, the strong resemblance between the simulated and observed data was numerically verified by calculating the Bayesian P values for each dataset and target force level, consistently finding values around 0.5 both in *Study 2* subgroups (young $0.46 \leq \text{all } P \text{ values} \leq 0.52$; middle $0.47 \leq \text{all } P \text{ values} \leq 0.53$; older $0.47 \leq \text{all } P \text{ values} \leq 0.52$), and *Study 3* subgroups (healthy controls: $0.45 \leq \text{all } P \text{ values} \leq 0.52$) and patients: $0.44 \leq \text{all } P \text{ values} \leq 0.54$).

Mean and Variability of the Attenuation Factor

Figure 5A displays the observed posterior distributions for the population-level mean attenuation factors from the divisive model. We summarized these posterior distributions by calculating the median values. In *Study 1*, the estimated (posterior median) mean attenuation factor was 1.61 [90% highest density interval (HDI) = [1.56, 1.66], describes the range of values within which the parameter of interest lies with 90% probability]. This value indicates that, on average at the population level, a force generated indirectly through the response device was perceived as 61% more intense than the same force when directly self-generated; or equivalently, that perception of the directly self-generated force was attenuated by 38% ($1 - 1/1.61 = 0.38$). Across the different subgroups of *Studies 1–3*, the posterior distributions of




Figure 4. Posterior predictive densities for *Study 1*. Gray histograms show matching forces pooled across all participants. Colored distributions show simulated matching forces based on posterior predictive density. The circles with error bars show the median and 90% HDI for data (black) and simulated trials for the same condition (cyan: direct, red: indirect).




Figure 5. Posterior densities for the mean attenuation factor (A) and the SD of attenuation factor (B), based on population-level central tendency parameters of the divisive attenuation model fit separately to each dataset and subgroup. Black lines indicate the posterior median.

the mean attenuation factors exhibited clear heterogeneity. Importantly, none of the reported HDIs include the value of 1, indicating that in all examined datasets, the perception of force in direct self-generation was attenuated in comparison to the indirect condition.

The posterior distributions of the standard deviation of the attenuation factor in the divisive model are shown in Fig. 5B. In contrast to the mean attenuation factor, estimates of this parameter exhibited much greater uniformity between datasets. Specifically, in *Study 1*, the median parameter was found to be 0.152 (90% HDI [0.128, 0.174]), indicating that the perceived intensity difference between conditions varied from trial to trial about its mean of 61% with an SD of 15%. This is approximately equivalent to an attenuation of the directly self-generated force by $38\% \pm 6\%$. Similar values were found for the different subgroups of *Studies 2* and *3*.

Attenuation Scales with Age

Study 2 compared the mean exaggeration of the target force in the direct condition across the three groups based on age difference, providing an opportunity to examine how sensory attenuation varies with age. Estimated mean attenuation factors were 1.45 (90% HDI [1.365, 1.536]) in the sample of young adults, 1.71 (90% HDI [1.63, 1.793]) in the sample of middle-aged participants, and 1.97 (90% HDI [1.834, 2.096]) in the sample of older adults (Fig. 5A). This suggests an increase in mean attenuation factor with age, consistent with conclusions of the source study. To corroborate this observation statistically, we calculated the mean differences (MD) between posterior distributions from these groups. We found that young adults had a lower mean attenuation factor compared with both middle-aged adults (MD = -0.259 , 95% HDI [-0.402 , -0.12]) and older adults (MD = -0.514 , 95%

HDI $[-0.698, -0.322]$), while middle-aged participants had a lower estimate compared with older adults ($MD = -0.255$, 95% HDI $[-0.438, -0.067]$).

In contrast, the standard deviation of the attenuation factor exhibited greater similarity among the three subgroups, with a trend to decreasing SD with age (Fig. 5B). To quantify this, we again calculated MDs between posteriors in different groups and, in all cases, found that the 95% HDI encompassed zero, indicating no evidence for a difference in estimated parameters. Specifically, this was the case when comparing young and middle-aged adults ($MD = 0.022$, 95% HDI $[-0.033, 0.079]$), young and older adults ($MD = 0.042$, 95% HDI $[-0.025, 0.11]$), and middle-aged and older adults ($MD = 0.02$, 95% HDI $[-0.045, 0.082]$).

Patients with Schizophrenia Have Reduced Sensory Attenuation

The data in *Study 3* comprised groups of schizophrenic patients and neurotypical controls. Estimates of the mean attenuation factor were 1.708 (90% HDI $[1.498, 1.928]$) for controls and 1.372 (90% HDI $[1.213, 1.526]$) for patients (Fig. 5A). This suggests that, on average, healthy controls exhibited stronger sensory attenuation. To confirm this, we calculated the mean difference between the posterior distributions, revealing that the attenuation factor was indeed statistically larger in the group of healthy participants ($MD = 0.3361$, 95% HDI $[0.007, 0.649]$).

Comparing the estimated standard deviation of the attenuation factor between the two groups (Fig. 5B), we found that a trend for smaller SD in patients did not reach conventional significance, as indicated by a 95% HDI of MD containing zero ($MD = 0.111$, 95% HDI $[-0.008, 0.234]$). Similarly to the comparison between samples in *Study 2*, a clear difference in the estimated attenuation factors was not accompanied by a significant difference in their variability. Note, however, that the posterior estimates from this study exhibited relatively higher uncertainty in parameter estimates (i.e., broader posterior) compared to other subgroups, most likely due to the smaller number of participants in each subgroup.

Testing a Nonpredictive Gating Account of Attenuation

A key component of the proposed model is a predictive mechanism (the forward model) that estimates sensory consequences of motor commands, enabling self-generated sensory inputs to be attenuated at an early stage of processing. Alternatives to this account propose that sensitivity to the matching force in the passive finger is decreased through a nonpredictive gating mechanism triggered by muscle activation or tactile feedback in the active finger (18–20). Such nonpredictive gating has been clearly demonstrated in active or moving body parts (43–48) but not in passive ones. As the right index finger is active in indirect as well as direct conditions (moving the slider or joystick), this account additionally requires that the gating is weaker in the indirect condition, perhaps because less force is required to adjust the response device.

To investigate the role of motor activity and feedback in the active finger in force matching, we applied our model to the data from *Study 4*, which manipulated the transmission of force from the active to the passive finger (5). Figure 6A

illustrates forces produced by the active finger and forces transmitted to the passive finger across the target force levels in the three gain conditions.

For the hierarchical fitting, as our interest was in comparing the three different conditions of direct reproduction, we fixed parameters corresponding to indirect force generation for each participant and condition to the same set of population-level parameters estimated from *Study 1* (which had matching target force levels), and performed fitting on direct condition data separately for each gain condition. The observed population-level posterior distributions are shown in Fig. 6B. Although the gating hypothesis would predict less attenuation as the factor multiplying the input force increases, we obtained highly consistent estimates of the mean attenuation factor across the three conditions. Specifically, the estimated (posterior median) mean attenuation factors were 1.175 (90% HDI $[1.066, 1.281]$) for the gain factor of 0.5, 1.204 (90% HDI $[1.099, 1.312]$) for the gain factor of 1, and 1.191 (90% HDI $[1.108, 1.269]$) for the gain factor of 2. We then calculated the mean differences between posterior distributions of these groups and, in all cases, found that the 95% HDI encompassed zero, indicating no differences in estimated parameters between these groups ($MD_{0.5-1} = -0.03$, 95% HDI $[-0.211, 0.156]$; $MD_{0.5-2} = -0.016$, 95% HDI $[-0.176, 0.147]$; $MD_{1-2} = 0.014$, 95% HDI $[-0.148, 0.176]$).

A similar conclusion can be drawn when comparing the posterior distributions of the standard deviation of the attenuation factor (Fig. 6C). The obtained posterior estimates were 0.113 (90% HDI $[0.044, 0.18]$) for the gain factor of 0.5, 0.091 (90% HDI $[0.031, 0.138]$) for the gain factor of 1, and 0.117 (90% HDI $[0.065, 0.163]$) for the gain factor of 2. Calculating MDs between these posterior distributions, here we also found that the 95% HDI in all cases encompassed zero, indicating no differences in estimated parameters between these groups ($MD_{0.5-1} = 0.024$, 95% HDI $[-0.079, 0.129]$; $MD_{0.5-2} = -0.001$, 95% HDI $[-0.103, 0.099]$; $MD_{1-2} = -0.026$, 95% HDI $[-0.112, 0.064]$).

DISCUSSION

Here, we proposed a formal computational model of the processes leading to exaggeration of target forces in the force matching task (2). According to this model, matching forces recorded in both direct and indirect reproduction conditions are subject to identical sources of perceptual, memory, and comparison noise, differing only in the presence or absence of predictive attenuation. On this basis, we isolated a pure measure of attenuation strength by comparing the matching forces between the two conditions. This analysis confirmed and extended previous conclusions that attenuation is enhanced in older compared to younger individuals (23) and weaker in patients with schizophrenia compared to age-matched control participants (22).

Previous studies suggested that a network including cerebellum, secondary somatosensory cortex (SII), and premotor cortex is involved in sensorimotor prediction and sensory attenuation (e.g., Refs. 21, 49, and 50). Reduced connectivity in cerebellar-midbrain and cerebellar-thalamic networks has been proposed as a plausible functional basis for the role of the cerebellum in psychotic disorders such as schizophrenia (51, 52). The reduction in sensory attenuation previously observed




Figure 6. Evidence against a gating account. **A:** forces generated by active finger (empty symbols and dashed line) and forces applied to passive finger (filled symbols and solid line) during direct matching in the three different gain conditions of *Study 4* (passive and active forces are identical for gain = 1, *middle panel*). Error bars indicate $\pm$ SE; shaded areas indicate 95% CI. **B** and **C:** posterior densities for the mean attenuation factor and the SD of attenuation factor, based on the divisive attenuation model fit separately to each gain condition. Horizontal lines indicate the posterior median. SD, standard deviation.

in patients with schizophrenia (22), and further delineated by our analysis, could be related to impaired cerebellar function, weakening the ability to predict movement consequences or use the prediction to attenuate matched sensory input. On the contrary, the increased attenuation found in older participants suggests that cerebellar impairment cannot provide a straightforward account of all modulations of sensory attenuation.

Removing the confounding effect of memory bias proved particularly important for older individuals: in indirect force matching, we observed a bias toward the average target force that grew stronger with age. This is consistent with a classical contraction bias (37, 38) in which memory of the target force on the current trial is drawn toward the average of target forces on previous trials. The increase in bias with age may reflect a corresponding decline in working memory fidelity.

Our model made a specific prediction of excess trial-to-trial variability in the direct compared with the indirect condition, attributed to variation in attenuation level due to noise in motor prediction. We experimentally confirmed this prediction in every dataset, observing higher variability in the direct condition compared with the indirect condition, even when accounting for the different magnitudes of the mean matching forces. One possible counterargument is that, rather than prediction noise, the higher variability in the direct condition could stem from higher motor noise, since participants use their right index finger to press in the direct condition but use it to manipulate a response device in the indirect one. However, the force-matching studies were designed to minimize contributions of motor noise by averaging the force applied to the passive finger on each trial over a period of 500 ms. Moreover, if motor noise was responsible for the excess variability, the higher motor variability typically observed in older adults (53) would then be expected to produce greater

excess variability, whereas we observed comparable variability across all age groups. Similarly, the trend toward smaller excess variability observed in the patients with schizophrenia would imply that patients have lower motor variability compared to controls, contradicting previous findings (54).

We proceeded to test two alternative accounts of how attenuation affects predicted force sensations. According to a subtractive account of attenuation, a self-generated force is perceived as weaker than an externally generated force by a fixed amount that does not scale with the force amplitude. Supporting this notion, Walsh et al. (14) and Bays and Wolpert (5) found that attenuation in the direct condition mainly consisted of a fixed “offset” from the target force. However, neither of these studies compared the direct condition to an indirect condition, raising the possibility that their findings might reflect biases unrelated to prediction, e.g., the above-mentioned contraction bias. According to the alternative divisive account, the difference in perceived force scales with the magnitude of the force, with lighter forces being subject to less attenuation compared with stronger forces. In other words, these two alternative hypotheses would have different effects on the perception of the weakest tactile stimuli. Formal model comparison decisively supported this divisive account of attenuation, based on the differing predictions for the distributions of matching forces. The divisive account of attenuation straightforwardly explains why even the lightest self-generated touch remains perceptible, despite being attenuated relative to an equivalent external force.

As the direct and indirect conditions differ in the amplitude of force applied by the active finger, it has been argued (18–20) that overestimation of forces in the direct condition could arise from a non-predictive mechanism based on this activity, similar to the sensory gating that suppresses tactile input in a

moving effector regardless of its predictability (43), but now spreading to affect the passive finger of the other hand. Applying our modeling methods to previous data from a gain manipulation experiment (5) provided evidence against the gating hypothesis. Despite a fourfold variation across gain conditions in the force output in the active finger, we observed attenuation of similar magnitude and variability across all conditions. This result appears incompatible with nonpredictive gating due to activity in the hand applying the force, which varied widely across conditions, but is consistent with a mechanism that attenuates the sensation based on a prediction of the force received on the passive finger, which remained approximately constant across conditions. These results also appear to rule out any nonpredictive mechanism based on tactile feedback in the active finger, which similarly varied fourfold in amplitude across gain conditions.

These findings align with studies that found attenuation and gating are governed by separate mechanisms (13, 55, 56), and a range of previous observations supporting a predictive basis for attenuation of self-generated touch. These include the finding that attenuation is present when fingers are expected to make contact, even if they fail to do so (4) (recently replicated in Ref. 6), whereas attenuation is not observed when the same movement is made without an expectation of contact, or when equivalent forces are presented passively to both fingers (3). The constant attenuation across different gain conditions implies that the predictive mechanism could update to reflect the new gain relationship, consistent with previous evidence for adaptability of the attenuation process to delays in force transmission (57, 58).

In a previous study based on the “free-energy” model by Friston (59), Brown et al. (15) proposed that sensory attenuation is a necessary component of action. According to this account, voluntary movements are the result of an internal prediction that movement is occurring; the disparity between the proprioceptive feedback expected on this basis and the actual feedback is resolved by the muscular system generating the predicted movement. These authors argued that strong evidence from the other senses that movement is not occurring will prevent movement initiation, and proposed to resolve this by increasing internal uncertainty about these sensory signals around the time of the predicted movement. To account for the attenuation of self-generated forces observed in the force-matching task, the authors further proposed that the perceived magnitude of a force is fixed to the 90% lower confidence bound on the internal estimate of magnitude, i.e., that forces that have higher uncertainty associated with them are perceived as weaker.

The free-energy account has some common elements with our proposed model of attenuation, in particular, the mechanism in both cases is based on predicting the sensory consequences of self-action. However, there are also significant differences. In the present model, the goal of predicting sensory consequences is to downweigh them relative to sensations with an external cause; excess variability is an inevitable but incidental consequence of noise in the predictive process. In the free-energy model, the goal is to increase internal uncertainty about whether a force is being self-generated, because this is necessary to permit a force to be self-generated: attenuation is an incidental by-product of this increase in uncertainty, solely due to their assumption that

more uncertain forces are perceived as weaker. Uncertainty could be increased by injecting noise into sensory inputs, which would predict excess variability in estimating self-generated forces; however, this does not appear to be a necessary component of the free-energy model, as uncertainty could instead be added to the prior or directly to the posterior without increasing variability. Predictions for variability in force estimation were not discussed in the previous paper.

Importantly, Brown et al. (15) modeled attenuation of self-generated forces in the context of force matching, but did not model the sequence of events in a force matching trial [in fact in Brown et al.’s (15) simulations, external forces were generated to match the lower confidence bound on a preceding self-generated force] or consider how variability in perception, memory or sensory predictions would affect reproduction of an external force. Attenuation in the free-energy model is due to a seemingly ad hoc assumption relating perception of sensation to uncertainty, which conflicts with the wider literature on sensory perception, including Bayesian models of perception (60). In contrast, the principle that self-generated sensations are actively cancelled, as in the present model, is supported by neurophysiological studies in multiple model systems (7–10).

More recently, the free-energy model inspired two works where a robotic arm was controlled by a hierarchical recurrent neural network (RNN) designed to follow the free-energy minimization principle (16, 17). The RNN was trained to make movements with feedback consistent with self-generated and externally generated contexts. After training, the precision of the sensory prior was reduced in the self-produced context and increased for externally generated sensations, similarly to Brown et al. (15). However, these studies did not address the perceived intensity of force sensation and did not model the force matching task.

A key assumption of our model of the force-matching task is that the attenuation mechanism is inactive in the indirect condition. This might seem counterintuitive, given that the force in the passive finger is “self-generated,” in the sense that it is controlled by the participant moving the joystick or slider and, at least in principle, predictable based on movement of the active hand. In common with previous accounts (2, 5, 21, 22), we suggest attenuation is absent because the causal relationship between action and sensory consequence in this condition is novel and lacks an ecological mechanism. This interpretation is consistent with previous findings that physically separating the finger applying a force from the finger receiving it reduces the level of attenuation [although does not eliminate it (5)], and that forces that are predictable based on active finger movement but not associated with a contact event do not show attenuation (3, 6, 20). Previous evidence that the predictive mechanism can adapt to changes in the causal relationship (57, 58) suggests that attenuation might develop in the novel context with sufficient exposure, but this may take considerably longer than a typical experimental session. Note that sensory prediction is not necessary for the task of force reproduction using a joystick or slider (or indeed a finger press), as this can be achieved with a simple feedback control loop.

A possible alternative account for the absence of attenuation in the indirect condition is based on the fact that selective attention is divided between two spatial locations (the

passive finger is located under the force lever while the active finger is at the location of the response device), whereas in the direct condition the fingers are approximately collocated. Although focused attention is not generally associated with a reduction in perceived intensity, it might be argued that the forces in active and passive fingers are more difficult to disambiguate when they are collocated and this confusion in some way leads to a lessening of the perceived force. However, this account would be inconsistent with the substantial body of evidence for predictive attenuation in “tapping” tasks, where the task-relevant events all occur in the same physical location, and *test* taps are compared with externally applied *reference* taps of varying amplitude to determine the point of subjective equality (3, 4, 55, 58). For example, in Bays et al. (3), when a self-generated force was produced by an active finger tapping on a passive finger, the perceived force in the passive finger was attenuated, but when equivalent forces were externally applied simultaneously to both fingers, no attenuation was observed, despite the spatial locations and postures of the fingers being identical.

Taken as a whole, the observations above imply that the predictability of a sensory input is necessary but not sufficient for it to be attenuated. Predictive attenuation may have evolved to enhance sensitivity to sensations with an external origin, perhaps because they are more likely to indicate a threat or otherwise require a behavioral response. For this purpose, downweighing inputs based purely on their predictability would be inadequate, as many sensory events with

external causes can be predicted (e.g., the sound of a door shutting can be predicted if you observe someone pushing it). Moreover, attenuating predictable signals could impair our ability to learn further causal relationships involving them. If attenuation requires a determination of self (vs. external) cause in addition to predictability, it may be related to a sense of agency (61). Future studies could examine whether the adaptation of attenuation to changed causal relationships (e.g., delays) is coupled to increases in the perception of self-agency in the interaction.

APPENDIX

Hierarchical Bayesian Modeling

Model specification.

Figure A1 shows a graphical illustration of the hierarchical Bayesian model we fit to empirical force-matching data. In brief, for each participant and force level, we fit two parameters (mean and SD) to account for matching forces in the indirect condition. These participant-level parameters were samples from normal and gamma hyperdistributions, respectively, which were unique to each force level but common to all participants. For each participant, we fit two further parameters to account for matching forces in the direct condition across all force levels: these were the mean and SD of the attenuation factor, the key parameters of interest for our





Figure A1. Graph describing the Bayesian hierarchical model. Unshaded and shaded circular nodes indicate unobserved and observed variables, respectively. Single-bordered nodes denote stochastic variables, while double-bordered nodes denote deterministic variables (as the mean direct matching force and standard deviation are fully determined by the stochastic means and standard deviations of the indirect force and attenuation factor). Plates indicate repetitive structures within the model. The additive version of the model is shown: the divisive model has an additional deterministic node corresponding to the skewness of direct matching forces.






Figure A3. Matching forces and trial-to-trial variability in matching force in *Study 3*. **A:** mean matching force as a function of target force, for direct (cyan) and indirect (red) conditions. Data are shown as symbols and linear fits as lines. **B:** matching force in the direct condition plotted against matching force in the indirect condition, with linear fit. Symbols indicate mean matching forces for each target force level. **C:** standard deviation of matching force as a function of target force, for direct (cyan) and indirect (red) conditions, with linear fits. **D:** standard deviation of matching force as a function of mean matching force, with linear fits. Each symbol corresponds to a single target force level. **A–D:** Healthy subjects. **E–H:** same as **A–D** but for patient subjects. Error bars indicate ± 1 SE, and shaded areas indicate 95% CI. CI, confidence interval; SE, standard error.

analysis. These participant-level parameters were again sampled from normal and gamma hyper-distributions common to all participants. Each parameter associated with a hyperdistribution (nodes outside the participant plate) was constrained by a hyperprior distribution. Later, we provide a full description of the model.

Indirect matching forces.

We simultaneously fit the data from all conditions within each experiment. Our aim for the indirect reproduction conditions (*I*) was to measure the matching forces rather

than to predict them, so we modeled the matching forces on these trials as samples drawn from a normal distribution, with a unique pair of mean and SD parameters for each target force level t and participant:

$$F_t^i \sim \text{Normal}(\bar{F}_t^i, \sigma_{F_t^i}). \quad (\text{A1})$$

Note that for readability, we suppress the participant index on all variables. For each participant, we set individual-level priors for the mean and standard deviation within each force level:

Figure A2. Matching forces and trial-to-trial variability in matching force in *Study 2*. **A:** mean matching force as a function of target force, for direct (cyan) and indirect (red) conditions. Data are shown as symbols and linear fits as lines. **B:** matching force in the direct condition plotted against matching force in the indirect condition, with linear fit. Symbols indicate mean matching forces for each target force level. **C:** standard deviation of matching force as a function of target force, for direct (cyan) and indirect (red) conditions, with linear fits. **D:** standard deviation of matching force as a function of mean matching force, with linear fits. Each symbol corresponds to a single target force level. **A–D:** young subjects. **E–H:** same as **A–D** but for middle-aged subjects. **I–L:** Same as **A–D** but for older subjects. Error bars indicate ± 1 SE, and shaded areas indicate 95% CI. CI, confidence interval; SE, standard error.




Figure A4. Posterior predictive densities for Study 2. *A*: young subject. *B*: middle-aged subjects. *C*: older subjects. Gray histograms show matching forces pooled across all participants. Colored distributions show simulated matching forces based on posterior predictive density. The circles with error bars show the median and 90% HDI for data (black) and simulated trials for the same condition (cyan: direct, red: indirect). HDI, highest density interval.




Figure A5. Posterior predictive densities for *Study 3*. *A*: healthy subject. *B*: patients subjects. Gray histograms show matching forces pooled across all participants. Colored distributions show simulated matching forces based on posterior predictive density. The circles with error bars show the median and 90% HDI for data (black) and simulated trials for the same condition (cyan: direct, red: indirect). HDI, highest density interval.

$$\begin{aligned}\bar{F}_I^t &\sim \text{Normal}(\mu_{F_I^t}, 1/\sqrt{\tau_{F_I^t}}) \\ \sigma_{F_I^t} &\sim \text{Gamma}(s_{\sigma_{F_I^t}}, r_{\sigma_{F_I^t}}).\end{aligned}\quad (\text{A2})$$

These individual-level parameters were, in turn, dependent on population-level hyperparameters, with hyperpriors:

$$\begin{aligned}\mu_{\bar{F}_I^t} &\sim \text{Normal}(F_T^t, 1/\sqrt{0.3}) \\ \tau_{\bar{F}_I^t} &\sim \text{Gamma}(0.001, 0.001).\end{aligned}\quad (\text{A3})$$

The shape (s) and rate (r) parameters of the gamma prior distribution on $\sigma_{F_I^t}$ had priors defined in terms of the central tendency and width of the distribution:

$$\begin{aligned}\mu_{\sigma_{F_I^t}} &\sim \text{Gamma}(0.01, 0.01) \\ \sigma_{\sigma_{F_I^t}} &\sim \text{Gamma}(0.001, 0.001).\end{aligned}\quad (\text{A4})$$

We then used a simple reparametrization to convert the mean and standard deviation into the shape and rate parameters (35):




Figure A6. Subtractive model posterior predictive densities for *Study 1*. Gray histograms show matching forces pooled across all participants. Colored distributions show simulated matching forces based on posterior predictive density. The circles with error bars show the median and 90% HDI for data (black) and simulated trials for the same condition (cyan: direct, red: indirect). HDI, highest density interval.

$$s_{\sigma_{F_l^t}} = \frac{\mu_{\sigma_{F_l^t}}^2}{\sigma_{\sigma_{F_l^t}}^2} \quad (A5)$$

$$r_{\sigma_{F_l^t}} = \frac{\mu_{\sigma_{F_l^t}}}{\sigma_{\sigma_{F_l^t}}^2}$$

Parameterizing a gamma distribution in terms of the mean and standard deviation, as done here, results in a much more intuitive interpretation of the prior and resulting posterior distributions than when using shape and rate parameters. However, gamma distributions are typically positively skewed, which can affect their means due to their asymmetric tails. In this case, a potentially better measure of central tendency is the median. As the median of a gamma distribution has no simple closed form, we opted to define priors over the mean instead. This approach offers a pragmatic compromise, particularly when using weakly informative priors. After fitting the model and obtaining the posterior distributions, we estimated the median numerically by first generating 10^7 random samples using the accepted parameters on each MCMC iteration, and then calculating the median of those samples to estimate the posterior distribution of the standard deviation of the attenuation. For completeness, we additionally computed the mode using the fitted shape and rate parameters of the gamma distribution (35). Comparisons conducted on the mode estimates were largely consistent with those performed on the median estimates.

Direct matching forces.

Subtractive model. In the subtractive model, direct matching forces are considered samples from a normal distribution with mean and SD determined by combining the

indirect force matching parameters corresponding to the same target force ($\bar{F}_l^t, \sigma_{F_l^t}$) with the attenuation factor parameters ($\bar{K}, \sigma_K$), according to *Eqs. 1 and 2*:

$$F_D^t \sim \text{Normal}(\bar{F}_D^t, \sigma_{F_D^t}) \quad (A6)$$

We set individual-level priors for the attenuation parameters:

$$\bar{K} \sim \text{Normal}(\mu_{\bar{K}}, 1/\sqrt{\tau_{\bar{K}}}) \quad (A7)$$

$$\sigma_K \sim \text{Gamma}(s_{\sigma_K}, r_{\sigma_K})$$

which are again dependent on population-level hyperparameters, with the following hyperpriors:

$$\begin{aligned} \mu_{\bar{K}} &\sim \text{Normal}(1.1, 1/\sqrt{0.1}) \\ \tau_{\bar{K}} &\sim \text{Gamma}(0.001, 0.001) \\ \mu_{\sigma_K} &\sim \text{Gamma}(0.01, 0.01) \\ \sigma_{\sigma_K} &\sim \text{Gamma}(0.01, 0.01) \end{aligned} \quad (A8)$$

$$s_{\sigma_K} = \frac{\mu_{\sigma_K}^2}{\sigma_{\sigma_K}^2}$$

$$r_{\sigma_K} = \frac{\mu_{\sigma_K}}{\sigma_{\sigma_K}^2}$$

Divisive model. In the divisive model, attenuated output force on an individual trial is considered a sample from a skew-normal distribution (see *Skew-Normal Approximation to the Product of Normals*) with mean, SD, and skewness defined as in *Eqs. 3–5*:

$$F_D^t \sim \text{SkewNormal}(\bar{F}_D^t, \sigma_{F_D^t}, \text{Skew}[F_D^t]). \quad (A9)$$

Individual- and population-level priors for the attenuation parameters and hyperparameters were set identically to the subtractive model (see *Subtractive model*), except for the population mean hyperprior for $\bar{K}$ which was set to:




Figure A7. Subtractive model posterior predictive check. **A:** Study 1. **B:** Study 2 (from top to bottom: young, middle, and older subjects). **C:** Study 3 (from top to bottom: healthy and patient subjects). The solid and dashed lines show fits of the subtractive and divisive model, respectively. The circles with error bars show the mean and 95% CI. CI, confidence interval.

Table A1. Summary of subtractive model

No.	Sample	Attenuation Factor		SD of Attenuation Factor		Bayesian <i>P</i> Value
		Median	90% HDI	Median	90% HDI	
1	Study 1	1.36	[1.24, 1.49]	0.726	[0.657, 0.803]	$0.34 \leq \text{all } P \text{ values} \leq 0.61$
2	Study 2 young	0.76	[0.62, 0.90]	0.48	[0.403, 0.569]	$0.43 \leq \text{all } P \text{ values} \leq 0.54$
3	Study 2 middle	1.11	[1.00, 1.23]	0.526	[0.467, 0.586]	$0.40 \leq \text{all } P \text{ values} \leq 0.57$
4	Study 2 older	1.41	[1.26, 1.55]	0.646	[0.567, 0.726]	$0.39 \leq \text{all } P \text{ values} \leq 0.60$
5	Study 3 healthy	1.15	[0.77, 1.53]	0.727	[0.520, 0.955]	$0.39 \leq \text{all } P \text{ values} \leq 0.55$
6	Study 3 patients	0.60	[0.36, 0.84]	0.393	[0.261, 0.528]	$0.46 \leq \text{all } P \text{ values} \leq 0.52$

$$\mu_{\bar{K}} \sim \text{Normal}(1.75, 1/\sqrt{0.1}). \quad (\text{A10}) \quad \text{for } -1 < \delta < 1.$$

To illustrate the fitted parameters of the divisive model, we plot the mean population-level attenuation factor (μ_K) in Fig. 5A. Following the described reparametrization of the gamma distribution, we display the median population-level standard deviation of the attenuation factor (m_{σ_K}) in Fig. 5B.

RESULTS

Linear Regression

Figs. A2 and A3 show data and linear fits obtained for Studies 2 and 3, respectively, in the same format as shown for Study 1 in Figs. 2 and 3. Corresponding linear regression coefficients are reported in Fig. 2C and Fig. 3C.

Divisive model.

Figs. A4 and A5 show the results of fitting the divisive model to data from Studies 2 and 3, respectively, in the form of posterior predictive distributions. We performed the model fitting separately in the three age subgroups of Study 2 and in patients and healthy controls for Study 3. Analogously to Study 1 (Fig. 4), these model predictions accurately reproduced the empirical distributions (gray histograms) of matching forces in direct (cyan) and indirect (red) conditions at each force level.

Subtractive model.

We also fit the subtractive model to the data from Studies 1–3. Figure A6 shows the results for Study 1 in the form of posterior predictive distributions. Compared with the divisive model (Fig. 4), the subtractive model reproduced the distributions of matching forces across the five force levels less accurately. This is further illustrated in Fig. A7, which compares empirical and posterior predicted mean matching forces for the subtractive and divisive models. A summary of estimated values of the attenuation factor and its standard deviation for the subtractive model is reported in Table A1.

Skew-Normal Approximation to the Product of Normals

$$\text{If } X \sim N(\mu_x, \sigma_x) \text{ and } Y \sim N(\mu_y, \sigma_y), \text{ for } Z = XY, \quad (\text{A11})$$

$$E[Z] = \mu_x \mu_y, \quad (\text{A12})$$

$$\text{Var}[Z] = \mu_y^2 \sigma_x^2 + \mu_x^2 \sigma_y^2 + \sigma_x^2 \sigma_y^2 = (1 + \delta_x^2 + \delta_y^2) \sigma_x^2 \sigma_y^2, \quad (\text{A13})$$

$$\text{Skewness}[Z] = \frac{6\mu_x \mu_y \sigma_x^2 \sigma_y^2}{\text{Var}[Z]^{3/2}} = \frac{6\delta_x \delta_y}{(1 + \delta_x^2 + \delta_y^2)^{3/2}}, \quad (\text{A14})$$

$$\text{where } \delta_x = \frac{\mu_x}{\sigma_x} \text{ and } \delta_y = \frac{\mu_y}{\sigma_y}. \quad (\text{A15})$$

This distribution can be approximated by a skew-normal distribution with the same moments. The skew-normal has pdf,

$$p(x) = \frac{2}{\omega} \phi\left(\frac{x - \xi}{\omega}\right) \Phi\left(\alpha \frac{x - \xi}{\omega}\right). \quad (\text{A16})$$

Defining

$$\delta = \frac{\alpha}{\sqrt{1 + \alpha^2}} \Rightarrow \alpha = \frac{\delta}{\sqrt{1 - \delta^2}}, \quad (\text{A17})$$

$$E[x] = \xi + \omega \delta \sqrt{2/\pi}, \quad (\text{A18})$$

$$\text{Var}[x] = \omega^2 (1 - (2/\pi) \delta^2), \quad (\text{A19})$$

$$\text{Skewness}[x] = \gamma_1 = \frac{4 - \pi}{2} (2/\pi)^{3/2} \frac{\delta^3}{(1 - (2/\pi) \delta^2)^{3/2}}, \quad (\text{A20})$$

$$\Rightarrow \delta = \text{Sign}(\gamma_1) \sqrt{\frac{\pi}{2} \frac{|\gamma_1|^{2/3}}{|\gamma_1|^{2/3} + ((4 - \pi)/2)^{2/3}}}, \quad (\text{A21})$$

$$\omega = \sqrt{\frac{\text{Var}[x]}{1 - (2/\pi) \delta^2}}, \quad (\text{A22})$$

$$\xi = E[x] - \omega \delta \sqrt{2/\pi}. \quad (\text{A23})$$

DATA AVAILABILITY

Data and code related to this study are available at <https://doi.org/10.17863/CAM.113018>.

ACKNOWLEDGMENTS

N.V. thanks Professor Nicola Smania (University of Verona, Italy) for helpful and insightful discussions. The authors thank all authors of source studies and the Cambridge Centre for Ageing and Neuroscience (CamCAN; <https://cam-can.mrc-cbu.cam.ac.uk>) for sharing data.

A preprint is available at <https://doi.org/10.1101/2024.06.20.599826>.

Present addresses: Nicola Valè, Dept. of Engineering for Innovation Medicine, Univ. of Verona, Verona, 37134, Italy; Ivan Tomić, Dept. of Psychology, Faculty of Humanities and Social Sciences, Univ. of Zagreb, Zagreb, Croatia.

GRANTS

This research was supported by the European Research Council (ERC) starting grant TICKLISHUMAN (Grant No. 101039152 to K.K.) and the Wellcome Trust (Grant No. 106926 to P.M.B.).

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

N.V., D.M.W., and P.M.B. conceived and designed research; K.K. performed experiments; N.V., I.T., and Z.G. analyzed data; N.V., I.T., D.M.W., and P.M.B. interpreted results of experiments; I.T. and P.M.B. prepared figures; N.V. and I.T. drafted manuscript; N.V., I.T., Z.G., D.M.W., K.K., and P.M.B. edited and revised manuscript; N.V., I.T., Z.G., D.M.W., K.K., and P.M.B. approved final version of manuscript.

REFERENCES

1. Blakemore S-J, Wolpert D, Frith C. Why can't you tickle yourself? *Neuroreport* 11: R11–R16, 2000. doi:10.1097/00001756-200008030-00002.

2. Shergill SS, Bays PM, Frith CD, Wolpert DM. Two eyes for an eye: the neuroscience of force escalation. *Science* 301: 187, 2003. doi:10.1126/science.1085327.
3. Bays PM, Wolpert DM, Flanagan JR. Perception of the consequences of self-action is temporally tuned and event driven. *Curr Biol* 15: 1125–1128, 2005. doi:10.1016/j.cub.2005.05.023.
4. Bays PM, Flanagan JR, Wolpert DM. Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biol* 4: e28, 2006. doi:10.1371/journal.pbio.0040028.
5. Bays PM, Wolpert DM. Predictive attenuation in the perception of touch. In: *Sensorimotor Foundations of Higher Cognition*. Oxford University Press, 2008, p. 339–358. doi:10.1093/acprof:oso/9780199231447.003.0016.
6. Job X, Kiltner K. Action does not enhance but attenuates predicted touch. *eLife* 12: e90912, 2023. doi:10.7554/eLife.90912.
7. Bell CC. Memory-based expectations in electrosensory systems. *Curr Opin Neurobiol* 11: 481–487, 2001. doi:10.1016/S0959-4388(00)00238-5.
8. Sperry R. Neural basis of the spontaneous optokinetic response produced by visual inversion. *J Comp Physiol Psychol* 43: 482–489, 1950. doi:10.1037/h0055479.
9. von Holst E, Mittelstaedt H. The principle of reafference: interactions between the central nervous system and the peripheral organs. *Perceptual Processing: Stimulus Equivalence and Pattern Recognition*, Appleton-Century, 1971, p. 41–72.
10. Cullen KE, Brooks JX, Sadeghi SG. How actions alter sensory processing. *Ann NY Acad Sci* 1164: 29–36, 2009. doi:10.1111/j.1749-6632.2009.03866.x.
11. Wolpert DM, Ghahramani Z. Computational principles of movement neuroscience. *Nat Neurosci* 3: 1212–1217, 2000. doi:10.1038/81497.
12. Blakemore S-J, Goodbody SJ, Wolpert DM. Predicting the consequences of our own actions: the role of sensorimotor context estimation. *J Neurosci* 18: 7511–7518, 1998. doi:10.1523/JNEUROSCI.18-18-07511.1998.
13. Kiltner K, Ehrsson HH. Predictive attenuation of touch and tactile gating are distinct perceptual phenomena. *Isience* 25: 104077, 2022. doi:10.1016/j.isci.2022.104077.
14. Walsh LD, Taylor JL, Gandevia SC. Overestimation of force during matching of externally generated forces. *J Physiol* 589: 547–557, 2011. doi:10.1113/jphysiol.2010.198689.
15. Brown H, Adams RA, Parees I, Edwards M, Friston K. Active inference, sensory attenuation and illusions. *Cogn Process* 14: 411–427, 2013. doi:10.1007/s10339-013-0571-3.
16. Idei H, Ohata W, Yamashita Y, Ogata T, Tani J. Emergence of sensory attenuation based upon the free-energy principle. *Sci Rep* 12: 14542–2322, 2022. doi:10.1038/s41598-022-18207-7.
17. Suzuki R, Idei H, Yamashita Y, Ogata T. Hierarchical variational recurrent neural network modeling of sensory attenuation with temporal delay in action-outcome. *2023 IEEE International Conference on Development and Learning (ICDL)*. IEEE, 2023, vol. 11, p. 244–249. doi:10.1109/icdl55364.2023.10364405.
18. Press C, Kok P, Yon D. The perceptual prediction paradox. *Trends Cogn Sci* 24: 13–24, 2020. doi:10.1016/j.tics.2019.11.003.
19. Press C, Thomas ER, Yon D. Cancelling cancellation? sensorimotor control, agency, and prediction. *Neurosci Biobehav Rev* 145: 105012–107634, 2023. doi:10.1016/j.neubiorev.2022.105012.
20. Thomas ER, Yon D, de Lange FP, Press C. Action enhances predicted touch. *Psychol Sci* 33: 48–59, 2022. doi:10.1177/09567976211017505.
21. Kiltner K, Ehrsson HH. Functional connectivity between the cerebellum and somatosensory areas implements the attenuation of self-generated touch. *J Neurosci* 40: 894–906, 2020. doi:10.1523/jneurosci.1732-19.2019.
22. Shergill SS, Samson G, Bays PM, Frith CD, Wolpert DM. Evidence for sensory prediction deficits in schizophrenia. *Am J Psychiatry* 162: 2384–2386, 2005. doi:10.1176/appi.ajp.162.12.2384.
23. Wolpe N, Ingram JN, Tsvetanov KA, Geerligs L, Kievit RA, Henson RN, Wolpert DM, Rowe JB, Cam-CAN. Ageing increases reliance on sensorimotor prediction through structural and functional differences in frontostriatal circuits. *Nat Commun* 7: 13034, 2016. doi:10.1038/ncomms13034.
24. Shafto MA, Tyler LK, Dixon M, Taylor JR, Rowe JB, Cusack R, Calder AJ, Marslen-Wilson WD, Duncan J, Dalgleish T, Henson RN, Bayne C, Matthews FE, Cam-CAN. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC Neurol* 14: 204, 2014. doi:10.1186/s12883-014-0204-1.
25. Taylor JR, Williams N, Cusack R, Auer T, Shafto MA, Dixon M, Tyler LK, Henson RN. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage* 144: 262–269, 2017. doi:10.1016/j.neuroimage.2015.09.018.
26. Liang F, Paulo R, Molina G, Clyde MA, Berger JO. Mixtures of g priors for Bayesian variable selection. *J Am Stat Assoc* 103: 410–423, 2008. doi:10.1198/016214507000001337.
27. Wagenmakers E-J, Love J, Marsman M, Jamil T, Ly A, Verhagen J, Selker R, Gronau QF, Drophmann D, Boutin B, Meerhoff F, Knight P, Raj A, van Kesteren E-J, van Doorn J, Šmíra M, Epskamp S, Etz A, Matzke D, de Jong T, van den Bergh D, Sarafoglou A, Steingrover H, Derks K, Rouder JN, Morey RD. Bayesian inference for psychology. Part II: example applications with JASP. *Psychon Bull Rev* 25: 58–76, 2018. doi:10.3758/s13423-017-1323-7.
28. Oliveira A, Seijas-Macias A. An approach to distribution of the product of two normal variables. *Discuss Math Probab Stat* 32: 87, 2012. doi:10.7151/dmps.1146.
29. Lee MD. How cognitive modeling can benefit from hierarchical Bayesian models. *J Math Psychol* 55: 1–7, 2011. doi:10.1016/j.jmp.2010.08.013.
30. Lee MD, Wagenmakers E-J. *Bayesian Cognitive Modeling: A Practical Course*. Cambridge University Press, 2013, p. 264.
31. Baribault B, Collins AGE. Troubleshooting Bayesian cognitive models. *Psychol Methods* 30: 128–154, 2025. doi:10.1037/met0000554.
32. Plummer M. JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. In: *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. Vienna, Austria, 2003, vol. 124, p. 1–10.
33. Steyvers M. Matjags 1.3: A MATLAB Interface For Jags. <https://github.com/msteyvers/matjags> [2025 Jul 2].
34. Gelman A, Rubin DB. Inference from iterative simulation using multiple sequences. *Statist Sci* 7: 457–472, 1992. doi:10.1214/ss/1177011136.
35. Kruschke JK. *Doing Bayesian data analysis: a tutorial with R, JAGS, and Stan* (2nd ed.). Academic Press, 2015, p. 759.
36. Kruschke JK. Posterior predictive checks can and should be Bayesian: comment on Gelman and Shalizi, 'Philosophy and the practice of Bayesian statistics'. *Br J Math Stat Psychol* 66: 45–56, 2013. doi:10.1111/j.2044-8317.2012.02063.x.
37. Hollingworth HL. The central tendency of judgment. *J Phil Psych Sci Meth* 7: 461–469, 1910. doi:10.2307/2012819.
38. Ashourian P, Loewenstein Y. Bayesian inference underlies the contraction bias in delayed comparison tasks. *PLoS One* 6: e19551, 2011. doi:10.1371/journal.pone.0019551.
39. Salthouse TA, Babcock RL. Decomposing adult age differences in working memory. *Dev Psychol* 27: 763–776, 1991. doi:10.1037//0012-1649.27.5.763.
40. McNab F, Zeidman P, Rutledge RB, Smittenaar P, Brown HR, Adams RA, Dolan RJ. Age-related changes in working memory and the ability to ignore distraction. *Proc Natl Acad Sci USA* 112: 6515–6518, 2015. doi:10.1073/pnas.1504162112.
41. Slifkin AB, Newell KM. Noise, information transmission, and force variability. *J Exp Psychol Hum Percept Perform* 25: 837–851, 1999. doi:10.1037//0096-1523.25.3.837.
42. Christou EA, Grossman M, Carlton LG. Modeling variability of force during isometric contractions of the quadriceps femoris. *J Mot Behav* 34: 67–81, 2002. doi:10.1080/00222890209601932.
43. Chapman CE, Bushnell MC, Miron D, Duncan GH, Lund JP. Sensory perception during movement in man. *Exp Brain Res* 68: 516–524, 1987. doi:10.1007/BF00249795.
44. Cohen LG, Starr A. Localization, timing and specificity of gating of somatosensory evoked potentials during active movement in man. *Brain* 110: 451–467, 1987. doi:10.1093/brain/110.2.451.
45. Colino FL, Buckingham G, Cheng DT, van Donkelaar P, Binsted G. Tactile gating in a reaching and grasping task. *Physiol Rep* 2: e00267, 2014. doi:10.1002/phy2.267.
46. Papakostopoulos D, Cooper R, Crow HJ. Inhibition of cortical evoked potentials and sensation by self-initiated movement in man. *Nature* 258: 321–324, 1975. doi:10.1038/258321a0.

47. **Rushton DN, Rothwell JC, Craggs MD.** Gating of somatosensory evoked potentials during different kinds of movement in man. *Brain* 104: 465–491, 1981. doi:[10.1093/brain/104.3.465](https://doi.org/10.1093/brain/104.3.465).
48. **Pertovaara A, Kemppainen P, Leppänen H.** Lowered cutaneous sensitivity to nonpainful electrical stimulation during isometric exercise in humans. *Exp Brain Res* 89: 447–452, 1992. doi:[10.1007/BF00228261](https://doi.org/10.1007/BF00228261).
49. **Shergill SS, White TP, Joyce DW, Bays PM, Wolpert DM, Frith CD.** Functional magnetic resonance imaging of impaired sensory prediction in schizophrenia. *JAMA Psychiatry* 71: 28–35, 2014. doi:[10.1001/jamapsychiatry.2013.2974](https://doi.org/10.1001/jamapsychiatry.2013.2974).
50. **Shergill SS, White TP, Joyce DW, Bays PM, Wolpert DM, Frith CD.** Modulation of somatosensory processing by action. *NeuroImage* 70: 356–362, 2013. doi:[10.1016/j.neuroimage.2012.12.043](https://doi.org/10.1016/j.neuroimage.2012.12.043).
51. **Abram SV, Hua JPY, Ford JM.** Consider the pons: bridging the gap on sensory prediction abnormalities in schizophrenia. *Trends Neurosci* 45: 798–808, 2022. doi:[10.1016/j.tins.2022.08.008](https://doi.org/10.1016/j.tins.2022.08.008).
52. **Moberget T, Ivry RB.** Prediction, psychosis, and the cerebellum. *Biol Psychiatry Cogn Neurosci Neuroimaging* 4: 820–831, 2019. doi:[10.1016/j.bpsc.2019.06.001](https://doi.org/10.1016/j.bpsc.2019.06.001).
53. **Hunter SK, Pereira HM, Keenan KG.** The aging neuromuscular system and motor performance. *J Appl Physiol* (1985) 121: 982–995, 2016. doi:[10.1152/japplphysiol.00475.2016](https://doi.org/10.1152/japplphysiol.00475.2016).
54. **Schröder J, Essig M, Baudendistel K, Jahn T, Gerdson I, Stockert A, Schad LR, Knopp MV.** Motor dysfunction and sensorimotor cortex activation changes in schizophrenia: a study with functional magnetic resonance imaging. *NeuroImage* 9: 81–87, 1999. doi:[10.1006/nimg.1998.0387](https://doi.org/10.1006/nimg.1998.0387).
55. **Kilteni K, Engeler P, Ehrsson HH.** Efference copy is necessary for the attenuation of self-generated touch. *iScience* 23: 100843, 2020. doi:[10.1016/j.isci.2020.100843](https://doi.org/10.1016/j.isci.2020.100843).
56. **Palmer CE, Davare M, Kilner JM.** Physiological and perceptual sensory attenuation have different underlying neurophysiological correlates. *J Neurosci* 36: 10803–10812, 2016. doi:[10.1523/JNEUROSCI.1694-16.2016](https://doi.org/10.1523/JNEUROSCI.1694-16.2016).
57. **Kilteni K, Ehrsson HH.** Dynamic changes in somatosensory and cerebellar activity mediate temporal recalibration of self-touch. *Commun Biol* 7: 522, 2024. doi:[10.1038/s42003-024-06188-4](https://doi.org/10.1038/s42003-024-06188-4).
58. **Kilteni K, Houborg C, Ehrsson HH.** Rapid learning and unlearning of predicted sensory delays in self-generated touch. *eLife* 8: e42888, 2019. doi:[10.7554/eLife.42888](https://doi.org/10.7554/eLife.42888).
59. **Friston K.** The free-energy principle: a unified brain theory? *Nat Rev Neurosci* 11: 127–138, 2010. doi:[10.1038/nrn2787](https://doi.org/10.1038/nrn2787).
60. **Ma WJ, Kording K, Goldreich D.** *Bayesian Models of Perception and Action: An Introduction*. MIT Press, 2023.
61. **Haggard P, Chambon V.** Sense of agency. *Curr Biol* 22: R390–R392, 2012. doi:[10.1016/j.cub.2012.02.040](https://doi.org/10.1016/j.cub.2012.02.040).