Title: THE NEURONAL BASIS OF BIASES
Project number: No 117/18
Researchers: _Ruben Moreno Bote_____
Institution: _Universitat Pompeu Fabra_____
Project duration: starting 01/01/2019 ; finishing 31/03/2021
Keywords: decision-making, biases, neuronal population, brain

Abstract:
Goal-directed behavior often requires integrating current sensory information with the prior history of stimuli, actions and outcomes. However, if the past is uninformative for the task at hand, using prior history can induce disadvantageous behavioral biases. Here, we report the existence of history-dependent choice biases in highly-trained monkeys performing a motion direction discrimination task where only the current stimulus was behaviorally relevant. The observed biases fluctuated at two separate time scales: slow, spanning tens to hundreds of trials, and fast, involving variables from the previous trial. These biases were predictive of upcoming choices, and more so on trials with weak stimuli. We found that pre-stimulus activity of neuronal ensembles in prearcuate gyrus represented biases and was also predictive of the monkey's upcoming choices. We show that similar axes in the neural population state space represent both bias and choices. Further, biases are incorporated into the decision-making process as an offset of baseline activity along the decision axes that persists throughout the integration process of sensory evidence.

Background

A large body of research in psychology and behavioral economics has identified multiple biases in perception and decision-making, such as overestimating small probabilities and ignoring previous useful information (Kahneman, 2011). Very often these biases hinder performance. In other cases, however, biases do not substantially impair performance, as when they arise from simple heuristics that in turn reduce the cost of neuronal processing (Gigerenzer, 2008), or even improve behavior, as when allowing more accurate or efficient perceptual processes (Averbeck et al., 2006; Hanks et al., 2011; Hesselmann et al., 2008; Moreno-Bote et al., 2008, 2011). Despite the importance of understanding the nature and sources of biases to improve human performance (and to design artificial systems that do not show harmful biases), it is still largely unknown how they are represented in the brain and arise mechanistically.

Several past studies have investigated the nature of behavioral biases and their neuronal representation in perceptual decision making. Using a random-dot direction-discrimination task (Gold et al., 2008), it was found that monkeys develop intrinsic choice biases in a condition in which trials of either motion direction are randomly interleaved in equal proportions. In particular, a form of choice bias with a slow time scale was described, where monkeys' preference towards one of the options alternated over the course of hundreds of trials. However, the authors failed to find a neuronal representation of this intrinsic bias in either the motion-sensitive middle temporal area (MT) or in lateral intraparietal area (LIP), where transformation of sensory signals into saccadic eye movements is represented. Other studies have reported the representation of such biases in responses of single neurons in LIP (Eskandar and Assad, 1999; Shadlen and Newsome, 2001) but slow dynamics of choice biases were not found. Neuronal representations of choice biases have been reported in sensory areas. For example, Lueckmann and colleagues demonstrated a dependence of current and past choices and their representation in V2 (Lueckmann et al., 2018). Finally, in the posterior parietal cortex of rodents, optogenetic inactivation reduced some types of perceptual biases (Akrami et al., 2018).

A key brain region that could contribute to the representation of choice biases is the prefrontal cortex (Miller and Cohen, 2001). Indeed, it has been shown (Padoa-Schioppa, 2013) that in value-based decision-making, single neurons in monkey orbitofrontal cortex (OFC) modulate their activity as a function of upcoming choice, suggesting that those neurons hold a representation of intrinsic choice bias that could be used for the formation of the choice. Another recent study (Nogueira et al., 2017) has found

neuronal signals in rat OFC that are predictive of upcoming choices, even in cases in which this bias is not adaptive.

Within the prefrontal cortex, an area that holds promise for a neuronal representation of bias in tasks involving eye movements is the frontal-eye-fields (FEF) and neighboring pre-arcuate gyrus (area 8Ar, PAG). This area shows strong choice representation of impending oculomotor movements (Bruce and Goldberg, 1985; Kiani et al., 2015; Schall and Hanes, 1993), similar to LIP (Churchland et al., 2008; Roitman and Shadlen, 2002). Indeed, past studies have established that FEF and pre-arcuate gyrus carry strong choice related signals when decisions are reported with saccades (Kiani et al., 2015; Kim and Shadlen, 1999; Mante et al., 2013). Indeed, it is possible to predict choices and even changes of mind with very high accuracy based on neural responses before the saccade is made (Kiani et al., 2014). But little is known about if and how neurons in these areas represent choice biases.

All in all, it is largely unknown how biases are formed in the brain, whether bias and choice representations coexist, and how biases enter into, and affect, the decision-making process –for instance, by being represented in similar axes in the state space of neural population responses. Further, while all previous reports were based on single-cell analyses, addressing the similarity between biases and choice representations very likely requires the study of large pools of simultaneously recorded neurons due to the weakness of behavioral choice bias typically found in well-trained subjects.

Here, we report the existence of behavioral choice biases in highly trained monkeys performing a direction discrimination task where there is no stimulus imbalance or any other incentive for the monkey to develop such biased behavior. We found that choice biases are shaped by fast, single-trial history dependence and slow-timescale processes spanning hundreds of trials. Both biases were represented in the population activity of pre-arcuate gyrus neurons prior to the stimulus onset. We also found that pre-stimulus neural responses were predictive of the upcoming choice. Finally, we found that the neuronal representations of biases and choices were aligned in the state space, and that biases entered into the decision making process as an initial offset in the accumulation of evidence. Our results suggest that prearcuate gyrus represents the computation of prior history signals and its integration into the oculomotor decision-making process to guide behavior.

Aims

Study the dynamics and neuronal basis of sequential biases in decision-making.
Study the neuronal representation of biases in the prefrontal cortex of non-human primates.
Testing a sequential-processing theory of decision making and its link to biased behavior.

Method

Experimental Procedures

We will simultaneously record extracellular activity of populations of neurons in the prearcuate gyrus (PAG) of two macaque monkeys performing a direction discrimination task. All training, surgery, and recording procedures conform to the National Institutes of Health Guide for the Care and Use of Laboratory Animals and are approved by NYU's Animal Welfare Committee. Our study will also use 16 recording sessions from the same task and brain region that were used previously (Kiani et al., 2014; Kiani et al., 2015) in the context of different scientific question. The mixture of old datasets and the new ones will increase the number of tested animals and recorded neurons, boosting the statistical power of our analyses.

Behavioral Tasks

The direction discrimination task begins with the appearance of a fixation point (FP) at the center of the monitor. After the monkey fixates on the FP, two targets (T1 and T2) appear, followed by a patch of randomly moving dots. The dots are commonly presented at the center of the screen and do not overlap with the targets. The fraction of coherently moving dots (stimulus strength) in the direction of one of the

targets defines trial difficulty (Britten et al., 1992; Kiani et al., 2008). The motion direction and strength, given by signed coherence *C*, is chosen randomly on each trial from a predefined set (ref: Kiani). The stimulus is followed by a delay period of variable duration (>300ms) randomly selected on each trial. At the end of the delay period, the FP disappears (Go cue), signaling the monkey to report the perceived direction of motion by making a saccade towards the corresponding target and maintaining gaze on the target until the outcome of the trial is revealed. Correct choices are rewarded with a drop of juice. The gaze position is monitored (1KHz sampling rate) with a scleral search coil (CNC Engineering, Seattle) or an infrared camera (Eyelink, SR-Research, Canada).

Neural Recordings

While monkeys perform the direction discrimination task, we will simultaneously record extracellular activity of many PAG neurons using 96 channel microelectrode arrays (Blackrock Microsystems). The electrode array is implanted between the anterior bank of the concavity of the arcuate sulcus and the posterior tip of the principal sulcus. Neural signals are saved online at 30 kHz sampling rate. The spike waveforms are sorted offline manually (Plexon Inc., Dallas) or using custom-made, automatic sorting algorithms. We typically record 100+ units in each session (ref: Kiani). We use the term ‘‘units’’ to refer to both well-isolated single neurons and multiunits with poorer separation. Monkeys typically perform more than 1000 trials per session.

Measuring Biases

We will measure psychometric curves, defined as the fraction of trials in which the monkey chooses target T1 for a given stimulus coherence *C*. We will fit a logistic function for the psychometric curve to study the overall bias and the presence of lapses using maximum-likelihood fitting methods. For each session separately and for each trial *i*, we will compute the monkey's preference to choose target T1 over target T2, denoted PF, as the ratio between T1 choices vs the sum of T1 and T2 choices in a time window of one hundred trials. We will use PF as a measure of "slow bias." In addition to the slow bias, we define another variable, called "fast bias," which combines the effect of previous choice and outcome on the current choice. Fast timescale history dependencies longer than just the previous trial will be tested too.

We will use both the slow and fast biases in a logistic model to predict the upcoming choice. If the model can predict choice with an accuracy larger than 50%, we can claim that the biases play a role in the decision-making process. Preliminary results suggest that these two sources of bias together can predict the upcoming choice with a probability of 54%, above the 50% chance level, suggesting that monkeys' choices are influenced by past history, even when task design does not include such dependencies and penalizes such influences. We will measure the effect of biases on behavior by computing the fraction of reward lost compared to the ideal scenario in which there are not biases.

Neuronal Data Analysis

In the present study we focus on the activity of PAG units before the appearance of random dots. The representation of biases in neural responses could be most easily detected in this period because the stimulus has not yet been presented and thus neuronal signals are not "contaminated" by it. For each recorded unit, we compute the firing rate (the number of spikes per unit of time) per trial using an 800 ms window that terminates 10 ms before the dots onset. Because of the large number of recorded units and limited number of trials, we might need to reduce the dimensionality of the neuronal population (consisting of 100+ dimensions, each corresponding to the firing rate of a single neuron) by using principal component analysis (PCA) and taking the first components that in total explain at least 50% of the variance. Preliminary results show that this selection is appropriate, but it will be further tuned.

The critical analysis of our proposal corresponds to comparing the axes in the reduced neuronal space where biases and choices are represented. To determine the axis in which biases are encoded, we will use a linear regression analysis in which the bias (measured behaviorally as described above) will be fit to a linear combination of dimensionality-reduced neuronal activities. To determine the axis in which choice is encoded, we will use the same neuronal activity but this time to predict the upcoming choice

using logistic regression. The bias regression will be performed with the pre-stimulus activity to avoid contamination by stimulus-driven choice signals that become quite strong after dot-motion onset. The choice regression is applied both to the pre-stimulus and peri-stimulus activity.

The above analysis will result in the identification of choice and bias encoding axes dimensions in neuronal activity space at different times. These encoding dimensions will be compared using two sets of statistical tests: comparison of projections of one axis over the other in the state space, and comparison of choice prediction accuracy from each axis. Significance will be assessed using permutation tests, which are non-parametric and thus does not assume normality or any other response model (ref: Nogueira).

Finally, the results of the previous analysis will be contrasted with models in which the recorded neuronal population responses are simulated and various degrees of bias and contribution to the formation of choice are added to the simulations. These simulated models will provide a foundation for interpreting our experimental observations and analyses.

Results

Using a 96-channel multi-electrode array, we recorded neural population activity from the prearcuate gyrus (area 8Ar, PAG), while monkeys (n = 2) performed a direction discrimination task (16 sessions; (Kiani et al., 2014, 2015)). Each trial began with the monkey fixating on a central fixation point on the screen, following by the appearance of two targets, and a circular patch of random dot kinematogram (Britten et al., 1992). The percentage of coherently moving dots (coherence or motion strength) and the net motion direction of motion varied randomly trial to trial. The motion stimulus was shown for 800 ms and was followed by a delay period. The monkey reported motion direction at the end of the delay period with a saccadic eye movement to the corresponding targets. We use signed motion coherence (Britten et al., 1992; Kiani et al., 2008; Shadlen and Newsome, 2001) to jointly represent the stimulus strength and direction with a single variable (positive for motion toward T1 and negative for motion toward T2).

Highly trained monkeys exhibit slow and fast choice fluctuations

Monkeys were extensively trained in the task and showed stable performance prior to neural recordings. Despite their extensive training and trial-to-trial independence of stimulus conditions, both monkeys demonstrated slow fluctuations in their choice preference, where monkeys chose one target more frequently than the other for tens to hundreds of trials before reversing their preference. These fluctuations were spontaneous and could not be explained by fluctuations of motion direction across trials because motion direction was largely balanced in those periods. To further ensure that the slow choice preference fluctuation did not merely reflect random fluctuations arising from spurious unbalance of motion directions and coherence, we repeated our analysis and replicated our results by subsampling trials to equalize the number of trials with stimuli moving toward T1 or T2 for each coherence. Across sessions, the correlation coefficient between slow choice preference fluctuations and signed motion coherence was weak and not statistically significant (mean ± SEM 0.02 ± 0.04, permutation test p-value = 0.26). Additionally, the mean auto-correlogram of slow choice preference fluctuations calculated on coherence-balanced trial history showed a statistically significant broad central peak, indicating the presence of slow fluctuations of response preference, irrespective of fluctuations of stimulus statistics.

On a finer timescale, we also observed that the monkey's choices were influenced by recent choices and outcomes. The monkey tended to choose the opposite target after error trials in this session. Fluctuations of both slow and fast choice preference were reflected as a shift in the psychometric curve when it was calculated conditioned on the direction of slow choice preference, previous choice, or recent outcome.

Improvement of choice prediction accuracy with slow and fast choice fluctuations

To quantify how monkeys' decisions were affected by the slow and fast fluctuations of choice preference, we measured whether and how much they would improve the prediction accuracy of upcoming choice beyond that given by motion stimulus alone.

We built a logistic regression model to predict choices based on three variables: stimulus coherence, slow choice preference fluctuation, and fast choice preference fluctuations expressed as a combination of previous choice and reward. The cross-validated model prediction accuracy was assessed using a *leave–one–out* procedure. To measure the role of choice preference fluctuations in determining choices, we compared the prediction accuracy with a baseline obtained by fitting the model using shuffled fast and slow choice preferences across trials. Because shuffling destroys the statistical relationship between current choices and choice preference fluctuations, the comparison with the baseline isolates improvement of the prediction accuracy conferred by the latter. The mean prediction accuracy improvement across all sessions was small but significantly larger than zero (0.016 ± 0.002, one sample *t*-test, *p*-value = 10^{-6}). Importantly, when we focused only on difficult trials, where the stimulus is less informative and biases could have a larger influence on choice, the improvement of choice prediction accuracy doubled (mean ± SEM: 0.032 ± 0.004; one sample *t*-test, *p*-value = 10^{-7}). Consistently, the improvement tested only on easy trials was not different from zero (mean ± SEM: 0.0001 ± 0.0002; one sample *t*-test test, *p*-value = 0.69), which is expected because prediction accuracies based on stimulus strength alone are already close to ceiling. Overall, biases had a tangible effect on upcoming choices especially for more difficult decisions.

It is possible to use similar choice prediction models to quantify the temporal extent of fast and slow choice preference fluctuations. For fast choice preference fluctuations, information about choice and reward of the previous trial significantly improved choice prediction accuracy. However, including information about more distant past (from two to five trials back) did no improve the prediction accuracy further (paired *t*-test on prediction accuracies, *p*-value > 0.33), indicating that in our task, fast choice fluctuations were shaped in a time scale that was not longer than one single trial in the past.

We studied how the size of the trials window used to calculate slow choice preference improved performance of a model that included both slow and fast choice preference fluctuations compared to a model in which only fast choice preference fluctuations was used (in both models coherence was also used as a regressor). If slow fluctuations were defined using trial windows of less than 130 trials, there was not statistically measurable effect. However, when slow fluctuations were calculated in larger windows (130 - 400 trials), there was a statistically significant increase in prediction accuracy. Therefore, we conclude that slow choice preference fluctuations estimated in a window of ~130 trials and fast choice fluctuations capturing the immediately preceding choice and reward were sufficient to explain the behavioral biases observed in our experiment.

An important question is to know the relative strength of the effect of fast and slow choice preference fluctuations. To compare their effects, we first expressed both types of fluctuations in the same units (log-odd units) by using the logistic regression model described above. Thus, "slow bias" was defined as the product of slow choice preference fluctuation and its corresponding weight in the model plus the model offset. Similarly, "fast bias" was defined as the product of fast choice preference fluctuation and its model weights. To compare the relative strength of slow and fast biases on choice prediction, we calculated the mean ratio of the absolute value of each divided by the sums of absolute values of both biases and stimulus strength in log-odds space. We found that both fast and slow biases were effective in shaping the choice but, on average, the fast bias had a lower impact on choice (mean ± SEM 0.18 ± 0.01) compared to the slow bias (mean ± SEM 0.25 ± 0.01; difference of fast and slow = $- 0.07 \pm 0.016$, one sample *t*-test *p*-value = 0.006).

To capture the overall effect of slow and fast biases on behavior, we defined "total bias" on each trial as the sum of slow and fast biases. On average, the contribution of total bias to the monkey's choices was 0.25 ± 0.01 in units of log-odds. For comparison, it was more than twice smaller than the contribution of motion stimulus alone (mean ± SEM 0.58 ± 0.02), reflecting the fact that monkeys made their decision largely based on the presented stimuli but they were also slightly impacted by their current bias.

Since our task was designed such that the stimulus sequence across trials did not have temporal correlations, the existence of behavioral biases described above can only impair the monkey's performance. As expected, periods with larger than average total bias correlated with periods with lower than average animal accuracy (percent of correct choices; Pearson correlation coefficient -0.08 ± 0.03,

permutation test, p-value = 0.001). However, the reduction in accuracy was very small (0.007 ± 0.003%; one sample t-test, p-value = 0.02), suggesting that monkeys may not have noticed the adverse effect of biases on their performance or they did not find enough incentive to fully abolish them.

Representation of bias in the pre-stimulus neural responses of PAG

Since the slow and fast biases influenced the monkey's choice, the information about them should be present in brain regions involved in the decision-making process. An area of interest could be PAG, where neural responses related to the accumulation of evidence have been found (Kiani et al., 2014; Kim and Shadlen, 1999; Mante et al., 2013). Also, since the observed behavioral biases had a history dependent component, there should be neurons sensitive to the bias even prior to the stimulus presentation.

To investigate whether these biases were represented in the responses of PAG neural population, we used a linear model in which each type of bias (slow, fast, or total) was regressed against pre-stimulus activity of simultaneously recorded PAG neurons. The analyses were performed on the top variance-predictive PCA components of the neural population responses that collectively explained 50% of the variance. Here, we used PCA to reduce overfitting (Mante et al., 2013; Yu et al., 2009), but qualitatively similar results were obtained without dimensionality reduction. All three biases were significantly represented in the pre-stimulus responses of PAG population.

Fitting a logistic regression model to the firing rate of the population of simultaneously recorded PAG neurons revealed that previous choices could be decoded up to three trials back in the past. A similar model could predict the outcome (reward) of the preceding trial, but not further back.

Predicting choices from pre-stimulus neural responses of PAG

Given that slow and fast biases influenced monkeys' decisions and that they were represented in the pre-stimulus PAG activity, we asked whether pre-stimulus activity was also predictive of monkeys' upcoming choices. To quantify this activity modulation at the population level, we fit a logistic regression model to predict upcoming choices using the PCA-dimensionality-reduced PAG population responses in the 800 ms before stimulus onset. For many sessions (44%; 7 out of 16), the cross–validated prediction accuracy was above chance level (0.55). Similar to the behavioral model, the cross-validated prediction accuracy was higher when we focused on difficult trials (white bars; 0.54 ± 0.008 permutation test, p-value = 0.001), and not significantly different from chance level for easy trials.

Prediction of choices based on pre-stimulus neural activity is due to the representation of slow and fast biases

A key question is whether choice predictive neural responses prior to stimulus onset are due to the representation of the fast and slow biases that we have defined behaviorally. One possibility is that encoding of our behaviorally-defined biases fully explains the representation of choice prior to the stimulus onset, which would imply that total bias and choice representations are "aligned" in neuronal activity space. Alternatively, choice predictive neural responses could arise from factors not fully captured by fast and slow biases, which would cause misalignments between choice and total bias representations. To differentiate these two possibilities, we asked if the neural representation of biases was as predictive of behavior as the neural representation of choice.

For each trial, we used the remaining trials in the session to find the best hyperplanes that explained the choice and total bias based on pre-stimulus responses. Then, we calculated the distance of pre-stimulus responses of the left-out trial from those two hyperplanes (d_{choice} and d_{bias}). If the pre-stimulus choice and bias representations were aligned, predicting the upcoming choice based on d_{bias} would be as accurate as using both d_{choice} and d_{bias}. Indeed, this was what we observed. Across sessions, the difference in predicted accuracy was negligible and not significant (paired t-test, p-value = 0.42), suggesting that the neural representation of total bias was functionally indistinguishable from the neural representation of choice prior to stimulus onset. Consistent with these results, predicting the choice based on the sign of d_{bias} or d_{choice} was comparable, with slightly better accuracies for d_{bias}, further supporting our conclusion.

Additional insight about choice predictive neural responses prior to stimulus onset is gained from comparing the geometry of choice and bias decoder hyperplanes. Vectors in high-dimensional spaces tend to be orthogonal (Hall et al., 2005). However, if the representation of choice and bias are functionally aligned, one would expect that the angle between the norms of their respective hyperplanes is less than 90 deg. In fact, we found that the weight vectors that defined the norm of the choice and bias hyperplanes were positively correlated and the correlation coefficients were significant for the majority of sessions, supporting the notion that the hyperplane norms were not orthogonal.

Further supporting the alignment hypothesis, we found that sessions with stronger representation of the total bias (higher cross-validated R^2) in pre-stimulus activity of PAG also had a higher cross-validated choice prediction accuracy. These results suggest that across-session variability in choice predictive power could be explained by the across-session variability in the representation of total bias.

The integration of bias into the accumulation of evidence

Given the presence of the bias signal prior to stimulus presentation, the question arises of whether and how this bias impacts the accumulation of the sensory evidence during the stimulus-viewing period. Our result about the alignment of bias and choice decoders before stimulus onset suggests that the bias could be implemented as an initial offset in baseline activity before accumulation of sensory information begins in each trial. This initial offset is best visible in the neuronal activity axes where the decision is encoded (orthogonal to the choice decoder hyperplane). Thus, we plotted how the decision variable (projection of neuronal activity onto the choice axes) evolved over time.

We found, consistent with our expectations, that there was an offset in the decision variable before stimulus onset. The offset was positive when the bias favored the final choice and negative when the bias was against the final choice (red). This offset was roughly constant for the whole duration before stimulus onset and persisted during the first few hundreds of milliseconds of stimulus viewing period, when the decision was formed. Interestingly, the offset lasted longer for more difficult stimuli, where monkeys integrated the sensory evidence longer. However, toward the end of the stimulus presentation, the offset vanished because the monkey had likely reached a decision on the majority of trials. Our results suggest that the initial bias was integrated into the decision-making process and contributed to the formation of the choice.

The initial offset and final convergence of the decision variables for positive and negative (red line) biases are consistent with predictions of bounded evidence-accumulation models for the decision-making process (Link, 1992; Ratcliff and Smith, 2004; Shadlen and Kiani, 2013). The gradual dynamics of the decision-variables and larger and longer-lasting effects of offset for weaker motion stimuli is compatible with these models too. This is because the accumulation of sensory evidences is slower for weaker stimuli, and the decision variable takes longer to hit the decision bound. This slower ramping lets the bias-induced offset (blue and red) survive longer during stimulus viewing. In contrast, sensory evidence accumulates quickly on easy trials, causing accelerated convergence of the decision variables and leaving minimal room for the bias-induced offset to influence the final decision.

Another interesting feature of the neurally inferred decision variable is that it continues to rise after the convergence of positive and negative bias. Because our model is designed to predict the choice, the magnitude of the model decision variable is influenced by any factor that improves its accuracy. Those include neural responses that may represent only the final choice but not necessarily the decision-making process that leads to the choice. Such choice-related responses have been shown to emerge in motor-planning regions toward the end of the stimulus viewing period in the dots task (Peixoto et al., 2018), and could be responsible for additional rise of the model decision variable after 700ms from stimulus onset, when the initial bias is no longer represented.

Conclusions

We have studied the dynamics and neuronal representation of biases in highly-trained monkeys performing a perceptual decision-making task while recording simultaneous responses of hundreds of

neurons in the prefrontal cortex. Despite trial-by-trial independence of the stimulus direction, monkeys exhibited weak but measurable behavioral biases. Observed biases emerged at two distinct time scales. Slow bias, reflecting the monkey's preference towards one of the targets, fluctuated at a time scale spanning tens to hundreds of trials, whereas fast bias was shaped by the choice and outcome of only the preceding trial. Using these biases improved prediction of upcoming choices beyond what could be achieved based on the stimulus alone, and this improvement was most evident on trials with weak stimuli. Further, we found that pre-stimulus responses of neuronal population in prearcuate gyrus represented the fast and slow biases and were also predictive of the monkey's upcoming choices. Critically, the axes that represented bias and choice in the neural population state space were similar, indicating that choice representation prior to stimulus onset was captured by the representation of the fast and slow biases. We also demonstrated that those biases were incorporated into the decision-making process as an offset of baseline activity that persisted throughout the integration of sensory evidence during viewing of the motion stimulus.

Optimal decisions should be based on all available information. In Bayesian decision theory, total information for sensory choices depends on both the stimulus and prior expectations (for review see (Summerfield and de Lange, 2014)). Such prior expectations might reflect for example the learned statistics of the environment and as such facilitate correct decisions especially in when sensory evidence is ambiguous (Hanks et al., 2011; Rao et al., 2012). However, when objective prior expectations are balanced and do not confer useful knowledge into the current task, fluctuations of subjective priors constructed based on the most recent experiences could lead to behavioral biases that change with time and could reduce accuracy. Since our task was designed to make locally-inferred prior uninformative, the fluctuating biases observed in our well-trained subjects are not learned mechanisms for reward optimization. Rather, they illustrate native mechanisms that shape history-dependent biases.

To investigate how such prior expectations arise on the neuronal level, several past studies introduced biases in the decision-making process by manipulating the frequency of stimuli associated with different choices (Abrahamyan et al., 2016; Hanks et al., 2011; Hermoso-Mendizabal et al., 2019; Rao et al., 2012). For example, Hanks et al (Hanks et al., 2011) trained humans and monkeys to perform a direction discrimination task with random dots stimuli, but instead of interleaving randomly equal number of trials moving in either direction, trials in which the stimulus moved in one direction were more frequent than in the opposite direction (for instance, the stimulus in 80% of the trials moved towards 'right' and 20% moved towards 'left'). In this condition, subjects adapted their strategy toward selecting the more frequent motion direction, thus developing a well-adapted behavioral bias that maximized reward outcome. The authors found that this stimulus bias was represented in the single neuron activity in the lateral intraparietal (LIP) area. This neuronal representation consisted in part of a pre-stimulus response offset that was positive for neurons that favored choices in the biased direction and negative in the opposite. Similar response offsets have also been reported in value-based decision-making in LIP, where initial neural responses are elevated when a target yields larger or more frequent rewards is in the neural response field (Platt and Glimcher, 1999; Rorie et al., 2010; Seo et al., 2009). Although these past studies revealed neuronal mechanism for the encoding of stimulus biases, it left open the question of whether choice biases spontaneously arise even in the absence of stimulus biases. In such condition the history of previous stimuli, choices or outcomes is uninformative and potential priors built up based on it can be detrimental.

Indeed, the impact of previous choice, outcome and stimulus on the upcoming choice has been extensively observed in behavioral experiments on humans, monkeys or rodents (Abrahamyan et al., 2016; Akaishi et al., 2014; Akrami et al., 2018; Busse et al., 2011; Fischer and Whitney, 2014; Hwang et al., 2017; Jasper et al., 2019; Nogueira et al., 2017). For example Abrahamyan and colleagues has reported that human subjects tend to repeat previous choices or, switch to the opposite target depending on the previous trial outcome (Abrahamyan et al., 2016). Similarly, the tendency to repeat the previous choice, especially in the trials with weak sensory evidence was reported for the task without feedback (Akaishi et al., 2014). Moreover, previous outcome (Busse et al., 2011) or stimulus (Akrami et al., 2018; Fischer and Whitney, 2014) also impacted current decisions. As it has been suggested, such harmful biases might

be a consequence of, a priori adaptive, mechanisms that take into account the relatively stable statistics of the natural environment. In such steady conditions the present stimulus likely will resemble the one occurring recently (Summerfield and de Lange, 2014)

Past history is not the unique source of biases. Several studies have also reported a behavioral bias with a much slower pace. Busse and colleagues reported preference towards one of the targets slowly changing across sessions and persisting across days. However according to authors such bias not necessarily originated spontaneously but might reflected minor imbalance in a reward size on one of the sides (Busse et al., 2011). Another type of slow bias was described in work by Gold and colleagues (Gold et al., 2008). This bias is similar to the slow bias that we have described in our study, but in their case it declined rapidly across sessions, probably reflecting training such that biases slowly disappear as the animal became more proficient in the task. However, the opposite trend, namely the monotonic increase of serial biases was observed in multiple data sets from behavioral experiments in which human subjects were engaged in delayed estimation of color (Barbosa and Compte, 2019). These findings lever the adaptive mechanism of serial biases or learning process as a source for some types of behavioral biases. In contrast, the slow bias reported in our current study did not have a monotonic build up or decay in most sessions, but rather fluctuated without any apparent trend throughout the experimental session.

While much work has been devoted to the behavioral characterization of biases, very little is known regarding their neuronal underpinnings (Akrami et al., 2018; Eskandar and Assad, 1999; Nogueira et al., 2017; Shadlen and Newsome, 2001). By recording the activity of hundreds of neurons, we have described how biases and choices are represented in neuronal population activity. We have found that neuronal activity in PAG correlates with the behaviorally-defined biases, and that both biases and choices are represented along similar axes in neuronal activity space. This result is not a necessary outcome of our analysis, as it is possible that additional sources of predictability were encoded in individual trial activity that do not correlate with behavioral variables of preceding or succeeding trials. Our study also demonstrates that bias signals seem to be directly incorporated into the decision process as an initial offset before sensory accumulation, in a similar axis where choice accumulation takes place. Although several theoretical studies have proposed this neuronal implementation (Drugowitsch et al., 2019; Hanks et al., 2011), it has been hard to experimentally test this prediction because of two main reasons: first, it is important to fully characterize the biases that are represented in the neuronal area under study and avoid missing both their fast and slow components; and, second, it is important to define the choice axes in neuronal population activity where sensory accumulation takes place to test whether biases signals project onto that or they are orthogonal, a question that is more naturally addressed by using neuronal ensemble analysis. By combining detailed behavioral analysis with neuronal ensemble analysis, we have been able to link these two levels of description and provide evidence leaning towards the bias-offset hypothesis.

To conclude, we have provided neuronal evidence that behavioral biases and choices are represented in the same neuronal circuits along similar directions of activity state space. The implications of these results can be multifarious. For instance, in a speculative vein, the fact that biases are directly incorporated into the decision process as an offset, just as veridical information would do, could speak about why it is so difficult to eliminate deleterious biases from our daily life behavior (Gigerenzer, 2008; Kahneman, 2011), and it is in line with current work on decision making proposing bottlenecks in sensory (Moreno-Bote et al., 2014) and value-based processing (Hayden and Moreno-Bote, 2018).

Published work resulting from the project supported with the present grant separated in abstracts (proceedings and meetings), full papers and contributions included in books:

Full papers:

Prefrontal cortex represents heuristics that shape choice bias and its integration into future behaviour, G Mochol, R Kiani, R Moreno-Bote, Current Biology, 31, 1-11, 2020

Sense of Agency for Mental Actions: Insights from a Belief-Based Action-Effect Paradigm, E Lopez-Sola, XD Arsiwalla, R Moreno-Bote, bioRxiv, 2020