Supplementary Information for

DOPAMINERGIC MEDICATION INCREASES RELIANCE ON CURRENT INFORMATION IN PARKINSON'S DISEASE PATIENTS

Iris Vilares[1,2,3,*] and Konrad Kording[2,3]

[1]Wellcome Trust Centre for Neuroimaging, University College London, London, UK.

[2]Rehabilitation Institute of Chicago, Chicago, IL, 60611, USA.

[3]Northwestern University, Feinberg School of Medicine, Chicago, IL, 60611, USA.

* *Correspondence should be sent to irisvilares@gmail.com*

SUPPLEMENTARY METHODS

Participants

23 Parkinson disease (PD) patients and 16 healthy controls were recruited for the study. From these, eight patients were discarded: six due to failure to follow the "on-off" medication schedule required for the experiment, one because the patient was still not taking meds and one that could not understand the task even after several explanations. One healthy control was discarded because she had already participated in a similar study in the lab before. Thus, in total, 15 PD patients (9 women) and 15 non-PD controls (8 women) participated in the experiment. This sample size was chosen to be in line with typical patient population papers sample sizes, at the time the experiment was designed, and because previous experiments in the lab using a similar paradigm indicated that it was possible to obtain significant effects using this sample size[1,2]. Participants were paid $20 for participating in each session of the experiment, so they got $40 for both sessions. Written informed consent was obtained for all participants. All protocols were approved by the Northwestern University IRB.

PD patients

All the PD patients were recruited from the Rehabilitation Institute of Chicago and the Northwestern Memorial Hospital. All PD patients had idiopathic Parkinson's disease as diagnosed by a neurologist. The 15 patients included in the study (9 women) had between 43 to 81 years old (62.8 ± 9.6 years old; mean ± s.d.). Only patients with bilateral involvement were recruited. Severity of disease was equivalent to Hoehn and Yahr stages ranging between II and IV. The duration of Parkinson's disease varied from 1.5 months to 18 years from the time of initial diagnosis (6.2 ± 4.8 years; mean ± s.d.). Eleven of the fifteen patients included in the study were taking daily L-dopa preparations (see Supplementary Table 1). Of those, eight received additional pharmacotherapy with dopamine receptor agonists, rasagiline or amantadine. Four patients were taking dopamine receptor agonists and/or rasagiline and/or amantadine. L-dopa equivalent units (LEU) of patients' regular daily dopamine replacement therapies were calculated as described elsewhere[3], and varied between 100 mg to 1733 mg (675 ± 436 mg; mean ± s.d.). Note that although the different dopaminergic medication may cause different effects on patients' behaviour, we were interested in studying the general effects of dopaminergic medication on decision-making under uncertainty, and not specific medication-related ones.

The half-life of L-dopa is about 90 minutes[4], so we would expect all the L-dopa medication effects to have washed out after overnight withdrawal medication. We would also expect that from rasagiline (a MAOB-inhibitor), whose half-life is 3 hours[5], although this medication may cause irreversible effects as well. The half-life of dopamine receptor agonists and of Amantadine may be longer, between 6 to 12 hours[6-8], hence after a 12h withdrawal only about 50-75% of the medication would have washed out (the specific values would change from patient to patient). Longer "off" times were quite often not possible, as patients would start being unable to move. We chose overnight/12h withdrawal to have the maximum difference between on and off medication states, while at the same time still having the patients willing and able to come perform the experiment.

Initially, when acquiring pilot data, the experimenter was blind to the medication state of the PD patient. However, due to the effects of dopaminergic medication (and withdrawal) on the patients' motoric behaviour, it was quickly realized that it was not possible to have a fully blind design. So, instead, the instructions were scripted (to minimize any variability in explanation). All the participants' behaviour was obtained and recorded in Matlab, hence the outcome assessment for each participant was done "blindly" (by a Matlab computer program).

Healthy controls

Age and gender-matched healthy volunteers, with no current known neurological problems, were recruited for the study, usually from amongst the patients' spouses or partners. The 15 healthy controls included in the study (8 women) had between 43 to 84 years old (63.4 ± 11.6 years old; mean ± s.d.). No significant difference in age existed between PD patients and Controls (*p-val* = 0.88, unpaired t-test; [-8.5,7.3] 95% confidence interval (CI) for mean difference between ages).

Inclusion and Exclusion Criteria

The inclusion and exclusion criteria were pre-established before any data collection, and are:

Inclusion criteria:

1. Adult, 18 years or greater
2. The participant is able to provide informed consent.
3. For the Parkinson's Disease population, definitive Idiopathic Parkinson's Disease as diagnosed by a Neurologist.
4. For the age-matched control population, no Parkinson's disease
5. For the Parkinson's Disease population, Hoehn and Yahr Stage 2 to 4
6. For the Parkinson's Disease population, on stable doses of Parkinson's medications for at least 2 weeks prior to study onset

Exclusion criteria:

1. Mini Mental Status Exam (MMSE) < 24
2. Anticipated change in Parkinson's medications in the duration of study
3. Symptomatic coronary artery disease
4. Other neurologic diagnosis, including Multiple Sclerosis
5. Other vestibular disease
6. Untreated severe depression (depression of greater than or equal 20 on the Geriatric Depression Scale)
7. Acute illness
8. History of, or current, alcohol abuse
9. Significant visual impairment that would inhibit ability to participate in study, with distance vision >20/40
10. Drug induced or inherited Parkinson's Disease
11. Significant camptocormia
12. Any medical condition which the investigator determines would compromise the safety for the subject.

Experiment (Coin-catching task)

General procedure

PD patients completed two test sessions, once about 1 hour after they had taken their regular dopaminergic medication (*on-medication*) and once after overnight withdrawal from dopaminergic medication (*off-medication*). Session order was counterbalanced across patients (randomly), so that in total eight patients started *on-medication*, and seven started *off-medication*. Control participants also completed two test sessions, and the results reported for controls were the average of the results obtained in the two sessions. For both PD patients and controls, the two sessions were generally done in two consecutive days. We chose overnight/12h withdrawal to have the maximum difference between on and off medication states, while at the same time still having the patients willing and able to come perform the experiment.

Coin-catching task

Participants performed a decision-making task in which they had to guess the position of a hidden coin on a screen[4] (see Fig. 1a). They were told the cover story of a coin being tossed into a pond and informed that their task was to guess where the coin had fallen. They could not see the coin, but they could see 5 blue dots that were the "splashes" produced by the coin falling in. They were told that the person who threw the coin aimed, albeit imperfectly, at the center of the screen (mean of prior). They were also told that, between blocks, the thrower changed, and the new one might be better or worse at throwing (what indirectly informed them that the variance of the prior changed). To estimate the coin position, participants could use (although they were never explicitly told so) both the coin position's likelihood, obtained from the "splashes", and its prior (the distribution of previous coin locations). There was no temporal deadline, so participants had all the time they needed to submit their response. The instructions were scripted and given to the participant to read, to minimize variability in explanation. If something was not understood, the researcher would repeat the main points until the participant acknowledged understanding.

Stimuli

The position of the coin was drawn from a Gaussian distribution which was centered on the middle of the screen, and with a standard deviation (s.d.) that was either low (σ_p = 2.5% of screen width) or high (σ_P = 8.5% of screen width). This distribution was the *prior* of the experiment. Participants were given the mean of the prior ("the coin throw is aimed at the screen center") but not its variance, which they could only estimate from the distribution of previous coin throws. The standard deviation of the prior was kept constant within blocks, but changed across blocks. On every trial, a cluster of five dots was shown on the screen. The x-position of each of these dots was drawn independently from a second Gaussian distribution in which the mean was the coin's horizontal location on that trial and the standard deviation was either low (σ_l = 6%) or high (σ_L = 15%). The distribution of these five dots

defined the *likelihood*. The y-position of the dots was drawn from a Gaussian distribution with the same low or high s.d. than in the x-position (although the specific sample s.d. may have varied), but in which the mean was always 0.5. This was done in order to separate better the dots and give a more realistic appearance of a splash. The s.d. of the likelihood was varied pseudo-randomly from trial to trial but counterbalanced across trials. We made the s.d. of the likelihood vary pseudorandom from trial to trial so that participants could not predict *a priori* the overall uncertainty that the trial would have. In total there were thus four conditions (see Fig. 1b): low prior uncertainty and low likelihood uncertainty (**pl**); low prior uncertainty and high likelihood uncertainty (**pL**); high prior uncertainty and low likelihood uncertainty (**Pl**) and high prior uncertainty and high likelihood uncertainty (**PL**). The s.d. values used were the same ones reported in [1].

The colours and appearance of the stimulus were also identical to the ones reported in [1] (see Fig. 1a). Both the target dot ("coin") and the stimulus dots had the same diameter (10 unit points). The screen units were normalized between 0 (the left edge) and 1 (the right edge). Stimulus presentation was performed using Matlab R2012a (MathWorks, Natick, MA).

Task procedure

Experimental blocks

At the onset of each trial, five blue likelihood dots were shown on the screen, where they remained until the end of the trial. Participants had to move a blue vertical bar (the "net") to estimate the coin position. Contrary to the task in [1], the initial position of the net was drawn randomly in every trial. Furthermore, in order for the participants to move the net they had to press either the "left" or the "right" keyboard key (which would move the net, respectively, to the left or the right). Once they made their decision, they would press the space bar. The choice to have the initial position of the net drawn randomly in every trial is to minimize bias and foment moving of the net, as preliminary data suggested that if the net was always started at the centre of the screen (mean of the prior), PD patients would tend to not move it. The change from the mouse to the keyboard was done to help participants perform the task, given that preliminary data suggested older participants had trouble using the mouse. Participants could take as long as they wanted to decide where to place the net but had to wait for at least 1.26 seconds (to minimize accidental double presses counting as a response). After they pressed the button, the true position of the coin was revealed and participants would get one point added to the score if the coin was inside the net. The cumulative score across the experiment was shown to the participants at the end of each trial. A new trial would then begin 1.5 seconds later (see Fig. 1a). Given that the net covered the entire height of the screen only the horizontal location was relevant, making this a one-dimensional estimation task. We chose to have the task effectively one-dimensional in order to make the task easier for participants, and keep them interested. Participants completed two blocks of 150 trials each for a total of 300 trials per experiment. The first block had low prior uncertainty and the second one had high prior uncertainty. We had to decrease the total number of blocks performed (in comparison with [1]) because preliminary data indicated that older participants were getting too tired and unwilling to do 600 trials. After finishing the two blocks

in the second session, participants also performed 100 trials of a control block. Each trial took an average of 8.5 seconds and the total experiment lasted on average approximately one hour.

Thus, this experiment is a decision-making experiment under uncertainty (as one does not know exactly where the coin is, but has to guess it based on uncertain/noisy information), in which what is being estimated by the participants is a continuous variable (coin location in the x-direction).

Control Block

The control block was performed immediately following completion of the two experimental blocks, on the second test session. The task was identical to the main experiment, with the only difference being that the coin's location was shown at the onset of each trial and could be seen throughout, so there was no uncertainty about its position. The 100 trials comprising the control block were selected by randomly sampling 25 trials per condition from the main experiment blocks. Each trial of the control block repeated one of those sampled experimental trials, showing the same likelihood dot display that was shown in the experimental trial and using as the coin position the actual position to which the person moved the net at that trial. As in the main experiment, participants were awarded one point for successfully moving the net to the coin's position. This control block was done in order to understand if participants were able to see the dots, accurately move the net to the required place and press the button. All participants included in the analysis were able to perform the control block without apparent problems.

Data Analysis

Measures

Bayesian modelling of behaviour: In order to perform the task successfully, for every trial the participants should place the net in the most likely location of the hidden coin. Bayes rule provides an optimal way to estimate this location[9]:

$$X_{est} = \sigma^2_L/(\sigma^2_L+ \sigma^2_P)\mu_P+\sigma^2_P/(\sigma^2_L+ \sigma^2_P)\mu_L \qquad (1)$$

where X_{est} is the estimated position of the coin, σ^2_L and σ^2_P are the variances of the likelihood (current sensory information) and of the prior, respectively, and μ_L and μ_P are the respective means. For our experiment, the real μ_P (the mean of the prior) is always a constant, in this case 0.5 (the center of the screen). The mean of the likelihood (μ_L) for each particular trial can be considered the centroid of the cloud of dots, and it changes from trial to trial. We can then make a linear regression of the participant's estimated coin position, X_{est}, as a function of the centroid of the cloud of dots. The slope of this linear regression, which we will call the **sensory weight** (sw), characterizes how much the participant is weighting the current sensory information (likelihood), and, if people perform according to the optimum prescribed by Bayesian statistics, its value should be equal to the perceived $\sigma^2_P/(\sigma^2_L+ \sigma^2_P)$. A slope (weight) of zero suggests that participants do not take into account likelihood information and may be just taking into account prior information, and a slope of one suggests that participants only use likelihood information (see also Fig. 2a). A slope between zero and one suggests that participants are using information from both prior and likelihood, and the larger the slope the more they rely on the likelihood and less on the prior.

Sensory weight per condition: For each of the four conditions (pl,pL,Pl and PL), per session, the corresponding trials were obtained (75 trials per condition). From these, any trials that fell within the first 20 trials of each block were removed (as, for most of the analyses, we were not interested in the initial task adaptation effects), as well as any trials in which the participant did not move at all. This lead to an average of about 64±2 trials (mean ± s.d.) per condition. We then calculated, per condition, a linear regression of the participant's estimated coin position as a function of the centroid of the cloud of dots (in the x-direction). The slope of this linear regression was thus the [sensory] weight (sw) for this condition.

Participants' average reliance on current vs. prior information (average sensory weight): the weights associated with each of the four conditions were averaged to get an estimate of an individual's average relative reliance on current/likelihood information (i.e. $(sw_{pl}+ sw_{pL}+sw_{Pl} +sw_{PL})/4$).

Participants' performance: The proportion of correct trials, for each condition, was calculated as the number of trials in which the participant accurately guessed the position of the coin, divided by the total number of trials in that condition.

Sensitivity to prior uncertainty: Is the difference between the weights in one type of prior and the weights in the other type of prior, namely the weights obtained in the high prior uncertainty conditions vs. the low prior uncertainty conditions (i.e. $((sw_{Pl} +sw_{PL})-(sw_{pl}+ sw_{pL}))/2$).

Alternative sensitivity to prior uncertainty measure: measures how much the weights changed between blocks of different prior uncertainties, comparing it to the change in weights that occurs

within a prior uncertainty block. In order to calculate this measure, one first calculates the initial and final weights within each prior block (about 15 trials used, not separated by likelihood type). Then, the difference between the weights at the beginning of a block and the end of the previous block are calculated, and compared to the difference between the weights at the beginning and end of a block.

Sensitivity to likelihood uncertainty: Is the difference between the weights in one type of likelihood and the weights/slopes in the other type of likelihood, namely the difference in the weights obtained in the more precise (less uncertain) likelihood conditions and the ones associated with the more uncertain likelihood conditions (i.e. $((sw_{pl}+sw_{Pl})-(sw_{pL}+sw_{PL}))/2$).

For these calculations (except the ones that directly measure behaviour at the start of the block), we discarded the first 20 trials of every block to minimize the effect of task learning, and also the trials in which the participant did not move at all.

Analysing learning - change in the weights across time: For each session, for each 10 or 15 trials starting at trial 1 (i.e. trials 1-10; 11-20; 21-30; etc.), we calculated the corresponding weight, i.e. the slope of the linear regression of the participant's estimated coin position as a function of the centroid of the cloud of dots in the x-direction. Note that these trials were not separated by likelihood type, and no-move trials were also not discarded (but discarding them lead to similar results). For the "statistically optimal weights assuming the imposed variances" displayed in Supplementary Figure 2 the average of the 2 likelihood conditions was displayed, per prior uncertainty block (see "models" subsection below for details on how these weights are calculated).

Estimating participants' subjective prior means: Although the prior mean was given to the participants (the coin thrower is aiming at the center of the screen), each participant's subjective prior mean can also be estimated based on where they placed the net on each trial. From equation (1), and noting that $\sigma^2_L/(\sigma^2_L+\sigma^2_P) = 1 - \sigma^2_P/(\sigma^2_L+\sigma^2_P)$, if again we do the linear regression of the participant's estimated coin position as a function of the centroid of the cloud of dots, if we assume participants are behaving in a statistically optimal way (i.e. following equation (1)), then the intercept of this regression should be equal to $\mu_P*\sigma^2_L/(\sigma^2_L+\sigma^2_P)$. We can then obtain the participant's estimated prior mean (μ_P) by: $\mu_P = b_0/(1-sw)$, where b_0 is the intercept of the linear regression and *sw* the slope of the regression (weight), which we assume is equal to the perceived $\sigma^2_P/(\sigma^2_L+\sigma^2_P)$. To analyse the change in prior means across time we adopted the same procedure as above (doing the linear regression for each 10-15 trials). Note that, for sensory weights close to 1, the estimated prior mean can give arbitrarily large values, which explains the sudden "peaks" in Supplementary Fig. 4.

Analysing learning – estimating a participant's prior variance across time: Again from equation (1), if we assume that people perform according to the optimum prescribed by Bayesian statistics, their sensory weights (sw) should be equal to the perceived $\sigma^2_P/(\sigma^2_L+\sigma^2_P)$. Hence, if $sw = \sigma^2_P/(\sigma^2_L+\sigma^2_P)$, this equation can be rearranged to $\sigma^2_P = \sigma^2_L*sw/(1-sw)$. The sw can be obtained in each 10/15 trial bin as usual (the slope of the linear regression of the participant's estimated coin position as a function of the centroid of the cloud of dots). If we furthermore assume that the variance in the likelihood (σ^2_L) can be approximated by the actual average sample variance of the cloud of dots (the likelihood), this gives us an estimated prior variance of $s^2_P = s^2_L*sw_{bin}/(1-sw_{bin})$, where s^2_L is the average sample variance of the cloud of dots for each time bin and sw_{bin} is the weight for the time bin. From this formula it can

also be seen that, like the estimated prior mean, for weights close to 1 the estimated participant's subjective prior variance (s^2_P) can give arbitrarily large values. Furthermore, because occasionally the experimental weights in a time bin can give values slightly above 1, this would lead to arbitrarily high negative values, which does not make sense for a variance (that is always positive). Hence, for the calculations we instead used a "transformed" binned slope, namely we calculated the logistic of the sw_{bin} = (1/ (1+e^{-swbin})) to have the binned weight vary from 0 to 1. The formula used to estimate each participant's prior variance was $s^2_P = s^2_L * sw'_{bin}/(1-sw'_{bin})$, where sw'_{bin} denotes the logistic-transformed sensory weight for that time bin.

As the data was transformed by a monotonically positive formula (the logistic), the relation between the different values is kept, so it is possible to compare the different populations. However, the specific numeric values have changed. Hence, to display the "imposed prior variances" (Supplementary Fig. 3), we transformed the variances using a similar procedure, in order to be directly comparable. Namely, we first calculated the "statistically-optimal" weight assuming the actual imposed prior variance for that block (σ^2_P) and the average sample variance of the cloud of dots for that time bin (s^2_L), so that $sw_{opt} = \sigma^2_P/(s^2_L + \sigma^2_P)$; and then we transformed this weight using the same procedure as above (sw'_{opt}=1/ (1+e^{-swopt})). Finally, we extracted back the transformed imposed prior variance by $s^2_L * sw'_{opt}/(1-sw'_{opt})$.

<u>*Trial-by-trial sensory weights*</u>: Equation (1) can be rewritten to calculate a "trial-by-trial" weight:

$$sw_{trial} = (X_{est} - \mu_P)/ (\mu_L - \mu_P) \qquad (2)$$

Where X_{est} is the participant's estimated position of the coin on that trial (in our experiment, where they placed the net), μ_P the mean of the prior (assumed always 0.5, the centre of the screen) and μ_L the mean of the likelihood (here calculated as the centroid of the cloud of dots for that trial). If the mean of the prior and the likelihood are very near each other this sw_{trial} can become arbitrarily large, hence we calculated the logistic of the sw_{trial} (1/ (1+$e^{-swtrial}$)) to have the trial-by-trial weight vary from 0 to 1.

<u>*Intra-individual variability in the sensory weights:*</u> This is simply the variance of the trial-by-trial weights (over the whole session).

<u>*Correlation between the trial-by-trial sensory weights and the sample standard deviation*</u>: We calculated the Pearson's correlation coefficient between the logistic-transformed trial-by-trial weights and the sample standard deviation (s.d.) of the cloud of dots (likelihood) associated with each trial. As only the x-dimension was relevant for the task, we only used the sample s.d. in the x (horizontal) location. To confirm that participants were indeed using mainly the sample s.d. on the x-direction, we calculated a similar correlation but in which we used the s.d. in the y-direction (vertical). We observed that the correlation is not as strong/significant when the s.d. in the y-direction is used, suggesting that, as expected, participants were using mainly the sample s.d. of the dimension of interest (see also Supplementary Results).

Models

Statistical-optimal model assuming the imposed variances: Bayesian theory tells us that the statistically-optimal sensory weight can be obtained from $\sigma^2_P/(\sigma^2_L+\sigma^2_P)$, where σ^2_P is the variance associated with the prior and σ^2_L is the variance associated with the likelihood. In this model we are assuming that the participants' perceived variances of the prior and the likelihood are equal to the experimentally-imposed variances, i.e. with a prior variance of $\sigma^2_p = 0.025^2$ or $\sigma^2_P = 0.085^2$ in unit-less screen coordinates and a sensory/likelihood variance σ^2_L estimated by σ^2_L = variance(cloud of dots)/ number of dots (so in our case, $\sigma^2_l = 0.06^2/5$ or $\sigma^2_L = 0.15^2/5$)[9]. Note that this is a simplification and that the participants' actual perceived variances are likely higher that these. However, as each participant's subjective variance is not directly observable, the experimentally-imposed variances can give us a good bench mark from which to compare actual behaviour.

Senses-only model: If participants only rely on current sensory information, not taking into account prior information, then the slope of linear regression of the participant's estimated coin position as a function of the centroid of the cloud of dots should be always equal to 1, regardless of the associated actual uncertainties[10]. This can also be seen as a Bayesian model in which participants, for some reason, have a subjective prior variance much higher than the subjective likelihood variance, making $\sigma^2_P/(\sigma^2_L+\sigma^2_P)$ approach 1.

Comparing models: to directly compare the performance of the Statistically-optimal model assuming the imposed variances and the Senses-only model we first calculated the mean absolute errors associated with each model (mean($|s_{model}-s_{obtained}|$), per participant, and then did a paired t-test comparing those values.

Statistical Analysis

Throughout the paper we used analysis of variance (ANOVA) and t-tests, being the specific test employed mentioned along with the result. We use ANOVA if more than one factor or more than two groups were analysed at the same time. Throughout, we use two-sided tests to obtain the p-value. If the assumptions of normality or homoscedasticity were not met, then we used the non-parametric equivalent test (e.g. Mann-Whitney U-Test instead of unpaired t-test).

For the ANOVA, if we were comparing two different populations (e.g. PD patients vs. Controls), then we used a mixed-effects ANOVA with subject as a random effect nested under population (which we called "repeated-measures" ANOVA), and the remaining factors as fixed effects (e.g. population type; prior; likelihood; session). If we were interested in analysing the interactions, first the full model is fit, and then the specific interactions are analysed in subsets of the data (the specific factors included are reported in parenthesis after each ANOVA result). If comparing the same participants in two different situations (e.g. on-medication *vs.* off-medication) then we used a mixed-effects ANOVA with subject as a random effect nested under population (also called repeated-measures ANOVA), and the remaining factors as fixed effects; or alternatively a paired t-test (whichever most appropriate, namely if only the effects of one factor are being analysed or more). Incorporating subject as a random effect was done to account for the dependencies within the sample.

Generally, for the PD population first the effect of medication is analysed (using ANOVA whenever possible), and if there is no significant effect of medication the results from the on-medication and off-medication sessions are pooled together. If medication has a significant effect, then when comparing PD patients with controls these sessions are analysed separately. When a new ANOVA is introduced, the fixed and random factors used are explained in parenthesis the first time a corresponding p-value is reported, and in the subsequent times that the results related to the same ANOVA are shown only the corresponding F-value and p-value are reported.

Data availability

Data and code can be made available, upon request to the corresponding author.

SUPPLEMENTARY TABLE

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
L-DOPA	0	1	1	0	1	1	1	0	1	1	1	1	1	0	1
DAG	1	1	0	1	0	0	1	0	1	1	0	0	1	0	0
MAOBi	1	1	0	0	0	0	1	1	0	0	0	1	1	1	0
Amant	1	0	0	0	1	0	0	1	0	0	0	0	0	0	1

Supplementary Table 1 – Dopaminergic medications taken by the PD participants. Represented are the dopaminergic medications that each Parkinson's Disease (PD) participant was taking, with 1 indicating that they were taking that particular kind of dopaminergic medication, and 0 indicating that they were not. Each participant is represented in a different column (n = 15 total). Legend: DAG: dopamine agonists; MAOBi: MAOB inhibitors; Amant: Amantadine.

SUPPLEMENTARY RESULTS & SUPPLEMENTARY FIGURES

Models

While we saw that PD patients show qualitative agreement with Bayesian predictions we may ask if they quantitatively agree. We can thus analyse if the weights obtained for each participant, per condition, are quantitatively different from the statistically optimal weights assuming the imposed variances (see Supplementary Methods for details). Here, for simplification, we are going to use the average behaviours per condition, and not potential trial-to-trial variations. We found no average significant difference between the PD participants' weights and the statistical optimal weights (t_{14}=1.89, *p-val*=0.08, one-sample t-test, n=15; [-0.016,0.255] 95% CI for mean difference; see Supplementary Fig. 1a; see also Supplementary Fig. 2 for variation across time). Participants were closer to the statistically optimal values when prior and current uncertainties are either both low or both high (pl or PL, Supplementary Fig. 1a), not being as extreme as predicted (very close to 0 or 1) in the other two conditions. In our task, the behaviour of PD patients could be quantitatively captured with a Bayesian-like approach assuming the imposed variances.

To see if medication state or disease affected how close participant's weights were from the statistically optimal weights assuming the imposed variances, we can compare the mean absolute distances between their weights and the statistically-optimal weights. PD patients off-medication were not further away from the optimum compared to when they were on-medication ($F_{1,13}$ = 0.02, *p-val* = 0.89 for main effect of medication state; [-0.041,0.046] 95% CI for mean difference; mixed-effects ANOVA with subject as a random effect (i.e. repeated-measures ANOVA) and session and medication state as fixed factors; Supplementary Fig. 1b; see also Supplementary Fig. 2). There was also no difference in the mean absolute deviation from the statistical optimum between PD patients and controls ($F_{1,29}$=1.33, *p-val*=0.26 for main effect of population type; [-0.029,0.102] 95% CI, mixed-effects ANOVA with subject as a random effect nested under population (repeated-measures ANOVA), and population and session as fixed factors; Supplementary Fig. 1b). Medication state and disease did not significantly affect how close participants' weights were from the Bayesian predictions assuming the imposed variances.

As session had a main effect on the weights, it could also have impacted the distance to the statistical optimum. Indeed, there was a significant effect of session ($F_{1,13}$ = 16.59, *p-val* = 0.001; [0.039,0.126] 95% CI; repeated-measures ANOVA with session and medication state as fixed factors; Supplementary Fig. 1). Similarly, when comparing PD patients with Controls, there was a significant effect of session ($F_{1,29}$ = 6.2, *p-val* = 0.02; repeated-measures ANOVA with population and session as fixed factors). Furthermore, while controls had similar distances in both sessions, PD patients were further away from the statistical optimum in the second session (difference in mean absolute distances between the 2nd and 1st sessions: t_{14} = -0.018, *p-val* = 0.99 for controls, [-0.048,0.047] 95% CI for mean difference; t_{14} = 4.22, *p-val* = 0.0008 for PD patients, [0.04,0.123] 95% CI; paired t-test). Indeed, while in the first session there was no average significant difference between PD patient's weights and the statistically optimal weights (t_{14} = 1.27, *p-val* = 0.22, one-sample t-test; [-0.05,0.194] 95% CI), these were significantly different in the second session (*p-val* = 0.048, w = 95, Wilcoxon signed-rank test). PD patients (but not controls) were further away from the statistical optimum in the second session, suggesting stronger carryover effects in the PD group.

In order to calculate the "statistically optimum" weights the experimentally imposed variances were used. This assumes that the participants actually knew the experimentally imposed variances. Realistically, the participants could only estimate those values, and hence their subjective uncertainty was probably at least slightly higher than the imposed variances. However, this provides a useful simplification, and gives us a benchmark from which we can try to infer the participant's underlying subjective uncertainties.

Another potential simple model would be if participants had extremely uncertain priors but relatively certain current sensory information (e.g. if they didn't learn the prior). If that were the case, PD patients should have a sensory weight around 1 in all conditions, irrespective of the experimentally imposed prior uncertainty ("senses only" model). The weights of PD patients were significantly different from the "senses only" weights ($t_{14} = 5.6$, $p\text{-}val<10^{-4}$, one-sample t-test; [0.221,0.491] 95% CI for mean difference; see Supplementary Fig. 1a), indicating that this model fails to describe the data. Comparing directly the mean absolute errors obtained using the "senses only" model versus the model using the statistically optimal weights assuming the experimentally imposed variances ("stat. opt. IV"), we can see that, for the data obtained during the first session, the model using the statistically-optimal weights produced smaller errors ($t_{14} = 2.23$, $p\text{-}val = 0.043$, paired t-test, [0.0055,0.287] 95% CI; see Supplementary Fig. 1c). In the second session there was no significant difference in the performance of the two models ($t_{14} = -0.34$, $p\text{-}val = 0.741$, paired t-test, [-0.202,0.147] 95% CI). Thus, altogether, the estimated weights of the participants were more closely approximated with the statistical optimal weights assuming the imposed variances than with weights that completely disregard prior information, although carryover session effects diminished this effect in the second session.

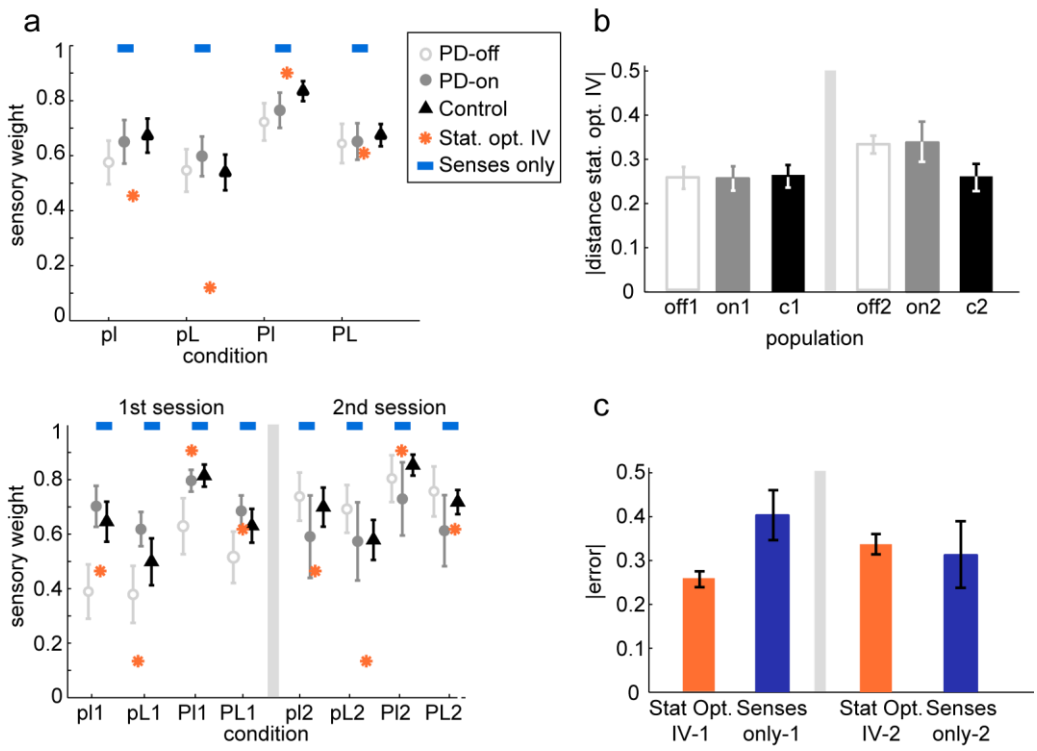

Supplementary Figure 1 – Model comparison. a) Average sensory weights ± standard error of the mean (s.e.m.) for PD patients off-medication (open light grey, ○), PD patients on-medication (filled dark grey, ●) and controls (black, ▲), divided per condition (pl,pL,Pl,PL). Overlaid are the statistically-optimal weights assuming the imposed variances (in orange, *) and the "senses-only" weights (in blue, ■), per condition. Top: average across both sessions; Bottom: separated by session. Similar to Fig. 2b (top) and Fig. 2d (bottom), but with the model predictions overlaid. PD patients' weights were not significant different from the statistical optimal weights (t_{14}=1.89, *p-val* = 0.08, one-sample t-test, n=15) but they were significantly different from the "senses-only" weights (t_{14} = 5.6, *p-val*<10^{-4}, one-sample t-test). **b)** Mean absolute distance between the participant weights and the statistically optimal weights, for PD patients off-medication (off), on-medication (on) and controls (c), averaged by condition but separated per session (data from 15 PD patients and 15 controls). There were no main effects of medication state ($F_{1,13}$ = 0.02, *p-val* = 0.89) but there was a significant effect of session ($F_{1,13}$ = 16.59, *p-val* = 0.001; repeated-measures ANOVA with session and medication state as fixed factors). Similarly, when comparing PD patients with Controls, there was no significant effect of population ($F_{1,29}$= 1.33, *p-val* = 0.26) but there was a significant effect of session ($F_{1,29}$ = 6.2, *p-val* = 0.02; repeated-measures ANOVA with population and session as fixed factors). Furthermore, while controls had similar distances in both sessions, PD patients were further away from the statistical optimum in the second session (difference in mean absolute distances between the 2nd and 1st sessions: t_{14} = -0.018, *p-val* = 0.99 for controls; t_{14} = 4.22, *p-val* = 0.0008 for PD patients; paired t-test). **c)** Mean absolute distance between the PD patients' weights (both on and off-medication) and the statistically optimal weights assuming the imposed variances (orange) or the "senses-only" weights (blue), separated by session. Error bars represent s.e.m.. During the first session, the model using the statistically-optimal weights assuming the experimentally imposed variances produced smaller errors compared to the senses-only model (t_{14} = 2.23, *p-val* = 0.043, paired t-test). In the second session there was no significant difference in the performance of the two models (t_{14} = -0.34, *p-val* = 0.741, paired t-test).

Analysing learning: Change in prior and sensory weights across time

We would like to know how PD patients learn. The sensory information is directly observable and does not need to be learned, and the mean of the prior was given to the participants (see Supplementary Methods, and also Supplementary Fig. 3). Therefore, changes in the weight across time should indicate changes in the participants' perceived prior variance. PD patients off-medication had lower initial sensory weights compared to both Controls and PD patients on-medication (Supplementary Fig 2). Looking within the first block (narrow prior, first session), there was a near significant effect of population type ($F_{2,0}$=2.56, *p-val*=0.096; repeated-measures ANOVA with population and time bin as factors; significant if only comparing PD patients *on* vs. *off*) and a significant effect of time bin ($F_{9,0}$=2.42, *p-val*=0.012), but no significant population by time bin interaction ($F_{18,0}$=0.76, *p-val*=0.75). All participants started with relatively higher weights, and decreased them shortly after. Initial learning happened quickly, as the main changes in weights, within a given block, happened within the first 15-20 trials.

We would also like to know how learning relates to the optimal behaviour assuming the imposed variances. All participants increased their weights when the prior variance increased, indicating that both PD patients and Controls were able to detect the changes in prior uncertainty and respond appropriately to it (Supplementary Fig. 2; see also Supplementary Fig. 5). For the large prior uncertainty block (*P* block), both controls and PD patients on-medication were close to the "statistical optimum", as well as the PD patients off-medication, although their average weights were slightly lower. Interestingly, for the small prior uncertainty block (*p* block) in the 1st session, the PD patients off-medication were the ones more near this "statistical optimum". Note that, although PD patients off-medication had lower initial weights, they were as sensitive to changes in the prior uncertainty as PD patients on-medication and controls. Thus, although the starting weights were different, all participants were able to learn the different prior uncertainties and, moreover, approached near-optimal weights.

Does learning happen with the same speed across the groups? In the first prior block all participants adapted with similar speeds (Supplementary Fig. 2b). Using a bin size of 15 trials we can see a significant difference between the weights in the first time bin of the second prior block (the high prior uncertainty block) and the weights of the last time bin of the first block (of low prior uncertainty), both for PD patients and Controls (t_{14} = 2.22, *p-val* = 0.043, [0.006,0.362] 95% CI, for PD patients; t_{14} = 2.22, *p-val* = 0.044, [0.005,0.277] 95% CI, for Controls; paired t-test). Furthermore, this weight update was similar between PD patients and Controls (t_{28} = 0.41, *p-val* = 0.68, [-0.171,0.257] 95% CI; unpaired t-test) and between PD patients on and off medication (*p-val* = 0.36, t_{13} = -0.94, [-0.517,0.203] 95% CI; unpaired t-test). Both PD patients and Controls could learn a new prior uncertainty relatively fast (within 15 trials there was already a significant difference), and there was no strong effect of medication status or disease on how fast participants adapted to the new prior uncertainty.

So how optimal is this process? Unfortunately, the optimal learning rate depends on how fast one thinks the world is changing[11-13]. In our task the changes in prior ("world" changes) only occur once per session, which is not enough to have a good estimate about the variability of the world (in terms of prior variance). For our experiment we were interested in knowing if PD patients could learn/adapt to different prior uncertainties, but not specifically interested in how exactly (e.g. how fast) this learning occurred. Future studies could take a similar paradigm but in which there are several prior blocks and different degrees of volatility in the world(like in [11,12]), and specifically look at how PD patients compare to controls in terms of updating their priors and choosing an appropriate learning rate.

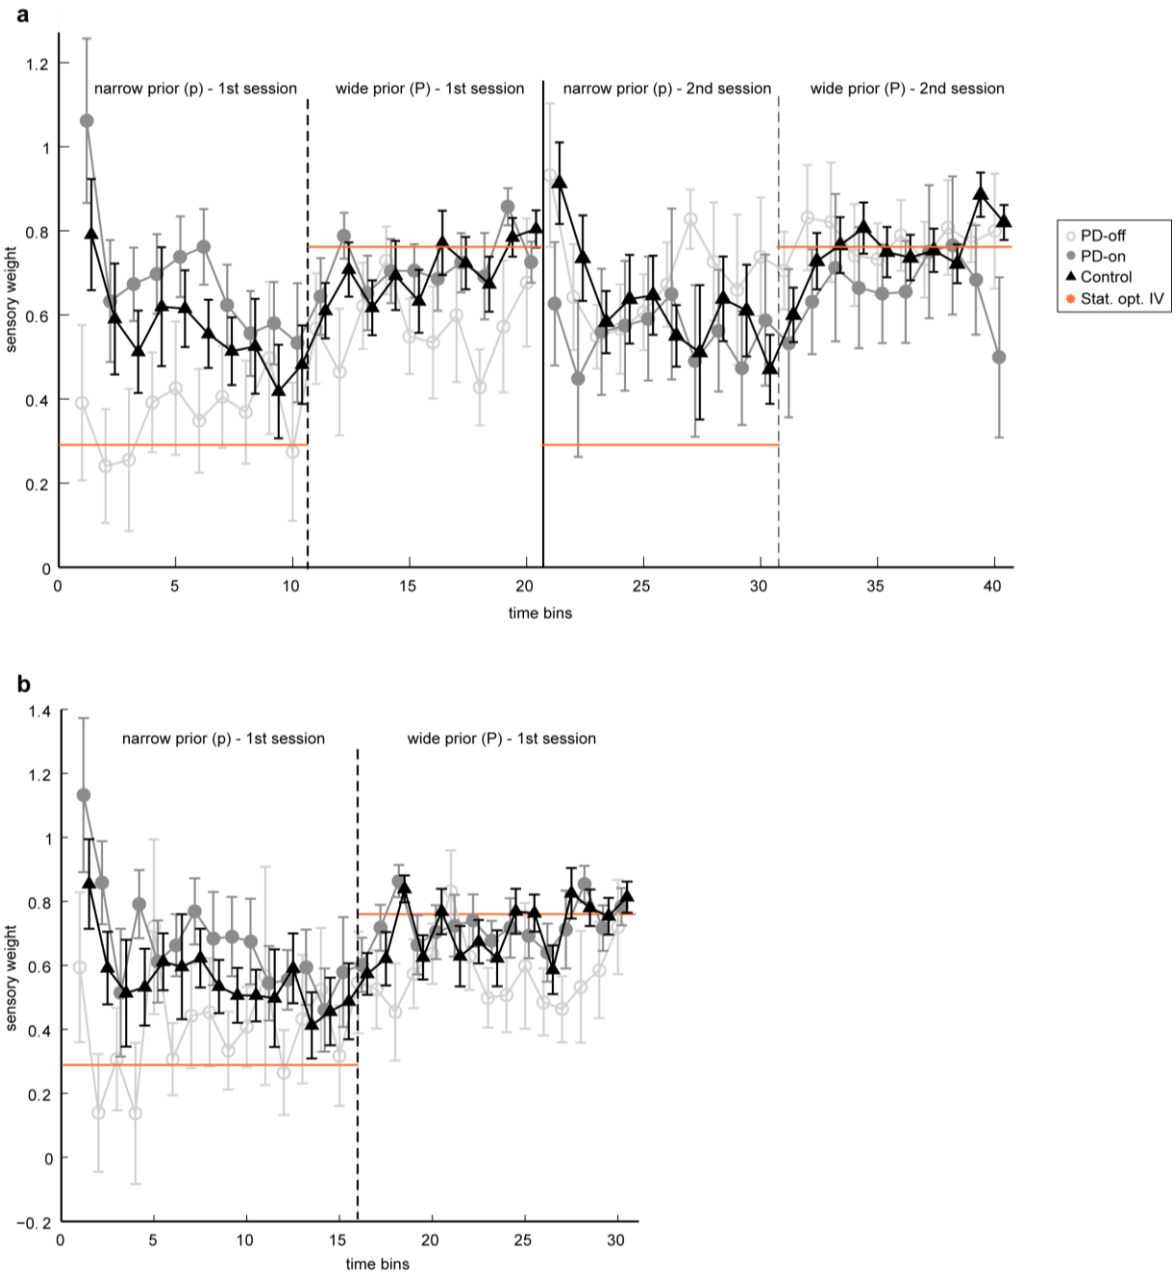

Supplementary Figure 2 – Changes in the sensory weight across time. a) Average sensory weights (± s.e.m.) across time, with each weight having been calculated over 15 consecutive trials (1-15, 16-30, etc.), not separated by likelihood type. Data shown for PD patients off-medication (open light grey,○), PD patients on-medication (filled dark grey,●) and controls (black,▲). The vertical dashed black lines represent changes in the prior uncertainty block within the same session, while the vertical black solid line represents a change in session/day. The first prior block of each session was the low uncertainty prior (p) and the second block was the higher uncertainty prior (P). There were 150 trials per block for a total of 300 trials per session, 600 trials overall. Data from n=15 controls and n=15 PD patients (from which n=8 were on-medication in the first session and n=7 started off-medication). Note that the PD patients that were on-medication in the first session were off-medication in the second session, and vice-versa, so the filled dark grey circles (●) represent the average of n=8 PD participants on-medication in the first session but are the average of n=7 participants in the second session (corresponding to the PD patients that were off-medication in the first session and are now on-medication); and vice-versa for the open light grey circles (○). The orange horizontal lines represent the statistically optimal weights assuming the imposed variances, averaged across the 2 likelihood conditions but separated by prior. **b)** Same as *a)*, but using only 10 trials per bin. Just the first session is shown.

If we assume that participants are behaving in a statistically-optimal way and that their perceived likelihood uncertainty is the experimentally observed one, then the participants' prior variances can be estimated directly from the sensory weights (see Supplementary Methods and Supplementary Fig 3). Our estimate follows the general shape of the change in weights, with PD patients off-medication starting with relatively lower estimated prior variances compared to PD patients on-medication and controls, and all participants increasing their estimated prior variance when they changed to a block of higher prior variance (see Supplementary Figs. 2 and 3). The prior variance update was similar between PD patients on and off medication, and between PD patients and Controls. Furthermore, both PD patients and controls estimated prior variances were close to the transformed experimentally imposed ones (see orange line in Supplementary Fig. 3 and Supplementary Methods). The main exception is in the small prior uncertainty block, especially in the second session. Altogether these results show that both PD patients and controls were able to adapt to the new prior uncertainty blocks, although they may have a harder time switching back to a lower prior uncertainty estimate when the previous prior block was more uncertain[2].

Note that this estimated subjective prior variance could not be observed directly (at least behaviourally) and instead was obtained assuming that the participant's subjective likelihood uncertainty was the observed veridical one (see Supplementary Methods). Hence, the estimated initial lower prior uncertainty in PD patients off-medication was a direct result of the observed lower sensory weights. However, it could as well be that the reason behind these lower weights is related to an overestimation of the current sensory uncertainty and not an underestimation of the prior uncertainty. What we can infer is that PD patients off-medication are still able to react to changes in prior uncertainty to the same level as controls, as they significantly changed their sensorial weights when the prior uncertainty changed.

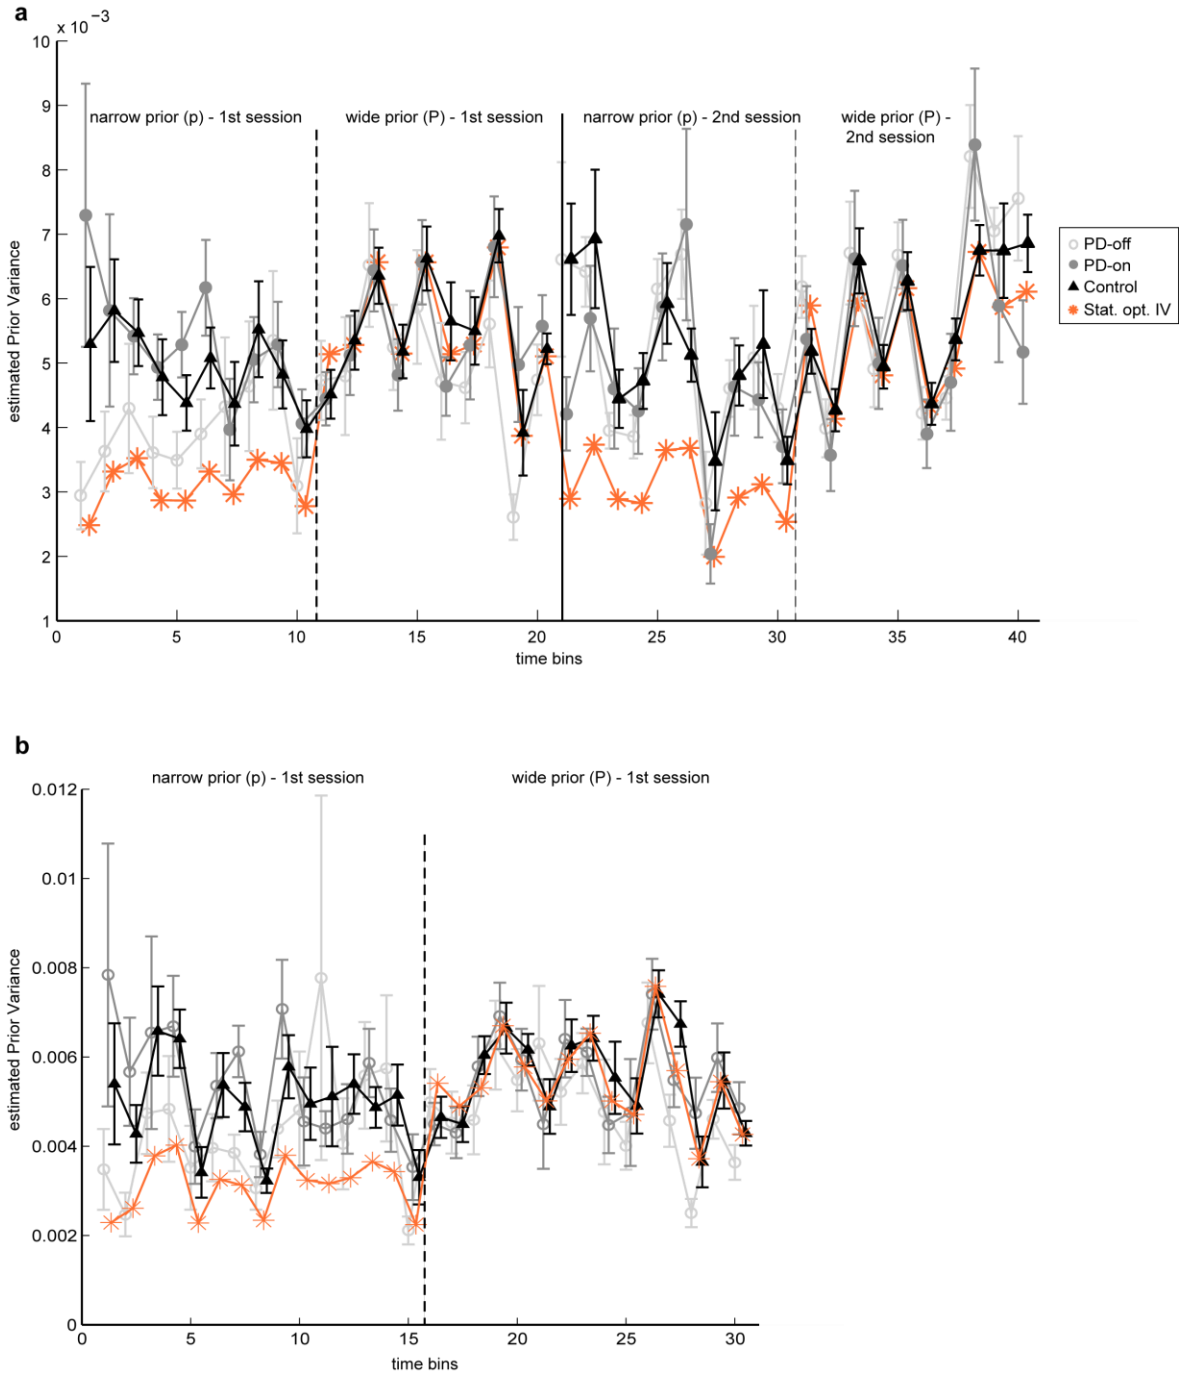

Supplementary Figure 3 – Changes in the estimated prior variances across time. a) Average estimated prior variances (± s.e.m.) across time, with each prior variance having been calculated over 15 consecutive trials (1-15, 16-30, etc.). To calculate each participant's subjective prior variance at each time bin, we assumed that participants were behaving in a statistically-optimal way and that their perceived likelihood uncertainty was the experimentally observed one during that time bin. Data shown for PD patients off-medication (open light grey, ○), PD patients on-medication (filled dark grey, ●) and controls (black, ▲). The vertical dashed black lines represent changes in the prior uncertainty block within the same session, while the vertical black solid line represents a change in session/day. Data from n=15 controls and n=15 PD patients (from which n=8 were on-medication in the first session and n=7 started off-medication). The orange lines represent the actual imposed prior variances, transformed in the same way as the participants' estimated prior variances (see Supplementary Methods for details). The random fluctuations are a result of the random variations in the trial bins' sample likelihood uncertainty. **b)** Same as *a)*, but using only 10 trials per bin. Just the first session is shown.

Prior mean

Participants were told the mean of the prior ("the coin thrower aimed at the centre of the screen"), but may not have used this information. To understand if participants were indeed using the correct prior mean, we estimated participants' subjective prior mean (see Supplementary Methods for details). Despite the sudden "peaks" sometimes observed in the estimated prior mean (which happen when the participant's sensory weight is very near 1, see Supplementary Methods), we can see that, on average, both Controls and PD patients had the correct prior mean (around 0.5, the centre of the screen; Supplementary Fig. 4). Furthermore, they had the correct prior mean from the beginning of the game, indicating that they indeed understood and correctly used the prior mean information.

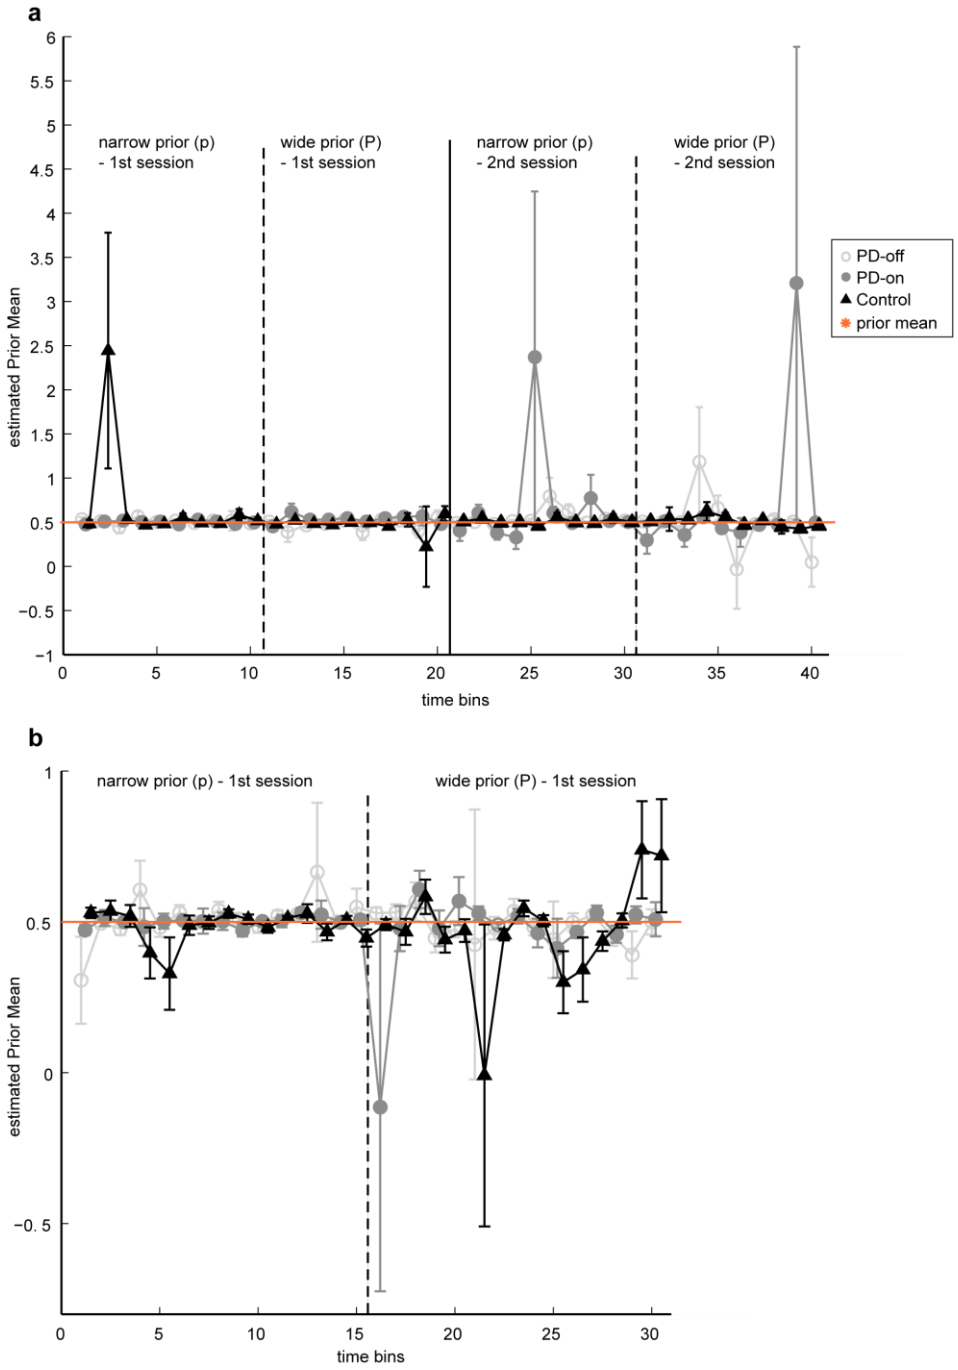

Supplementary Figure 4 – Changes in the estimated prior mean across time. a) Average estimated prior mean (± s.e.m.) across time, with each prior mean having been calculated over 15 consecutive trials (1-15, 16-30, etc.), not separated by likelihood type. To calculate each participant's subjective prior mean at each time bin, we assumed that participants were behaving in a statistically-optimal. Data shown for PD patients off-medication (open light grey, ○), PD patients on-medication (filled dark grey, ●) and controls (black, ▲). The vertical dashed black lines represent changes in the prior uncertainty block within the same session, while the vertical black solid line represents a change in session/day. Data from n=15 controls and n=15 PD patients (from which n=8 were on-medication in the first session and n=7 started off-medication). The orange line represents the actual prior mean (0.5, corresponding to the centre of the screen). The occasional strong deviations from 0.5 are likely a direct result from a participant's sensory weight being very near 1, which leads to arbitrarily high or low estimated prior mean values (see Supplementary Methods for details). **b)** Same as *a)*, but using only 10 trials per bin. Just the first session is shown.

Behaviour of age and gender-matched Controls

To understand if healthy older adults can also react to changes in prior and sensory uncertainty, we analysed the sensory weights within the age-matched control population (63.4 ± 11.6 years). We found that, within the control population, weights were affected by both prior uncertainty and likelihood uncertainty ($F_{1,102}$ = 22.41, *p-val* < 10^{-5}, P>p, [0.086,0.211] 95% CI, for main effect of prior; $F_{1,102}$ = 21.91, *p-val* < 10^{-5}, l>L, [0.085,0.209] 95% CI, for main effect of likelihood uncertainty; mixed-effects ANOVA with subject as a random factor and prior and likelihood uncertainty and session as fixed factors; see black triangles in Fig. 2b). Session number had also a significant effect ($F_{1,102}$ = 4.28, p = 0.041, 2>1, [0.003,0.127] 95% CI). Furthermore, similarly to the PD population, older adult (age-matched) controls had higher sensory weights when current information was more reliable or when the prior was more uncertain. This indicates that healthy older adults are able to react to changes in prior and sensory uncertainty and modify their corresponding weights in a way predicted by Bayesian statistics.

PD patients were sensitive to changes in prior uncertainty

To understand better how PD patients reacted to differences in prior uncertainty, we analysed directly the changes in weights that occurred between blocks of different prior uncertainties ("sensitivity to prior uncertainty"). We did not find a significant effect of medication on PD patients' sensitivity to prior uncertainty (t_{14} = -0.78, *p-val* = 0.45, n = 15, paired t-test; [-0.143,0.067] 95% CI for mean difference in sensitivity; Supplementary Fig. 5), thus we pooled the data from PD patients when on and off-medication together. There was no significant difference between PD patients and controls in their sensitivity to prior uncertainty ($F_{1,29}$ = 1.02, *p-val* = 0.32, [-0.136, 0.046] 95% CI, for main effect of population type; mixed-effects ANOVA with subject as a random effect nested under population type and session and population type as fixed effects; Supplementary Fig. 5). There was also no significant effect of session ($F_{1,29}$ = 0.58, *p-val* = 0.45, [-0.107, 0.049] 95% CI). PD patients, as well as controls, had sensitivity to priors significantly different from zero (t_{14} = 4.28, *p-val* = 0.0008, [0.052, 0.155] 95% CI; and t_{14} = 3.97, *p-val* = 0.001, [0.068, 0.229] 95% CI; respectively; one-sample t-tests), i.e. they significantly changed their weights when the corresponding prior uncertainty changed, having higher sensory weights (hence lower prior weights) when the prior was more uncertain. Together, this indicates that PD patients were able to learn the prior uncertainty to the same level as controls.

However, it could be that what was happening was just a constant increase in the sensory weight value over time, irrespective of the prior uncertainty in each block, which would nevertheless result in an apparent change just because one block happens after the other. To test this, we developed an alternate sensitivity to prior measure, comparing the change in weights between the beginning of a new prior uncertainty block and the end of the previous block with the change in weights happening within the same block. We found that, for PD patient, this alternate measure of prior uncertainty is also significantly different from zero (t_{14} = 2.46, *p-val* = 0.028, [0.02, 0.292] 95% CI, one-sample t-test, n = 15). This suggests that the observed change in behaviour with the change in prior uncertainty is not only due to a generalized increase in sensory weights throughout the experiment, reinforcing the idea that PD patients are indeed able to learn different prior uncertainties.

Although there was no significant effect of medication state on the sensitivity to prior or likelihood uncertainty, if anything these effects were opposite (see Fig. 3 and Supplementary Fig. 5). While the administration of dopaminergic medication resulted in an increase in the sensitivity to changes in likelihood uncertainty (Fig. 3), it resulted in a decrease in the sensitivity to prior uncertainty (Supplementary Fig. 5). Nevertheless, these effects were non-significant. Thus, while we observed significant effects of dopaminergic medication on the PD patient's average weights, the effects of the medication on the sensitivity to prior and likelihood uncertainty were less clear and potentially opposite.

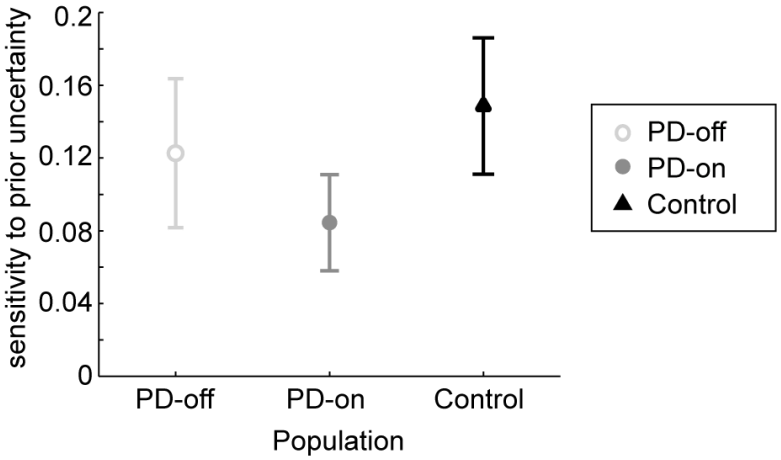

Supplementary Figure 5 - **Sensitivity to prior uncertainty.** Sensitivity to prior uncertainty, separated by population type. It is calculated as the average difference in the sensory weights between the high prior uncertainty conditions and the low prior uncertainty conditions. Error bars represent s.e.m. Represented is the average sensitivity to prior uncertainty (including both sessions) ± s.e.m. for PD patients off-medication (open light grey circles, ○), PD patients on-medication (closed dark grey circles, ●) and controls (black triangles, ▲). PD patients, as well as controls, had sensitivity to priors significantly different from zero (t_{14} = 4.28, *p-val* = 0.0008 and t_{14} = 3.97, *p-val* = 0.001, respectively; one-sample t-tests). There was no significant difference between PD patients and controls in their sensitivity to prior uncertainty ($F_{1,29}$ = 1.02, *p-val* = 0.32 for main effect of population type, mixed-effects ANOVA with subject as a random effect nested under population type and session and population type as fixed effects).

Lower sensitivity to likelihood uncertainty not only a motor or visual effect

It could be that the lower sensitivity to likelihood uncertainty observed in PD patients is related to some motoric deficit (e.g. direct saccade generation) or poor visual acuity. To rule out simple motoric and visual deficits we had participants perform a "control Block" task, in which they also saw the same cloud of dots and had to move the net to where the target dot ("coin") was, but in which there was no uncertainty because the coin location was shown from the beginning of the trial. We found no difference between either PD patients off and on medication ($U_{8,7}$ = 20.5, *p-val* = 0.39, Mann-Whitney U-Test) or between PD patients and controls ($U_{15,15}$ = 121.5, *p-val* = 0.663, Mann-Whitney U-Test) in their performance in the control block, with both groups being able to locate the coin and move the net to it with on average at least 98% accuracy (see also Supplementary Fig. 5c-d). This indicates that PD patients were able to see the target dot on the screen and correctly move the net to it, suggesting that it is not simply a motor/visual effect driving the results in this task.

Future studies could analyse in more detail what may be the specific difference in the processing of information PD patients have in relation to controls. They were still able to see the dots on the screen, suggesting it is not simply a vision/saccade generation problem. Participants had all the time they wanted for each trial, so time needed to perform a saccade should not be an issue. They could be less able to notice dots more on the periphery of the visual scene. However, then we would expect the deficit to be more pronounced in the wide prior uncertainty condition, as likelihood dots are more at the periphery of the visual scene then. Instead, we observed the opposite, with PD patients less sensitive to changes in current sensory uncertainty when in the narrow (less uncertain) prior condition. Alternatively, PD patients could be less able to estimate the centroid of the cloud of dots. However, then again we would not expect an asymmetry in the sensitivity to likelihood uncertainty depending on the prior type. A future experiment could show PD patients the cloud of dots and ask them to move the net to the sample mean location, which would be mainly the same as the current task but with prior information about previous target locations being irrelevant.

Which axis (X or Y) are participants paying most attention to?

While we assumed participants would only use the likelihood data from the x-dimension, it could be that they relied equally on the x and y dimension, or even that they just cared about the y-dimension of the data. Both PD patients and Controls showed consistent negative correlations between their trial-by-trial weights and the sample s.d. of the cloud of dots in the x-direction (see Results and Fig. 4). We can then compare these correlations with similar ones using the sample s.d. of the cloud of dots in the y-direction. If participants cared equally about the x and y dimensions then we expect these correlation values to be similar, while if they care more about one dimension than the other we expect those correlations to be stronger in that dimension. Participant's trial-by-trial weights were more significantly (negatively) correlated with the sample s.d. of the likelihood dots in the x-direction than in the y-direction (t_{14} = -3.26, *p-val* = 0.006 for PD patients, [-0.064, -0.013] 95%CI for difference in correlations; t_{14} = -3.77, *p-val* = 0.002 for controls, [-0.074, -0.02] 95%CI, paired sample t-test for difference in mean correlation values). Participants cared more about the variability in the X-direction than in the Y-direction, as one would expect given the design of the task.

Performance on the task

We also wanted to understand how performance was affected in the coin-catching task. Analysing within PD patients, we found that there was a significant main effect of prior uncertainty and of likelihood uncertainty ($F_{1,101}$=29, p-val<10^{-6}, p>P, [0.046, 0.1] 95%CI, for main effect of prior; $F_{1,101}$=108.5, p-val<10^{-6}, l>L, [0.115, 0.168] 95%CI, for main effect of likelihood uncertainty, mixed-effects ANOVA with subject as a random effect and medication state, session, prior and likelihood uncertainty as factors; see Supplementary Fig. 6a), with performance decreasing when the uncertainties were increased. This is not surprising, given that both types of uncertainty may be expected to change the precision of participants' estimates and thus their expected task performance. Session did not have significant effect ($F_{1,101}$=1.6, p-val = 0.21; [-0.044, 0.01] 95%CI). There was a significant main effect of medication state, with PD patients guessing where the hidden target was significantly fewer times when they were off-medication compared to when they were on-medication ($F_{1,101}$=4.7, p-val = 0.032; [0.003, 0.057] 95%CI). Thus, administration of dopaminergic medication had a positive impact in performance for this task.

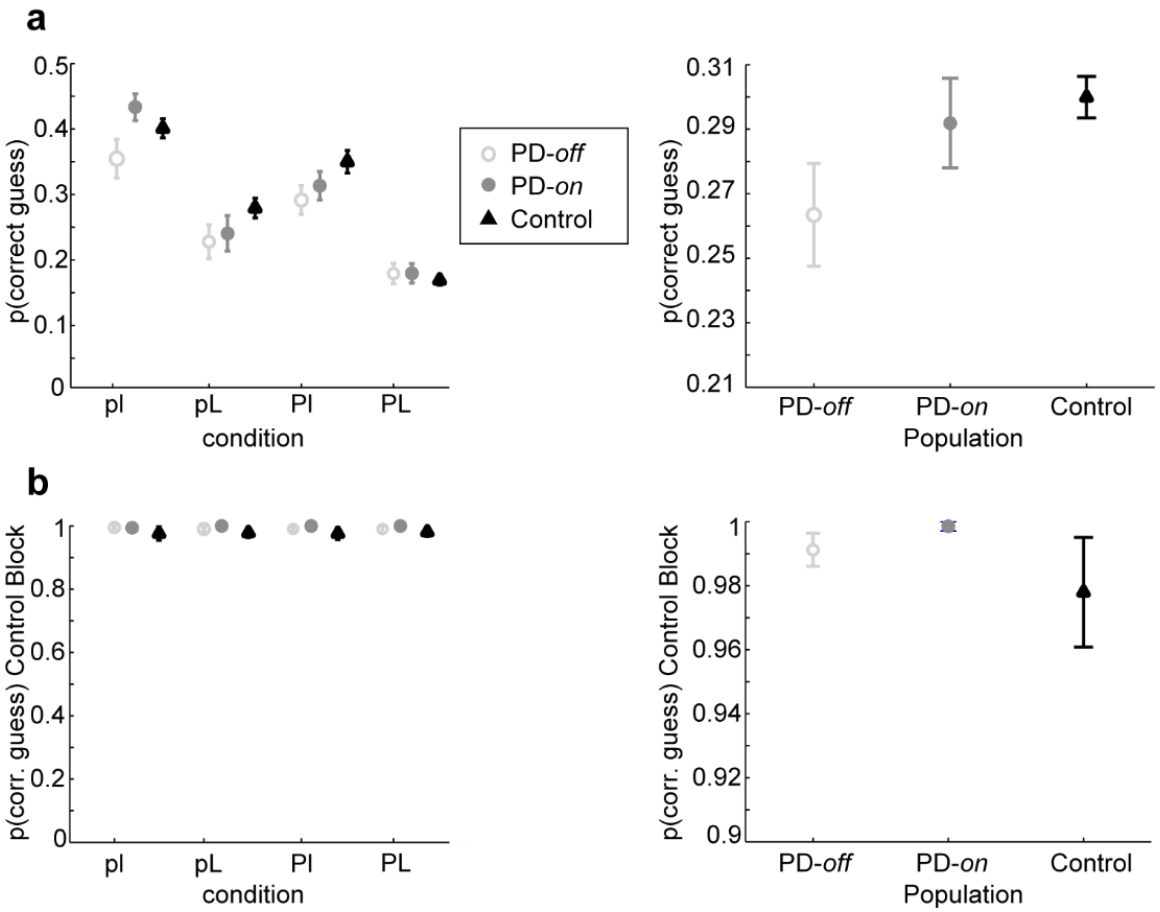

Supplementary Figure 6 - Performance in the task. a) Average proportion of trials in which participants, in the main task, correctly guessed the position of the target for PD patients off-medication (open light grey, ○), PD patients on-medication (filled dark grey, ●) and controls (black, ▲), separated by condition (left) or averaged across conditions (right). Within the PD population, there was a significant effect of medication state ($F_{1,101}$ = 4.7, *p-val* = 0.032), prior uncertainty ($F_{1,101}$ = 29, *p-val* < 10^{-6}) and likelihood uncertainty ($F_{1,101}$=108.5, *p-val*<10^{-6}), but no significant effect of session ($F_{1,101}$=1.6, *p-val* = 0.21, mixed-effects ANOVA with subject as a random effect and medication state, session, prior and likelihood uncertainty as factors), with performance being higher when participants were on-medication and when there was less uncertainty. **b)** Average proportion of trials in which participants, in the control task, correctly "guessed" the position of the target, i.e. participant moved the net to where the target was, separated by condition (left) or averaged across conditions (right). There was no difference between either PD patients off and on medication ($U_{8,7}$=20.5, *p-val* = 0.39, Mann-Whitney U-Test) or between PD patients and controls ($U_{15,15}$=121.5, *p-val* = 0.663, Mann-Whitney U-Test) in their performance in the control block, which was at ceiling. Note that, for visualization purposes, the graphics in the right do not start at zero. Error bars represent s.e.m..

Because PD patients off-medication differed from PD patients on-medication in their initial average reliance on current information (sensory weight), it is possible that their initial lower weights drove the difference in performance. However, as we saw above, the weights of PD patients off-medication were not further from the "statistically optimal weights" assuming the imposed variances (Supplementary Fig. 1). Also, for PD patients off-medication there was no significant correlation between their average weight value and their performance (spearman correlation r = 0.09, *p-val* = 0.75). For PD patients on-medication there was a significant correlation but it was negative (spearman correlation r = - 0.66, *p-val* = 0.01), suggesting that the higher weights in PD patients on-medication could even have hurt their performance. This both indicates that the lower weights in PD patients off-medication did not drive their lower performance, and also that PD patients on-medication have put more weight on current information despite it not improving (and indeed potentially hurting) their chances of catching the coin.

What, then, could be driving the performance gap? It could be a motoric or visual deficit, given that the task involved moving a "net" and searching through the visual space. However, and as we saw above, both PD patients on and off-medication performed at ceiling in the control block (Supplementary Fig. 6b). PD patients did not move the net less when they were off-medication ($F_{1,102}$=0.03, p-val = 0.86, [-0.007, 0.008] 95%CI, for main effect of medication, mixed-effects ANOVA with subject as a random effect and medication state, prior and likelihood uncertainty as factors), and the distance they moved was not correlated to performance (r = -0.22, *p-val* = 0.25). Together, this indicates that it was not a motoric or visual deficit driving the difference in performance.

Another possibility for the performance gap could be higher intra-individual variability. It has been proposed that PD patients have higher intra-individual behavioural variability, potentially related to losses in dopamine receptors, and that this could lead to differences in performance[14-17]. To obtain a measure of intra-individual variability, we calculated the trial-by-trial weight and analysed its corresponding variability (see Supplementary Methods for details). We found that PD patients had more variability in their trial-by-trial weights when off-medication ($F_{1,102}$=5.27, *p-val* = 0.024, [0.001, 0.012] 95%CI, mixed-effects ANOVA with subject as a random effect and medication state, prior and likelihood uncertainty as factors; see Supplementary Fig. 7a). Furthermore, this variability was strongly negatively correlated with performance across all participants (r = -0.78 spearman correlation, *p-val* < 10^{-5}, n = 30; see Supplementary Fig. 7b). If we consider only the PD patients, then this correlation becomes even stronger (r = -0.91 spearman correlation, *p-val* < 10^{-8}, n = 15; circles in Supplementary Fig. 7b). Our results indicate that PD patients, when off-medication, have higher intra-individual variability in the chosen weights, and suggest that this is the main reason for their lower performance off-medication.

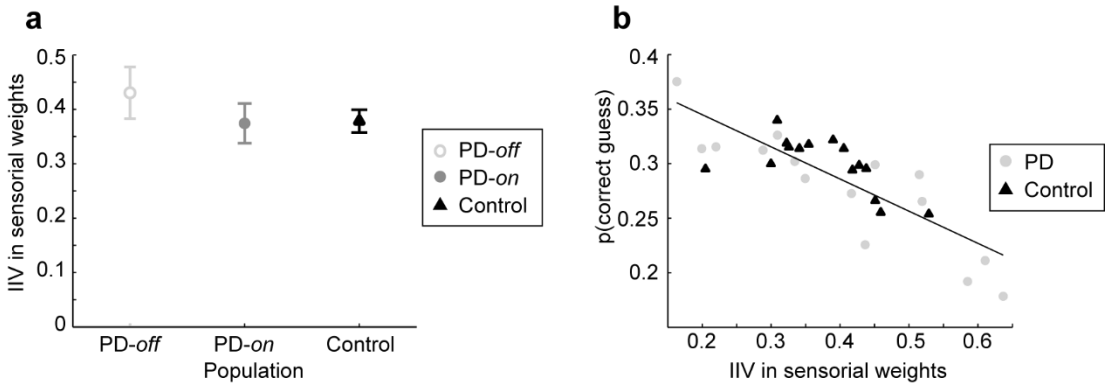

Supplementary Figure 7 – Intra-individual variability and its correlation with performance. a) Average Intra-individual variance (IIV) in the trial-by-trial weights for PD patients off-medication (open light grey,○), PD patients on-medication (filled dark grey,●) and controls (black,▲), averaged across conditions and sessions. PD patients had more variability in their trial-by-trial weights when off-medication ($F_{1,102}$=5.27, *p-val* = 0.024, mixed-effects ANOVA with subject as a random effect and medication state, prior and likelihood uncertainty as factors). **b)** Correlation between a participant's IIV in the trial-by-trial weights and their performance (proportion of correct guesses) in the main task. Represented are individual PD patients (light grey circle, ●) and Controls (black triangle,▲). The black line denotes a linear regression of a participant's performance as a function of their IIV in the trial-by-trial weights. The intra-individual variability in the sensory weights was strongly negatively correlated with performance across all participants (r = -0.78 spearman correlation, *p-val*<10^{-5}, n = 30).

SUPPLEMENTARY REFERENCES

1 Vilares, I., Howard, J. D., Fernandes, H. L., Gottfried, J. A. & Kording, K. P. Differential Representations of Prior and Likelihood Uncertainty in the Human Brain. *Current Biology* **22**, 1641-1648, doi:Doi 10.1016/J.Cub.2012.07.010 (2012).

2 Berniker, M., Voss, M. & Kording, K. Learning Priors for Bayesian Computations in the Nervous System. *Plos One* **5**, doi:ARTN e12686

10.1371/journal.pone.0012686 (2010).

3 Tomlinson, C. L. *et al.* Systematic Review of Levodopa Dose Equivalency Reporting in Parkinson's Disease. *Movement Disord* **25**, 2649-2653, doi:Doi 10.1002/Mds.23429 (2010).

4 Fabbrini, G., Juncos, J., Mouradian, M., Serrati, C. & Chase, T. Levodopa pharmacokinetic mechanisms and motor fluctuations in Parkinson's disease. *Ann Neurol* **21**, 370-376 (1987).

5 Thébault, J. J., Guillaume, M. & Levy, R. Tolerability, safety, pharmacodynamics, and pharmacokinetics of rasagiline: a potent, selective, and irreversible monoamine oxidase type B inhibitor. *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy* **24**, 1295-1305 (2004).

6 Wright, C. E., Sisson, T. L., Ichhpurani, A. K. & Peters, G. R. Steady-State Pharmacokinetic Properties of Pramipexole in Healthy Volunteers. *The Journal of Clinical Pharmacology* **37**, 520-525 (1997).

7 Kushida, C. A. Ropinirole for the treatment of restless legs syndrome. *Neuropsychiatric disease and treatment* **2**, 407 (2006).

8 Horadam, V. W. *et al.* Pharmacokinetics of amantadine hydrochloride in subjects with normal and impaired renal function. *Annals of internal medicine* **94**, 454-458 (1981).

9 Kording, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244-247, doi:10.1038/nature02169 (2004).

10 Fernandes, H. L., Stevenson, I. H., Vilares, I. & Kording, K. P. The Generalization of Prior Uncertainty during Reaching. *J Neurosci* **34**, 11470-11484, doi:10.1523/Jneurosci.3882-13.2014 (2014).

11 Behrens, T. E. J., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. S. Learning the value of information in an uncertain world. *Nat Neurosci* **10**, 1214-1221, doi:Doi 10.1038/Nn1954 (2007).

12 Castro, L. N. G., Hadjiosif, A. M., Hemphill, M. A. & Smith, M. A. Environmental consistency determines the rate of motor adaptation. *Current Biology* **24**, 1050-1061 (2014).

13 Burge, J., Ernst, M. O. & Banks, M. S. The statistical determinants of adaptation rate in human reaching. *Journal of Vision* **8**, 20-20, doi:10.1167/8.4.20 (2008).

14 MacDonald, S. W. S., Karlsson, S., Rieckmann, A., Nyberg, L. & Backman, L. Aging-Related Increases in Behavioral Variability: Relations to Losses of Dopamine D-1 Receptors. *J Neurosci* **32**, 8186-8191, doi:10.1523/Jneurosci.5474-11.2012 (2012).

15 MacDonald, S. W. S., Nyberg, L. & Backman, L. Intra-individual variability in behavior: links to brain structure, neurotransmission and neuronal activity. *Trends Neurosci* **29**, 474-480, doi:10.1016/j.tins.2006.06.011 (2006).

16 Li, S. C. & Lindenberger, U. Cross-level unification: A computational exploration of the link between deterioration of neurotransmitter systems and dedifferentiation of cognitive abilities in old age. *Cognitive Neuroscience of Memory*, 103-146 (1999).

17 Burton, C. L., Strauss, E., Hultsch, D. F., Moll, A. & Hunter, M. A. Intraindividual variability as a marker of neurological dysfunction: A comparison of Alzheimer's disease and Parkinson's disease. *J Clin Exp Neuropsyc* **28**, 67-83, doi:10.1080/13803390490918318 (2006).