Title: Dynamic task-belief is an integral part of decision-making

Authors: Cheng Xue, Lily E. Kramer, Marlene R. Cohen

Affiliations: Department of Neuroscience and Center for Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, PA

Abstract: Unlike in laboratory settings, natural decisions are often made under uncertain beliefs about task demands. To quantify the unexplored dynamics between task-belief and decisions, we trained macaque monkeys to make perceptual discriminations under implicitly evolving task rules. By analyzing task- and perception-related signals from simultaneously recorded neuronal populations in cortical areas 7a and V1, we demonstrated that fluctuating task-belief and perceptual decision-making are inextricably linked. Stronger task-belief is correlated with better perception, and in turn, response fluctuations in visual neurons affect task-belief changes. Our results demonstrate that combining large-scale inter-area recordings with rigorously controlled complex, realistic behaviors can bring new understanding of the relationship between cognition and behavior in health and disease.

One Sentence Summary: Reciprocal dependence of fluctuations in task-belief and perception explains behavior

Main Text: Humans and animals make countless decisions every day that affect their well-being or even survival. In the laboratory, decision-making has typically been studied by observing behaviors and neuronal activity while subjects perform simple, well-understood sensory-motor integration tasks (*1-3*). But real-life decisions usually need to contend with a more important problem even before making perceptual judgements: inferring the relevant task to solve in a certain situation (i.e., task-belief). Task-beliefs allow decision-makers to focus on a relevant subset of the huge amount of information in natural environments, and beliefs are flexibly adapted belief as the environment evolves. Flexibly adapting task-belief is critical but difficult: the inability to appropriately respond to changing conditions is a debilitating symptom of disorders including autism, dementia, and substance abuse (*4-6*).

Typically, task-belief is assumed to be a separate functional module that occurs before, and independent of, perceptual decision-making. In this view, the belief module (possibly involving parietal, prefrontal, and cingulate cortical areas (*7-11*)) identifies the relevant task and then the perception module (involving sensory areas such as visual cortex) performs perceptual judgements on the chosen task (*10, 12, 13*). However, recent studies suggest that even with experiments' best attempts to keep task-belief constant (with fixed stimuli, explicit instructions, and task statistics), internal belief states still have uncontrolled fluctuations (*14-17*), some with effects on visual cortical activity and perceptual performance that is supposedly confined to the perception module (*16-18*). These results suggest that beliefs and perception interact in complex ways. The biggest barrier to understanding such interactions is estimating task-belief during each decision, which is by definition internal and continually changing.

To address this challenge, we devised a novel two-feature discrimination task to assess perception and belief simultaneously. We trained animals to discriminate either the spatial location or spatial frequency of two Gabor patches presented in series. The animals indicated both the subjective belief about which feature was task-relevant and the corresponding perceptual judgment by making a saccade to one of four targets (Figure 1A, upper left panel). They were rewarded only when both task-belief and perception were correct (Figure 1A, right panel). The relevant feature was not cued and switched with a low probability from one trial to the next (Figure 1A, lower left panel). This design provides rich and easy measurements and

manipulations of subjects' behavior during dynamic belief-based decision-making. Meanwhile, we recorded from groups of neurons from which we could decode information about both visual features the monkeys discriminated (in visual cortical area V1) and task-belief (in parietal area 7a (*9, 19*)) (Figure 1B). Together, these measurements provide a unique window into belief updating and perceptual decision-making on every trial.

After training, the animals successfully discriminated the feature change they believed to be relevant, and largely ignored the feature believed to be irrelevant (Figure 1C). The animals also effectively updated their belief according to the evolving task requirements, switching tasks only a couple of trials after the (uncued) task changes occurred (Figure 1D). The number of trials the animals took to notice task changes was close to optimal given their perceptual sensitivity (Figure 1E).

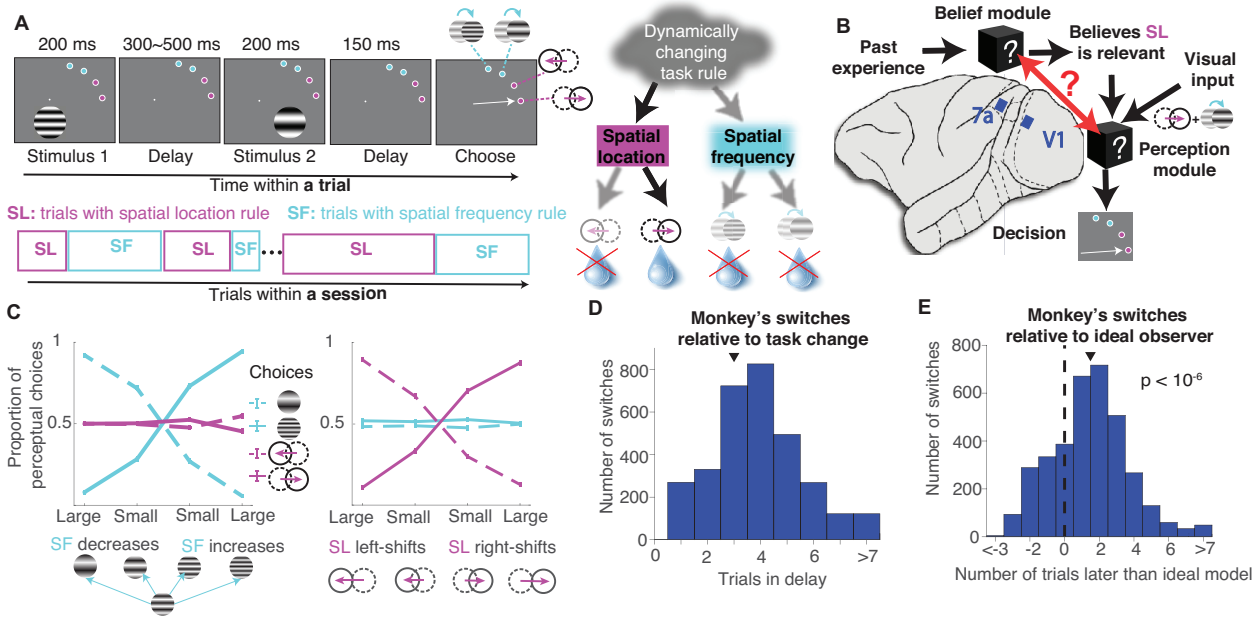

Fig. 1 Behavioral paradigm and electrophysiological recording. (**A**) Schematic of the two-interval, two-feature discrimination task with stochastic task switching. On each trial, monkeys discriminate the difference in either spatial frequency or spatial location between two subsequent Gabor stimuli; and are rewarded for correctly reporting the sign of the change in the relevant feature. The relevant feature is uncued and changes with 2.5% probability on each trial. The monkeys indicate their perceptual decision and the feature believed to be relevant by making a saccade to one of four choice targets. (**B**) Belief-based decisions could potentially be solved by independent hierarchical modules that compute belief and perception (black boxes). We simultaneously recorded population activity from one representative brain region for each module (7a and V1 respectively, blue squares show approximate implant locations) to test the hypothesis that these modules are non-independent (red arrow). (**C**) Psychometric curves showing the monkeys' perceptual choice proportion as a function of spatial frequency (left panel) and spatial location (right panel) differences. The flat curves for the irrelevant feature show that animals successfully ignored irrelevant visual information. (**D**) Distribution of number of trials it took the monkeys to adapt to task changes across experimental sessions (mean 3.1 trials). (**E**) Distribution of number of trials the monkey took to adapt to the task change relative to an ideal observer model (see Supplementary Methods). Positive values refer to

occasions where monkeys were slower than the model; negative values indicate that the monkeys were accidentally faster (mean 1.5 trials).

Our behavioral results demonstrate that dynamic task-belief strength affect the accuracy of perceptual decision-making. By design, the animals' perceptual choices are informed by stimulus information within the trial and should ideally be independent from trial history. However, across experimental sessions with same stimulus conditions, the animals had better perceptual performance (i.e., perceptual accuracy of whichever task the animal chose to perform) after rewarded trials (which reinforced task-belief) than after unrewarded trials (which introduce uncertainty to the monkey's belief state) (Figure 2A, Figure S1A). Correspondingly, information coding in V1 was also better after rewarded than unrewarded trials (Figure S1B).

These results suggest that strong beliefs are associated with better information coding in V1. To test this hypothesis on a trial-by-trial basis, we leveraged the fact that neuronal populations in parietal area 7a encode belief. We decoded this continuous measure of the animals' belief on each trial (Figure 2B). Consistent with the idea that rewards reinforce beliefs, the animals' task choice was better classifiable after a rewarded than an unrewarded trial (Figure S2A). Decisions to switch tasks were associated with a dynamic change in decoded task-belief away from the old task and toward the new task (Figure 2C).

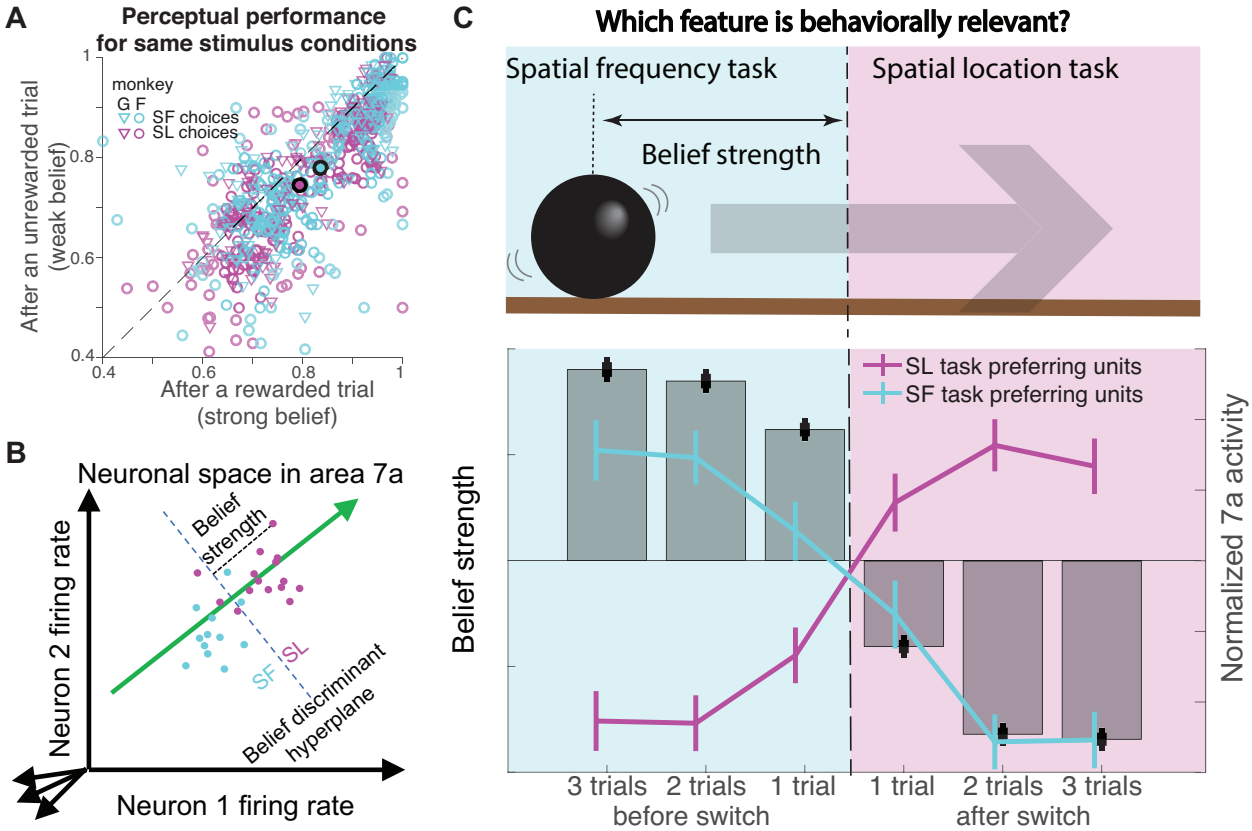

Fig. 2 Behavioral effect and neuronal measure of belief strength. (**A**) Perceptual performance is better following rewarded trials (abscissa) than unrewarded trials (ordinate). Each point represents one stimulus condition of an experimental session, and we compute perceptual performance for whichever task the monkey chose, regardless of whether that task-belief was correct. The distribution lies significantly below the unity line (all $p<10^{-6}$ for both monkeys and both features), showing lower perceptual performances following a non-rewarded trial than following a rewarded trial, with the

same perceptual difficulty. (**B**) In a high dimensional neuronal space expanded by the activity of 7a units during the delay period, we find the best hyperplane to discriminate the task the animal performed on the trial. We define our single-trial neuronal measure of belief strength as the Euclidean distance from 7a population activity on each trial to the hyperplane. (**C**) Belief strength is schematized as the distance from a rolling ball to a boundary. for trials leading up to the animals' decision to switch tasks, the average belief strength decreased monotonically, changed sign right at the point the monkey decided to switch tasks and recovered as the new task-belief was reinforced (histograms in bottom panel). Normalized activity of task-selective 7a units tracked the same dynamics as decoded belief around task switches (lines in bottom panel). Error bars indicate standard errors.

Likewise, we also estimated trial-by-trial feature discriminability using V1 population responses in the corresponding feature encoding dimensions (Figure 3A). As expected, with the same stimuli (i.e. same difficulty), trials with larger relevant feature discriminabilities yield better perceptual performances (Figure S3A-B). For each task-belief and stimulus condition, we look for potential correlation between belief strength measured from area 7a and perceptual discriminability measured in V1 (Figure 3B). Despite the fact that the resulting correlation is based on few trials and only a few dozen neurons across two very weakly connected areas (*20*), there is a positive correlation between belief and the encoding of the feature that is believed to be relevant, but not when the feature is believed to be irrelevant (Figure 3C-D). Together, our results indicate that belief-based decision making is an integrated system rather than a separable two-stage computation (first the categorical task-belief, then the corresponding perceptual decision).

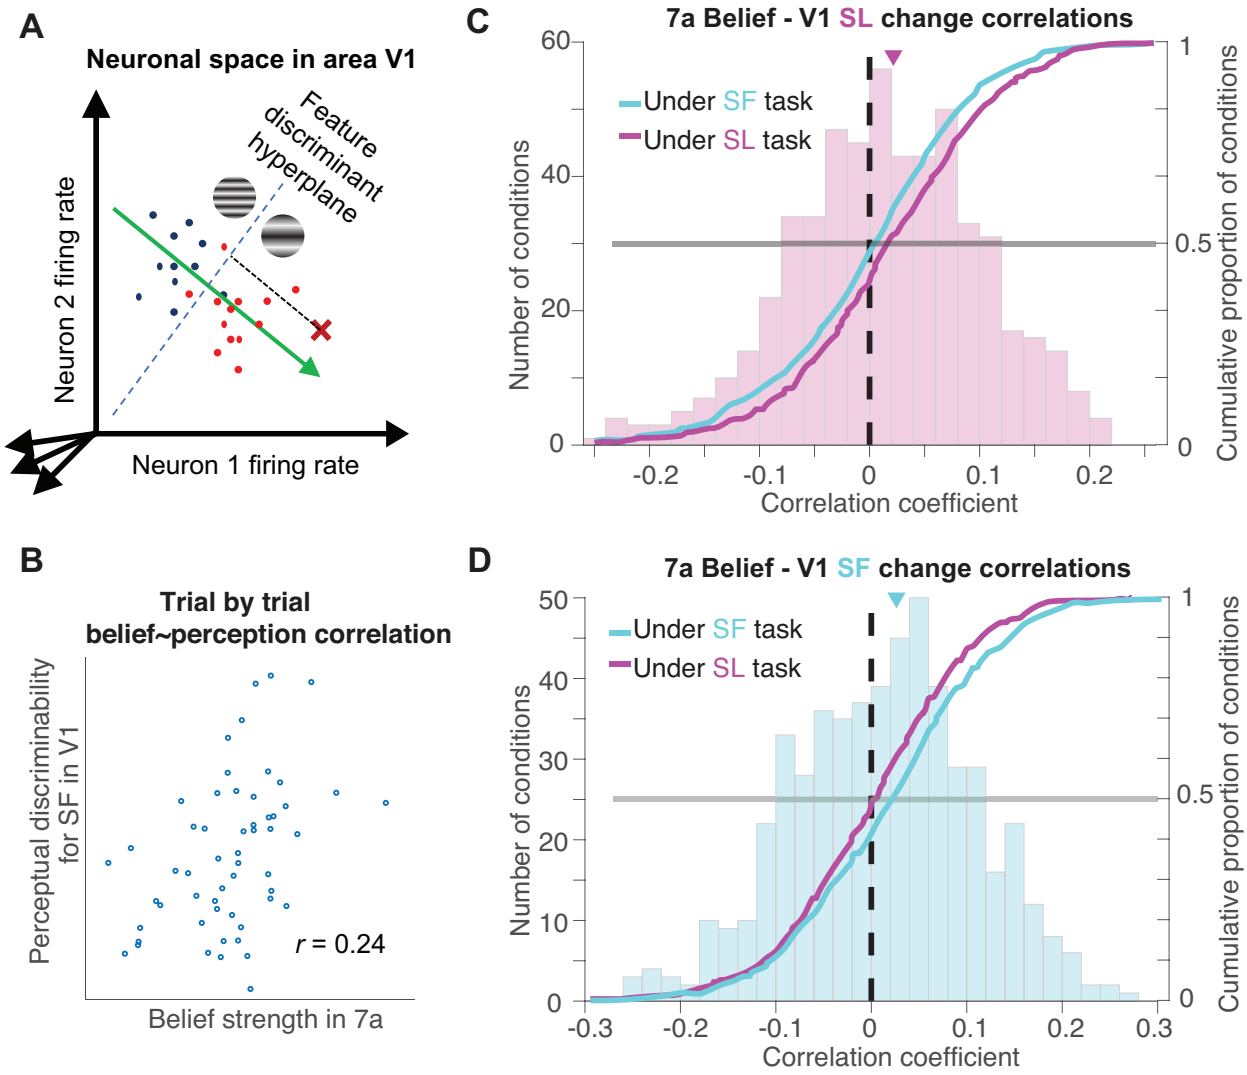

Fig. 3 Belief and perception are linked on a trial-by-trial basis. (**A**) Using a procedure similar to that described in Figure 2B, we define the perceptual discriminability of each stimulus feature change on each trial as the Euclidean distance from V1 population activity to the hyperplane that best classifies the stimulus change of that feature (e.g., higher vs. lower spatial frequency). (**B**) Trial-by-trial comparison between belief strength (abscissae, decoded from 7a) and perceptual discriminability (ordinates, decoded from V1) for an example stimulus/task condition. If belief decisions and perceptual decisions are implemented by separate functional modules of the brain, then internal fluctuations of the two systems should have no correlation. (**C**) The belief- spatial location discriminability correlation is positive when spatial location is believed to be relevant (histogram and magenta cumulative distribution curve, $p=4\times10^{-6}$), but not when it is believed to be irrelevant (cyan cumulative distribution curve, $p>0.05$). The two distributions are significantly different (Wilcoxon rank sum test, $p=0.014$). (**D**) Similarly, belief- spatial frequency discriminability is significantly positive when spatial frequency is believed to be relevant ($p=0.0015$) but not when it is believed to be irrelevant ($p>0.05$). The two distributions are significantly different (Wilcoxon rank sum test, $p=0.03$).

In addition to the trial-by-trial interaction between belief and perception, two pieces of evidence demonstrate that perception on the previous trial affects how task-belief is updated for the upcoming trial.

First, trial information beyond reward outcome affects how task-belief will be updated. On average, the monkeys were more likely to switch tasks after they missed rewards on trials with big changes in the stimulus feature believed to be relevant (Figure S5B). This reliance of belief updating on vision is captured by our ideal observer model which optimally updates belief to changes in the environment based on historical reward, stimulus, and choices (see Supplementary Methods and Figure S4). The ideal observer model consistently predicts the animals' behavior better than an alternative strategy in which every unrewarded trial affects belief independent of visual and choice experience (Figure S5A). These results demonstrate that, consistent with findings from studies with similar task structure (*10, 12*), confidence in historical choices inform belief updating.

Second, even when trial conditions (stimulus, choice, and reward) were identical, there is a trial-to-trial relationship between fluctuations in the representation of visual stimuli and belief updating. We captured uncontrolled fluctuations in perception by fitting estimates of each feature from V1 (Figure 3A) using logistic regression (Figure 4A). We used this population neurometric curve to estimate the monkey's confidence about each perceptual choice, and used this to predict the animals' task switching decisions (Figure 4B). This model predicts the animals' task switching decisions better than an alternative model that does not incorporate trial-by-trial variability in V1 (Figure S6). Furthermore, if we shuffle V1 responses among trials with identical trial conditions (Figure 4C), the model's switch prediction performance suffers significantly (Figures 4D, E). This difference likely reflects confidence fluctuations in past visual discrimination, since for identical trial conditions, the monkeys were more likely to switch tasks after they missed rewards on trials with larger relevant feature discriminability estimated from V1 (Figure S3C). Together, these results demonstrate that trial-to-trial fluctuations in perception affect belief updating on the subsequent trial, even though these fluctuations provide no benefit for estimating the relevant feature.

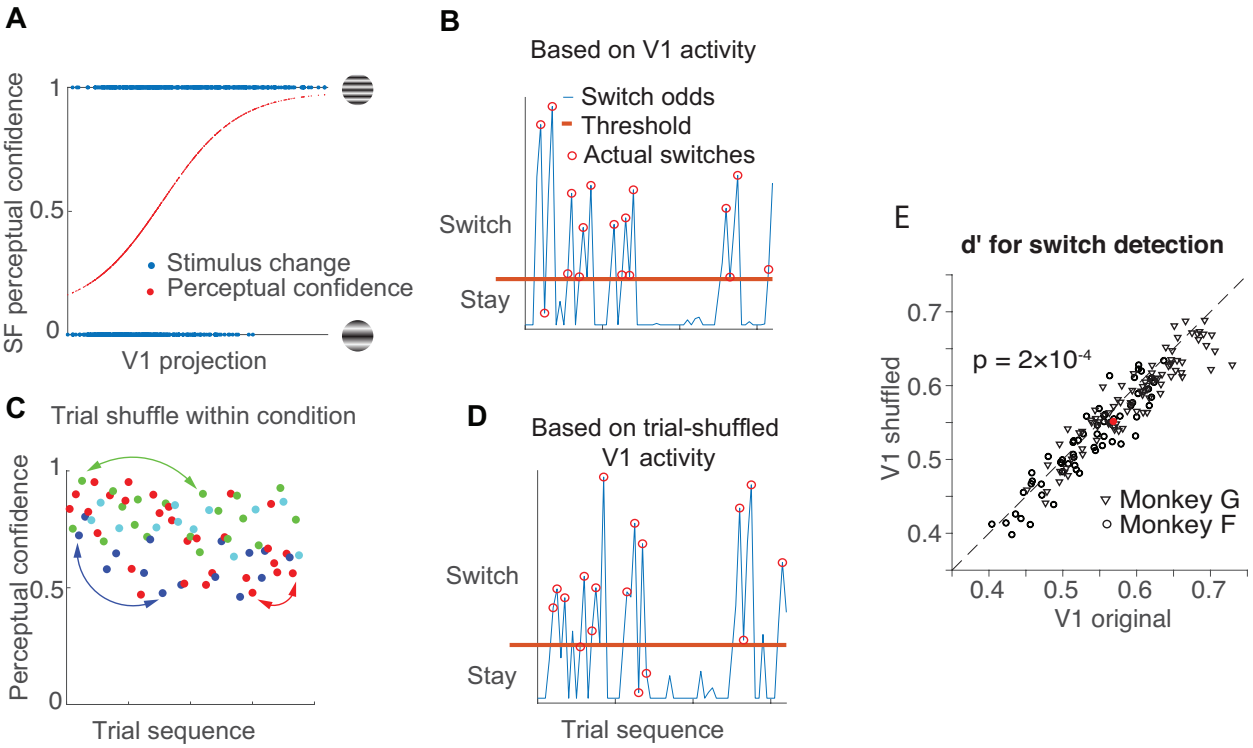

Fig. 4 Trial to trial variability in visual cortex affects belief updating. (**A**) Example neurometric curve showing the ability of a decoder to discriminate spatial frequency changes from the population of recorded V1 neurons using logistic regression on the perceptual discriminability of spatial frequency (as in Figure 3A). (**B**) Based on perceptual confidence on each trial (estimated from V1

population activity), a normative model determines whether the subject should switch tasks given the trial history (see Supplementary Methods). (**C**) Based on the V1 projections to the relevant feature subspace on each trial, we estimate from the neurometric curve, which represents the probability the monkey's behavioral choice is correct, and we take this value to be the confidence in the perceptual choice. In the trial-shuffle analysis, we randomly switch the confidence within trials with the same conditions (dots with same color). (**D**) Model predictions after trial-shuffle, conventions as in (B). (**E**) Trial-to-trial variability in V1 is related to belief. The model's ability to predict whether the monkey would switch tasks is better using the actual than trial-shuffled V1 activity. Each data point here represents an experimental session, and its coordinates show the sensitivity index (*d'*) of switch prediction for the model based on original (x-axis) or trial-shuffled V1 activity (y-axis).

Our findings demonstrate that there is no such thing as a standalone perceptual decision-making process: every aspect of perceptual decision-making is profoundly integrated with the dynamic belief states that dominate natural behavior. A foraging animal may frequently switch between evaluating food sources and searching for subtle signs of predators based on evolving beliefs about the safety of the environment. Using a combination of multi-neuron, multi-area physiology, complex but controlled behavior, and hypothesis-driven dimensionality reduction, we demonstrated that perception and task-belief are intimately intertwined such that weak task-beliefs are associated with poor perception of task-relevant information. This suggests that fluctuation in belief strength, instead of reflecting a homogenous process such as arousal, may have specific effects on the believed relevant information only. It will be interesting to determine whether fluctuations in other types of belief (e.g. those reviewed in (*21*)) interact with decision making in different ways.

The idea that dynamic task-beliefs and decisions are inextricably linked opens up exciting avenues for therapies that address deficits in flexible decision-making associated with neuropsychiatric disorders. For instance, our results imply that cognitive flexibility is mediated by interactions between neural populations responsible for perception and belief. As such, therapies that affect communication between brain areas (e.g. by affecting neurotransmitters like dopamine (*22, 23*) have the potential to improve cognitive flexibility in health and disease. Indeed, stimulants that affect the dopamine system like methylphenidate or amphetamines can change focus and flexibility (*23, 24*). Going forward, studying the highly integrated belief-based decision-making system will open up doors to potential treatments of conditions that affect cognitive flexibility and even solutions for healthy individuals to become better decision-makers in volatile environments.

Materials and Methods

Experimental subjects

The subjects in our study were two adult male rhesus monkeys (*Macaca mulatta*, monkey F weighed 12 kilograms, monkey G weighed 9 kilograms), who were both experimentally naïve prior to the current experiments. All animal procedures were approved by the Institutional Animal Care and Use Committees of the University of Pittsburgh and Carnegie Mellon University. After we implanted each animal with a titanium head post, they were trained to perform two-interval, two-feature discrimination with stochastic rule switching (Figure) (monkey 1 was trained for 11 months, monkey 2 for 9 months). We made sure the animals understood the essential requirements of the task based on their behavior (Figure 1), before implanting each animal with 6×8 microelectrode arrays (Blackrock Microsystems) in both parietal cortical area 7a and visual cortical area V1. Each array was connected to a percutaneous connector that allowed daily electrophysiological recordings. The distance between adjacent electrodes was 400 μm, and each electrode was 1 mm long. We identified areas 7a and V1 using stereotactic coordinates and by visual inspection of sulcal landmarks.

Behavioral task

To study perceptual decision making under evolving task-beliefs in dynamic environment, we trained the animals to perform a two-interval, two-feature discrimination task with stochastic task switching. A trial began when the subjects fixated their gaze on a central dot on the screen and they were required to maintain fixation as long as the dot remained on the screen, or the trial would be aborted and unrewarded. A Gabor stimulus was then displayed for 200 ms. After a random delay (300ms to 500ms), a second Gabor stimulus was displayed for 200 ms with a slightly different spatial location (shifted left or right) and a slightly different spatial frequency (higher or lower), with independently randomized change amounts in the two features. The ranges of change amounts are titrated at the beginning of each session so that the overall perceptual performances of the spatial location task and spatial location task are both approximately 75%. Following a subsequent delay of 150ms, the fixation dot disappeared, and the animals looked at one out of four peripheral targets to indicate both the inferred relevant feature and the direction of change in that feature. The two cyan targets correspond to the increase and decrease of spatial frequency when it was believed to be the relevant feature, while the two magenta targets correspond to the left-shift and right-shift of spatial location. The monkeys were rewarded only if they correctly reported the direction of change in the relevant feature. The visual stimuli throughout a trial contain no information about the behavioral relevance of features. The relevant feature switches on a randomly chosen 2.5% of trials. The monkeys therefore needed to infer the relevant feature based on their choice and reward history.

Electrophysiological recording

All visual stimuli were displayed on a linearized CRT monitor (1,024 × 768 pixels, 120-Hz refresh rate) placed 57 cm from the animal. We monitored eye position using an infrared eye tracker (Eyelink 1000, SR Research) and used custom software (written in Matlab using the Psychophysics Toolbox, (*25*) to present stimuli and monitor behavior. We recorded eye position and pupil diameter (1,000 samples per s), neuronal responses (30,000 samples per s) and the signal from a photodiode to align neuronal responses to stimulus presentation times (30,000 samples per s) using hardware from Ripple.

We recorded neuronal activity from Utah arrays during daily experimental sessions for several months in each animal (89 sessions from monkey F and 68 sessions from monkey G). We set the threshold for each channel at three times the standard deviation and used threshold crossings as the activity on that unit. We positioned the stimuli to maximize the overlap between potential stimulus locations and the joint receptive fields of V1 units, as determined using separate data collected while the monkeys fixated and Gabor stimuli were flashed across a range of retinal positions.

We included experimental sessions if they contained at least 480 completed trials (where monkeys successfully maintained fixation until they indicated their choice). We analyzed the activity of area 7a units during the first 300ms after the offset of the first stimulus, when there is no Gabor stimulus on the screen; and the activity of area V1 units during stimulus display periods, shifted with 34 ms visual latency. Units from area 7a were included if their average activity during the delay period was at least 5 sp/s. Units from V1 were included in the analyses if their average stimulus response was 1) at least 25% larger than baseline activity, measured 100ms before stimulus onset, and 2) larger than 5 sp/s. These procedures resulted in 89 sessions from Monkey F and 68 sessions from Monkey G; average 53 7a units, 46 V1 units, and 1053 completed trials per session.

Population analyses

To obtain a continuous neuronal measure of the animals' belief state, we analyzed the activity of the population of 7a neurons during the delay period in a high dimensional space in which the activity of each unit was one dimension. We used linear discriminant analysis to identify the best hyperplane to discriminate between 7a population activity on trials where monkeys chose spatial location targets from trials when they chose spatial frequency targets. We defined the belief strength on each trial as the Euclidean distance from the 7a population response to the discriminant hyperplane. Similarly, we obtained a continuous neuronal measure of the discriminability of stimulus change using V1 activity.

Normative behavioral model

We use a normative model to characterize belief updating of an ideal observer, given the trial history and the perceptual ability of the monkey (*10, 12*). Based on the monkeys' psychometric curve in an experiment session and the change amount of the chosen feature in each trial, we estimated the trial-by-trial probability that their perceptual choice was incorrect. For a non-reward trialed, the odds of likelihoods that the actual task is different from the monkeys' subjective belief is given by

$$\frac{p(diff|nr,\theta,c)}{p(same|nr,\theta,c)} = \frac{p(nr,\theta,c|diff)\cdot p(diff)}{p(nr,\theta,c)} \cdot \frac{p(nr,\theta,c)}{p(nr,\theta,c|same)\cdot p(same)} = \frac{p(nr,\theta,c|diff)}{p(nr,\theta,c|same)}$$

where θ and c refer to the stimulus change amount and perceptual choice in the feature the monkeys believed to be relevant; and nr refers to a non-reward trial outcome. We assumed that overall, the monkeys experienced an equal number of trials where the feature was the same or different from their current belief (i.e., $p(diff) = p(same)$). The monkeys were never rewarded on trials when their subjective task-belief was different from the actual task rule, so $p(nr,\theta,c|diff) = 1$. Meanwhile when subjective task-belief is consistent with the actual rule, the probability of perceptual error can be simply derived from the psychometric function associated with that choice:

$$\frac{p(diff|nr,\theta,c)}{p(same|nr,\theta,c)} = \frac{1}{1-p(\theta|c)}$$

where $p(\theta|c)$ is the psychometric function associated with the perceptual choice c (Figure 1C). In Figure 4, the psychometric function is replaced with the neurometric function (see *Population analyses* section).

For n consecutive non-reward trials, the likelihood ratio grows larger as perceptual evidence for a task switch grows as

$$\frac{1-\mathcal{L}_n^{same}}{\mathcal{L}_n^{same}} = \prod_{i=1}^{n} \frac{1}{1-p(\theta_i|c_i)}$$

where $\mathcal{L}_n$ is the likelihood that the task has not changed after n consecutive non-rewarded trials (examples in Figure S4A). Aside from perceptual evidence, the observer presumably also has prior knowledge about the volatility of the task environment. After n consecutive non-reward trials, the prior probability that the task stays the same with the last rewarded trial is

$$Pr_n^{same} = Pr_{n-1}^{same}\cdot(1-h) + (1-Pr_{n-1}^{same})\cdot h$$

where h represents the hazard rate of task change at each trial (for an ideal observer, $h = 0.025$, see *behavioral task* section), with

$$Pr_0^{same} = 1-h$$

Examples of Pr_n^{same} under different environment volatility are shown in Figure S4B. Taking both perceptual evidence and prior knowledge of environment volatility into account, the model shows that an

ideal observer should switch tasks if the posterior probability is higher for actual task rule being different from the subjective task-belief than when they are the same (Figure S4C), i.e.:

$$switch\ odds = \text{loglog}((1 - Pr_n^{same}) \cdot (1 - \mathcal{L}_n^{same})) - \text{loglog}(Pr_n^{same} \cdot \mathcal{L}_n^{same}).$$
$$when\ positive,\ \ switch\ task$$

Acknowledgments: We are grateful to Karen McKracken for providing technical assistance and to Josh Gold, Lori Holt, Douglas Ruff, Amy Ni, and Ramanujan Srinath for comments on an earlier version of this manuscript.

Funding: This work was supported by the Simons Foundation (Simons Collaboration on the Global Brain award 542961SPI to MRC), the National Eye Institute of the National Institutes of Health (award R01EY022930 to MRC), and the McKnight Foundation (McKnight Scholar award to MRC).

Author contributions: C.X., and M.R.C. designed research; C.X., and L.E.K. performed research; C.X. analyzed data; and C.X., and M.R.C. wrote the paper.

Competing interests: The authors declare no competing interests.

Data and Materials availability: Data and code are available upon reasonable request.

References:

1. J. I. Gold, M. N. Shadlen, The neural basis of decision making. *Annu Rev Neurosci* **30**, 535-574 (2007).
2. J. W. Kable, P. W. Glimcher, The neurobiology of decision: consensus and controversy. *Neuron* **63**, 733-745 (2009).
3. N. Uchida, A. Kepecs, Z. F. Mainen, Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat Rev Neurosci* **7**, 485-491 (2006).
4. K. T. Brady, K. M. Gray, B. K. Tolliver, Cognitive enhancers in the treatment of substance use disorders: clinical evidence. *Pharmacol Biochem Behav* **99**, 285-294 (2011).
5. A. Thapar *et al.*, Psychiatric gene discoveries shape evidence on ADHD's biology. *Mol Psychiatry* **21**, 1202-1207 (2016).
6. D. P. Dickstein *et al.*, Cognitive flexibility in phenotypes of pediatric bipolar disorder. *J Am Acad Child Adolesc Psychiatry* **46**, 341-355 (2007).
7. G. Stoet, L. H. Snyder, Neural correlates of executive control functions in the monkey. *Trends Cogn Sci* **13**, 228-234 (2009).
8. T. J. Buschman, E. L. Denovellis, C. Diogo, D. Bullock, E. K. Miller, Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron* **76**, 838-846 (2012).
9. T. Kamigaki, T. Fukushima, Y. Miyashita, Cognitive set reconfiguration signaled by macaque posterior parietal neurons. *Neuron* **61**, 941-951 (2009).
10. M. Sarafyazd, M. Jazayeri, Hierarchical reasoning by neural circuits in the frontal cortex. *Science* **364**, (2019).

11. R. Bartolo, B. B. Averbeck, Prefrontal Cortex Predicts State Switches during Reversal Learning. *Neuron* **106**, 1044-1054.e1044 (2020).
12. B. A. Purcell, R. Kiani, Hierarchical decision processes that operate over distinct timescales underlie choice and changes in strategy. *Proc Natl Acad Sci U S A* **113**, E4531-4540 (2016).
13. V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78-84 (2013).
14. B. A. Purcell, R. Kiani, Neural Mechanisms of Post-error Adjustments of Decision Policy in Parietal Cortex. *Neuron* **89**, 658-671 (2016).
15. R. B. Ebitz, J. C. Tu, B. Y. Hayden, Rules warp feature encoding in decision-making circuits. *PLoS Biol* **18**, e3000951 (2020).
16. M. R. Cohen, J. H. Maunsell, Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* **70**, 1192-1204 (2011).
17. M. R. Cohen, J. H. Maunsell, When attention wanders: how uncontrolled fluctuations in attention affect performance. *J Neurosci* **31**, 15802-15806 (2011).
18. S. Monsell, Task switching. *Trends Cogn Sci* **7**, 134-140 (2003).
19. G. Stoet, L. H. Snyder, Single neurons in posterior parietal cortex of monkeys encode cognitive set. *Neuron* **42**, 1003-1012 (2004).
20. N. T. Markov *et al.*, A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cereb Cortex* **24**, 17-36 (2014).
21. W. J. Ma, M. Jazayeri, Neural coding of uncertainty and probability. *Annu Rev Neurosci* **37**, 205-220 (2014).
22. M. Botvinick, T. Braver, Motivation and cognitive control: from behavior to neural mechanism. *Annu Rev Psychol* **66**, 83-113 (2015).
23. A. Mueller, D. S. Hong, S. Shepard, T. Moore, Linking ADHD to the Neural Circuitry of Attention. *Trends Cogn Sci* **21**, 474-488 (2017).
24. K. S. Bagot, Y. Kaminer, Efficacy of stimulants for cognitive enhancement in non-attention deficit hyperactivity disorder youth: a systematic review. *Addiction* **109**, 547-557 (2014).
25. D. H. Brainard, The Psychophysics Toolbox. *Spatial vision* **10**, 433-436 (1997).