Neuron
Report

Effects of Cortical Microstimulation on Confidence in a Perceptual Decision

Christopher R. Fetsch,[1,4] Roozbeh Kiani,[2,4] William T. Newsome,[3] and Michael N. Shadlen[1,*]
[1]Howard Hughes Medical Institute, Department of Neuroscience and Kavli Institute for Brain Science, Columbia University, New York, NY 10032, USA
[2]Center for Neural Science, New York University, New York, NY 10003, USA
[3]Department of Neurobiology and Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA
[4]Co-first author
*Correspondence: shadlen@columbia.edu
http://dx.doi.org/10.1016/j.neuron.2014.07.011

SUMMARY

Decisions are often associated with a degree of certainty, or confidence—an estimate of the probability that the chosen option will be correct. Recent neurophysiological results suggest that the central processing of evidence leading to a perceptual decision also establishes a level of confidence. Here we provide a causal test of this hypothesis by electrically stimulating areas of the visual cortex involved in motion perception. Monkeys discriminated the direction of motion in a noisy display and were sometimes allowed to opt out of the direction choice if their confidence was low. Microstimulation did not reduce overall confidence in the decision but instead altered confidence in a manner that mimicked a change in visual motion, plus a small increase in sensory noise. The results suggest that the same sensory neural signals support choice, reaction time, and confidence in a decision and that artificial manipulation of these signals preserves the quantitative relationship between accumulated evidence and confidence.

INTRODUCTION

Decision making refers to the process of deliberating toward a commitment to a proposition, hypothesis, or plan of action. Although decisions have a discrete, all or none character—true or false, left or right, option D—they are also associated with a degree of belief that the decision will turn out to be correct. This graded scale of choice certainty, or confidence, affects the way we express our decisions (forcefully or reservedly) and learn from our mistakes and successes. Confidence is critically important when making interrelated decisions without immediate feedback, or when reasoning about a sequence of choices vicariously (Tolman, 1948). For these and other reasons, psychologists have counted confidence among the three main observables of choice behavior (Vickers, 1979), along with the outcome of the decision (correct or incorrect; i.e., accuracy) and the time needed to complete it (reaction time [RT]).

The neural basis of assigning confidence in a decision is not well understood, in part because it is difficult to study in animal models. Recently, methods have been introduced that allow animals to report their confidence, often in the form of a postdecision wager (PDW) (Foote and Crystal, 2007; Hampton, 2001; Kepecs et al., 2008; Kiani and Shadlen, 2009; Middlebrooks and Sommer, 2011; Smith et al., 2008). In one type of PDW, animals indicate their degree of certainty by opting out of the primary behavioral report when a decision is less likely to be successful, instead choosing a guaranteed but smaller (or less preferred) reward (Foote and Crystal, 2007; Hampton, 2001; Kiani and Shadlen, 2009). Monkeys exercise this “sure-bet” option more frequently when the trial is difficult, and they are more accurate when the sure bet is offered and waived versus when it is not offered. This improvement holds within a particular level of stimulus difficulty and even for identical replays of the same stimulus (Kiani and Shadlen, 2009). It suggests that the decision to accept the sure bet is based on an assessment of the reliability of internal sensory evidence, rather than on a simple association with trial difficulty or some property of the stimulus (Smith et al., 2012).

A recent study (Kiani and Shadlen, 2009) reported the activity of decision-related neurons in the lateral intraparietal area (LIP) recorded while monkeys performed a direction discrimination task with PDW. They found that these neurons—previously shown to represent a decision variable (DV) that explains choice and RT (Gold and Shadlen, 2007)—also reflect the degree of confidence in the choice. The results raised the possibility of a common neural mechanism underlying choice, RT, and confidence. This hypothesis makes a clear prediction: if the representation of accumulated evidence used to guide a perceptual decision also supports a degree of confidence in that decision, then a causal manipulation of the evidence will affect PDW in a manner predictable from the effect on choices.

Here we test this hypothesis using electrical microstimulation (μS). Previous studies showed that μS of direction-selective neurons in the macaque visual cortex during a direction discrimination task causes monkeys to choose the preferred direction of neurons near the electrode tip more often (Salzman et al., 1990, 1992) and more rapidly (Ditterich et al., 2003). These effects on choice and RT can be quantified as an equivalent change in the motion strength, as though the stimulation effectively added to the visual evidence supporting the preferred direction. Thus, a change in PDW commensurate with the shift

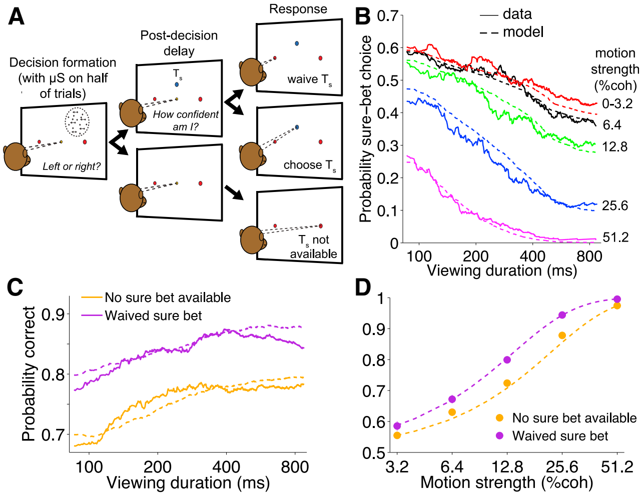

Figure 1. Postdecision Wagering Reflects Confidence in the Motion Decision

(A) Postdecision wagering (PDW) task sequence (see Experimental Procedures). Red spots indicate direction targets; blue spot is the "sure-bet" target (T_S). (B) Probability of choosing T_s as a function of viewing duration and motion strength (color coded). Combined data from no-μS trials for two monkeys (n = 26,924 trials). Solid traces are running means (proportions) of the data sorted by viewing duration. Dashed traces in all panels are fits to the bounded accumulation model (see text and Figure 4). (C) Improvement in decision accuracy on no-μS trials when the sure bet was offered but waived. Solid traces are running means using all nonzero coherences and directions. (D) Same format as (C) but broken down by motion strength (absolute value of coherence) and pooled across viewing durations. Symbols indicate the mean ± SE.

of the choice function would support the idea that the same sensory signals underlie both the choice and the confidence associated with it. On the other hand, since choice and confidence are known to be dissociated in a variety of settings (Del Cul et al., 2009; Drugowitsch et al., 2014; Kahneman et al., 1982; Komura et al., 2013; Lau and Passingham, 2006; Rahnev et al., 2012; Rounis et al., 2010), we might expect artificial stimulation to induce a discrepancy between the two. Indeed, the effect of μS on neuronal circuits is unlike anything elicited through natural vision (Histed et al., 2009; Logothetis et al., 2010). Nonetheless, here we show that μS affects confidence much like a change in the visual stimulus, consistent with a common mechanism characterized by bounded accumulation of evidence.

RESULTS

We trained two rhesus monkeys on a two-alternative direction discrimination task with PDW (Figure 1A; Experimental Procedures). The monkeys were required to decide between the direction preferred by neurons near the stimulating electrode and the opposite "null" direction and to indicate this choice after a memory delay. Monkeys were rewarded for correct choices and randomly on the neutral (0% coherence) stimulus. During the memory delay, the monkey was sometimes offered a third alternative (the sure-bet target [T_S]) to opt out of the high-stakes direction decision and receive a guaranteed but smaller reward. Monkeys chose T_S more frequently for shorter viewing durations and weaker motion strengths (Figure 1B) and showed greater accuracy on waived-T_S trials compared to when T_S was unavailable (Figures 1C and 1D). Figure 1C also reveals a saturation in performance with longer viewing durations, suggesting a bounded accumulation process (Kiani et al., 2008; Kiani and Shadlen, 2009).

On half of all trials, electrical μS (5–10 μA) was applied to area MT or MST during the presentation of random dot motion. Importantly, the presence or absence of μS did not alter the designation of correct and incorrect trials. Thus, if microstimulation of rightward preferring neurons were to cause the monkey to answer "right" on a trial in which leftward motion was shown, this would be regarded as an error, hence unrewarded.

Does Microstimulation Affect the Degree of Confidence?

Since microstimulation induces an artificial pattern of activity in the brain, we wondered whether monkeys would simply opt out of the direction decision on μS trials when given this opportunity. The answer is resoundingly negative. As shown in Figure 2A, monkeys varied their propensity to choose T_S from session to session, but such variation was highly correlated on μS and no-μS trials (Pearson's r = 0.88, $p < 10^{-20}$). Averaged within individual experiments, monkeys did not opt out more frequently on μS trials; indeed, the trend favors a small decrease in T_S choices (see Supplemental Experimental Procedures available online). The key point is that μS did not cause indiscriminate uncertainty about perceptual judgments, or we would have observed the opposite trend (more T_S choices on μS trials).

Although the average frequency of T_S choices was similar on μS and no-μS trials, microstimulation nevertheless exerted a substantial effect on confidence judgments, corresponding to a shift of the bell-shaped function along the motion axis (Figure 2B, top). Notice that for both μS and no-μS trials, the monkey accepted the sure bet most often for the stimulus conditions that led to the most equivocal choice proportions (i.e., 0.5 preferred-direction choices; Figure 2B, bottom). Across sessions, the shift of the sure-bet function was highly correlated with the shift of the choice function (Pearson's r = 0.87, $p < 10^{-18}$; Figure S1A). This close association, despite the wide range of magnitudes of both effects, is consistent with the idea that a common neural signal underlies choice and confidence.

Another way to frame this result is to consider each motion direction separately. For motion in the preferred direction (positive coherence), monkeys chose the sure bet less often when μS was present (two-proportion z test, $p < 10^{-18}$), suggesting that μS increased confidence by reinforcing the evidence from the visual stimulus. In contrast, for motion in the null direction (negative coherence) monkeys chose the sure bet more often when μS was present ($p < 10^{-9}$), suggesting that μS decreased confidence by contradicting the evidence for null-direction motion. The end result is a leftward shift of the curve, as if μS had injected a signal largely equivalent to a change in motion coherence. Lastly, Figure 2B (top) clarifies the subtle decrease in the number of T_S choices accompanying μS (noted above), which is most evident at the peaks of the sure-bet functions. We will explain this apparent increase in confidence using the model described below.

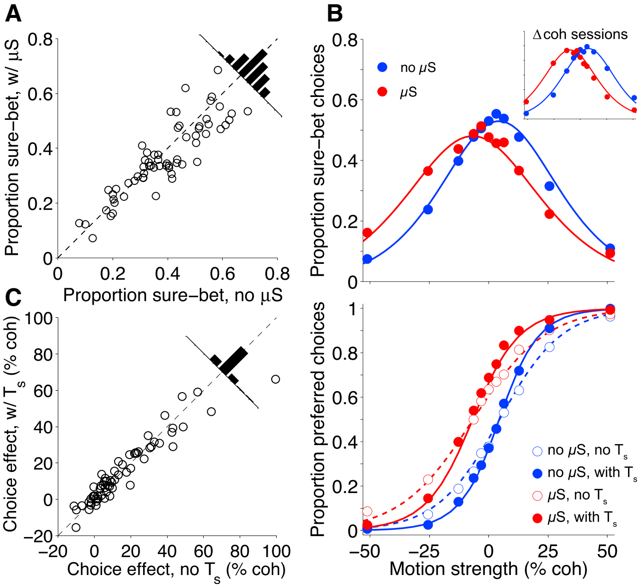

Figure 2. Effects of µS on PDW and Perceptual Decisions

(A) The proportion of trials in which the monkey opted out of the direction task and chose the sure bet, comparing µS and no-µS trials (n = 63 sites). (B) Top: proportion of sure-bet (T_s) choices as a function of motion strength (percent coherence; positive = preferred direction of neurons at the stimulation site) for all sessions in both monkeys (n = 53,134 trials). Red and blue data points indicate µS and no-µS trials, respectively, combining across all viewing durations. Error bars (SE) are smaller than the data points. Top, inset: proportion T_s choices in separate control sessions, for trials with (red) and without (blue) an offset added to the motion coherence in lieu of µS ("Δcoh," see text and Figure S3). Bottom: proportion of preferred-direction choices as a function of motion strength, plotted separately for the four conditions of the 2 × 2 design: µS present (red) or absent (blue), and T_s offered but waived (solid curves and filled symbols) or T_s not offered (dashed curves and open symbols). In both panels, smooth curves represent fits to the bounded evidence-accumulation model (see text), with the exception of the red solid and dashed curves in the bottom panel. These are the predicted µS choice functions based on a fit to the remaining observations. (C) Comparison of the effect of µS on choices (represented as an equivalent change in motion strength) on trials with and without the T_s option.

Microstimulation combined with PDW yields four choice functions: ±µS when the sure bet was not offered and ±µS when the monkey could have opted out but instead chose the preferred or null direction of the stimulated neurons. These four conditions are represented by the four curves in the bottom of Figure 2B. There are three salient observations. First, µS shifted the choice function by the same amount whether or not T_s was available, and this similarity was apparent across the 63 sites (Figure 2C; paired t test, p = 0.56). If on some trials µS had affected confidence in a manner unlike a change in visual evidence, its effect on choices might have been different depending on whether the monkey had the chance to opt out of such trials. The results did not support this possibility. Second, µS reduced monkeys' sensitivity to motion, consistent with previous studies (Ditterich et al., 2003; Salzman et al., 1992). This is only apparent as a subtle attenuation in the slope of the red curves compared with their blue (no-µS) counterparts, but the effect is reliable (17% ± 2% change, $p < 10^{-16}$, logistic regression; Equation 2). It suggests that µS occasionally weakened the directional signal and/or added a small amount of noise to the decision process. As shown below, an increase in noise can also explain the small decrease in the maximum rate of T_s choices on µS trials. Third and most importantly, µS did not abolish the improved sensitivity to motion on trials when T_s was available but waived, as indicated by the steeper slope of the solid compared to the dashed curves (Figure 2B, bottom; logistic regression, $p < 10^{-24}$ for both µS and no-µS conditions tested separately). Recall that this improvement is a sign that the monkey evaluated the reliability of the evidence and communicated its direction choice when the reliability seemed high (Kiani and Shadlen, 2009). Its presence on µS trials implies that such evaluation of evidence is not compromised by artificially manipulating the sensory representation. It also means that µS did not simply compel the monkey to choose the preferred direction with some probability, irrespective of the state of the perceptual evidence. Rather, it exerted its effects by changing the available evidence for the decision.

Controls: High-Current Stimulation and a Visual Perturbation

The data presented thus far suggest that µS does not reduce the monkey's overall degree of certainty but instead resembles a change in visual motion. However, it is possible that the absence of an increase in the overall frequency of T_s choices on µS trials was due to the monkey's inability or unwillingness to choose T_s beyond some rate throughout the experiment. One way to test this possibility is to apply a µS condition that impairs discrimination performance (e.g., by deliberately weakening the differential directional signals underlying choice). This kind of impairment can be achieved simply by increasing the current amplitude, thus activating indiscriminately a larger population of neurons with a broad distribution of preferred directions (Murasugi et al., 1993). Thus, at eight sites, after completing a block of trials with standard low-amplitude pulses (7.5 µA), we began a second block with 75 µA pulses while keeping all other parameters identical. High-current µS reduced the monkey's sensitivity to motion (logistic regression, $p < 10^{-19}$; Figure 3B) and also led to a greater proportion of T_s choices (no-µS = 0.45 ± 0.01; µS = 0.49 ± 0.02; $p < 0.05$; Figure 3A). The latter effect can be described as primarily a widening, rather than a shift, of the sure-bet function, driven by a pronounced increase in T_s choices for the highest motion strengths. The result implies that PDW does not lack the power to expose a decrease in confidence, and it reinforces the notion that decision accuracy and confidence are linked. Indeed this link was also present in the main experiments: across sessions, flatter choice functions were associated with wider sure-bet functions (Spearman's rank correlation, $\rho = -0.55$, $p < 10^{-5}$), and the modest changes in these two metrics caused by low-current stimulation were correlated ($\rho = -0.37$, $p < 0.004$; see Supplemental Experimental Procedures).

In a second control experiment, instead of stimulating the brain electrically on half the trials, we manipulated the visual stimulus in a manner that mimics the hypothesized effect of µS. We reasoned that if the brain interprets µS like a change in motion strength, we should approximate the effects of µS on choice and PDW by simply adding an offset to the motion

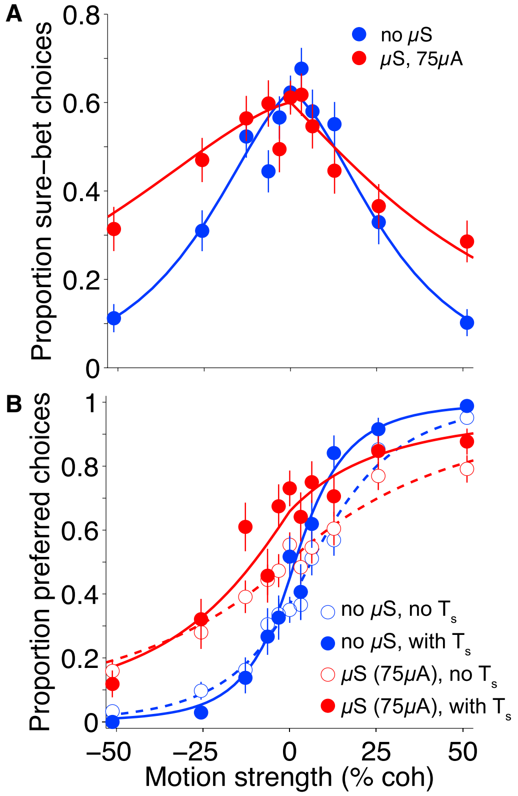

Figure 3. Stimulation with High Current Disrupts Both Accuracy and Confidence

(A and B) Combined data from eight experiments using 75 μA stimulation (n = 4,483 trials). Same conventions as Figure 2B. Smooth curves are best fits of the extended model described in the Supplemental Experimental Procedures (see also Figure S2).

coherence, termed Δcoh, in one arbitrary "preferred" direction (Salzman et al., 1992). We performed 42 such experiments, using a range of Δcoh values (varied across sessions) to approximate the range of μS effects in the main experiment. Like μS, Δcoh trials were rewarded based on the direction of motion that would have occurred in the absence of a coherence offset. As expected, the Δcoh manipulation shifted the pattern of direction and T$_s$ choices by an amount similar to the magnitude of added coherence. The results were largely comparable to μS sessions, including similar shifts of the T$_s$ curve and the choice function (Figures S1B and S3B), similar effects on direction choices with and without the T$_s$ offer (Figure S3C; $p = 0.26$, paired t test), and improved sensitivity on T$_s$-waived trials (Figure S3B, bottom; logistic regression, $p < 10^{-17}$). A notable difference from μS is the lack of an effect on the maximum rate of T$_s$ choices (compare Figure 2B [top] versus inset; see below and Supplemental Experimental Procedures).

A Common Mechanism for the Effect of Microstimulation on Choice and Confidence

In the absence of μS, both direction choices and PDW are explained by the accumulation of noisy evidence bearing on the direction of motion (Kiani and Shadlen, 2009) (Figure 4A). The model assumes that (1) a direction choice is based on the sign of the accumulated evidence and (2) a sure-bet choice supersedes a direction choice if the odds that the direction choice will be correct are less than a criterion, θ (Figure 4B). The latter requires the brain to have implicit knowledge of the association between the accumulated evidence, termed a decision variable (DV), and the likelihood that a decision based on this evidence will be correct. We fit this model to the monkey's direction and T$_s$ choices on no-μS trials (Figure 2B, top and bottom, blue symbols) and then incorporated the effect of μS as a perturbation of the evidence, equivalent to a change in motion coherence and/or a change in sensory noise (see Supplemental Experimental Procedures). The parameters implementing μS were fit using only the T$_s$ choices on μS trials (Figure 2B, top, red symbols). We then used the fitted parameters to predict the pattern of direction choices on μS trials, both with and without the T$_s$ option available (Figure 2B, bottom, red filled and open symbols). Importantly, we did not allow μS to alter the association between DV and

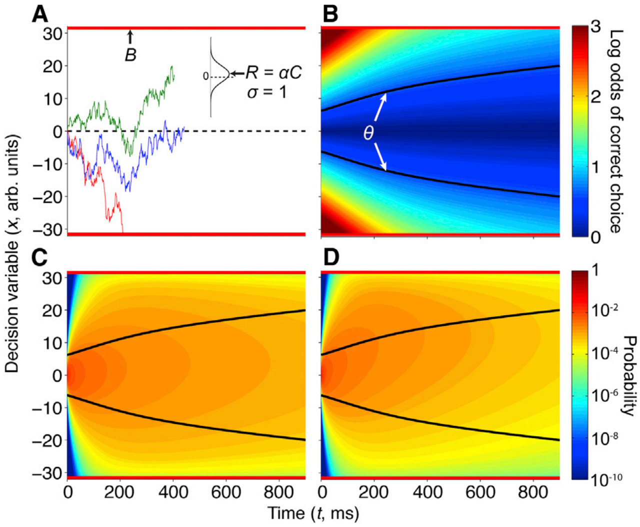

Figure 4. Bounded Evidence Accumulation Model Explains PDW and Effects of μS

(A) Colored traces represent the accumulation of noisy motion evidence (i.e., from MT/MST) on three individual (simulated) trials. Evidence is drawn from a Gaussian distribution with mean (R) proportional to motion strength and SD (σ) equal to 1. Decision formation terminates when the stimulus is turned off (green and blue trials) or when the accumulated evidence (the decision variable, x) reaches a bound at ±B (red trial). (B) The model prescribes a sure-bet choice when the logarithm of the odds of being correct is below a fixed threshold, θ, indicated by the black contours which divide the x, t plane into low- and high-confidence regions. For example, the blue trial in (A), but not the green trial, would have ended in a sure-bet choice. (C and D) The probability density of the decision variable across time is shown for a particular motion coherence (3.2%, i.e., weak preferred-direction motion), either without (C) or with (D) μS. Bias in this example was set to zero for simplicity. Microstimulation shifts the density upward, thereby decreasing the probability of a sure bet and increasing the probability of a preferred choice (see Figure S4). A key assumption of the model is that the brain applies the same mapping between accumulated evidence and the expected log odds of being correct (i.e., confidence), and the same criteria for opting out (black contours in B–D), irrespective of the presence of μS.

Table 1. Maximum-Likelihood Estimates of Model Parameters, ± SE

	μS Data Set (n = 53,134 Trials)	Δcoh Data Set (n = 43,054 Trials)
α	0.294 ± 0.001	0.291 ± 0.002
B	31.6 ± 1.30	31.2 ± 1.00
γ	−0.0134 ± 0.0002	−0.0224 ± 0.0002
θ	0.609 ± 0.004	0.507 ± 0.004
δ_C	0.112 ± 0.001	0.173 ± 0.001
δ_{σ^2}	0.237 ± 0.015	0.039 ± 0.006

confidence or the criterion for selecting T_s (i.e., θ). This strategy formalizes the qualitative assertion that μS-induced changes in neural activity are processed like vision-induced changes in neural activity. In other words, the brain does not know that it is being stimulated.

The model fits and predictions are shown by the smooth curves in Figure 2B (and dashed curves in Figures 1B–1D). They capture several key features of the data: (1) the relationship between T_s choices and trial difficulty (i.e., motion strength and viewing duration; Figures 1B and 2B, top, blue curve), (2) the improvement in sensitivity when T_s was offered but waived (Figure 2B, bottom: solid versus dashed curves; Figures 1C and 1D), and (3) the main effects of μS on choice and PDW (Figure 2B, horizontal shift of red versus blue curves).

Not surprisingly, the model explains the results from the Δcoh experiment as well (Figure S3B), but it also helps explain the key discrepancy between the effects of Δcoh and microstimulation. For the μS experiments, an adjustment to the variance of the DV (δ_{σ^2}; Equation S8; Supplemental Experimental Procedures) largely accounted for the small decrease in sensitivity to motion on μS trials (Table 1), whereas this adjustment was negligible for the Δcoh manipulation, which did not affect sensitivity to motion (p = 0.68, logistic regression). Interestingly, the added variance also explains the apparent increase in confidence (decrease in T_s choices) associated with stimuli near the point of maximum ambiguity (Figure 2B, top, red curve). It may seem counterintuitive that an increase in variance (i.e., sensory noise) would predict an increase in confidence, but it is readily explained in our framework because dispersion of the DV away from the starting level causes more of its density to lie outside the region for opting out of the direction decision (see Rahnev et al., 2012 for a similar explanation).

We considered and rejected several alternative models, the most important of which allow for the possibility that μS induces a change in either the mapping of the DV to confidence or the criterion—applied to this mapping—for opting out. Specifically, we relaxed the assumption that the criterion (θ) was unaffected by μS. This is important because allowing μS to affect θ is tantamount to accepting that μS induces a change in neural activity that is processed qualitatively differently than activity caused by visual motion. We found that this extension was not justified for the main data set (see Supplemental Experimental Procedures). Note that the effects of μS on choice (Figure 2B, bottom, solid and dashed red curves) were predictions of the model, based on the fit to the rest of the data. The impressive agreement to data leads us to conclude that microstimulation did not alter the quantitative relationship between the neural representation of evidence and its mapping to a degree of certainty. In this way, the modeling exercise supports a unified theory of choice, confidence, and—by extension to previous work (Ditterich et al., 2003; Hanks et al., 2006)—reaction time.

DISCUSSION

Cortical microstimulation in behaving monkeys has long been a fruitful approach for exposing causal relationships between neural activity and perception (Bartlett and Doty, 1980; Doty, 1965). The power of the technique lies in its ability to link the functional properties of sensory neurons (e.g., direction selectivity in MT/MST) with psychophysical performance, as shown previously in several brain areas and tasks (Afraz et al., 2006; DeAngelis et al., 1998; Gu et al., 2012; Romo et al., 1998; Salzman et al., 1990). However, subjects in these studies typically report only a primary decision about the sensory stimulus. Here we have stimulated direction-selective neurons while allowing monkeys to report something additional about the decision process: their confidence, or lack thereof, in the choice. We found that μS affected confidence as if there were an offset to the visual evidence supporting the choice. The results support a quantitative framework in which confidence emerges from the same basic mechanism—bounded evidence accumulation—that successfully accounts for choice and reaction time. Thus, combined with previous studies (Ditterich et al., 2003; Salzman et al., 1990), there is now experimental evidence that links the activity of neurons in extrastriate visual cortex in a causal fashion to all three pillars of choice behavior (Vickers, 1979).

Some might wonder whether this is in any way surprising, given what we know from previous work (Bisley et al., 2001; Ditterich et al., 2003; Romo et al., 1998; Salzman et al., 1992). In fact, our study could have turned out differently because microstimulation induces a pattern of activity that is quite different from that caused by visual stimulation. This pattern could have failed to engage the same networks that normally read out sensory information for the purpose of computing confidence (Bach and Dolan, 2012; Barttfeld et al., 2013). Had μS induced incongruous changes in choice and confidence—or interfered with the improvement in sensitivity achieved by opting out of select trials—it would not have called into question previous findings of the effects of microstimulation on choice and reaction time. In short, the linking hypothesis tested here was by no means a foregone conclusion. Indeed, one might expect confidence to rest heavily on factors (e.g., metacognitive or affective) beyond operations on evidence and its conversion to a decision, especially considering the proposed role of higher-order structures (Kepecs et al., 2008; Komura et al., 2013; Rounis et al., 2010). Our findings do not directly conflict with these previous studies, but they do support a relatively straightforward mechanism for computing confidence in a perceptual decision (Kiani and Shadlen, 2009)—one that is tightly linked to the decision process itself.

Monkeys can be trained to detect microstimulation in a number of brain areas (Histed et al., 2013; Murphey and Maunsell, 2007), and we cannot rule out the possibility that they could detect its presence in the current study. What we do know is that they are unable or unwilling to counteract the effects of μS on choice (and confidence), even though doing so would increase reward rate. The presence of a compensatory choice bias against the stimulated direction (Figure 2B, bottom, rightward shift of blue curves; Salzman et al., 1992) further argues that monkeys do not differentiate between μS and non-μS trials. They do not adjust their strategy solely on μS trials but instead adjust their bias on all trials, reducing errors caused by μS at the cost of more errors when μS was absent. Even if μS were detectable, our results suggest that such detection did not disrupt the critical aspect of the decision process that establishes a level of confidence. From the perspective of downstream brain areas, the additional perturbation caused by μS of MT/MST is largely equivalent to a change in the neural activity produced by a visual stimulus.

We were able to explain the monkey's PDW behavior using the same bounded drift-diffusion model used to explain direction choices and RT in previous studies (Gold and Shadlen, 2007; Kiani and Shadlen, 2009; Link, 1992; Palmer et al., 2005; Smith and Vickers, 1988). The model exploits the association between the DV and the probability that a choice based on that DV will be correct, predicting a sure-bet choice when this probability is below a fixed threshold. This model can explain the principal effects of μS by treating it as an offset to the motion strength (Figures 2B and S4). Importantly, the model explains the assignment of confidence in a single decision based on an evolving DV. An alternative is that the monkey identifies the motion coherence and opts out with some frequency based on a learned association between coherence and the probability of being correct. However, this interpretation is contradicted by the improvement in performance—for all motion strengths and durations—on trials where the sure bet was offered and waived. The improvement implies that the brain is opting out selectively, based on a prediction that the decision reached during motion viewing is likely to be correct. The observation is also incompatible with other alternatives, such as selecting T_s following lapses of attention or evading the motion decision entirely on some fraction of trials (i.e., wishing for T_s and simply guessing if it does not become available). Importantly, the model explains the degree of improvement with impressive fidelity (Figure 2B, bottom, blue curves) and is able to predict the similar pattern on μS trials (Figure 2B, bottom, red curves) based on a fit to the other features of the data (see Supplemental Experimental Procedures). The fact that this improvement is preserved on μS trials is notable and could not have been predicted from previous work. It suggests that a rather sophisticated capacity to assess the reliability of sensory evidence is maintained despite the unnatural pattern of neural activity induced by μS.

In addition to shifting the sure-bet curve, μS also slightly reduced the peak rate of T_s choices (Figure 2B, top). The lack of such an effect in the Δcoh control experiment (Figure 2B, top, inset) suggests that this is a consequence of μS itself rather than any analysis method or incidental feature of the task, such as reward contingencies or the compensatory bias. The change in peak T_s frequency can be explained if we assume that μS affects both the signal and the noise of the sensory representation. In the context of bounded evidence accumulation, adding noise effectively increases the likelihood that the DV will diffuse away from zero (i.e., neutrality) and beyond the threshold for waiving T_s. This explanation is also consistent with the small decrease in sensitivity associated with μS. Note that an effect on noise is distinct from the proposed mechanism by which high-current stimulation reduces perceptual sensitivity (Figures 3 and S2; see Supplemental Experimental Procedures). The latter is believed to result from the spread of current to multiple columns with different preferred directions (Murasugi et al., 1993)—a dilution of signal rather than an increase in noise. Even low-current stimulation may spread across columns in some cases, but the changes in confidence that we observed suggest an effect on noise per se, the mechanism of which remains unknown. This subtle effect notwithstanding, a key conclusion from the model is that μS does not influence higher-level aspects of decision strategy, such as the internal mapping between the state of the accumulated evidence and the likelihood of making a correct choice.

It seems uncontroversial that signals in the visual cortex would affect both choice and certainty, but it is remarkable that the coupling should be so well explained by a common mechanism. After all, certainty and confidence invite consideration of temperament, mood, and subjective experience about the decision process itself (e.g., metacognition). Thus, it is noteworthy that the monkeys did not exercise the option to indicate that something was peculiar about the decision process on trials accompanied by μS. In effect, the monkeys have communicated just the opposite: they "wager" as if they experienced a change in the visual stimulus. Moreover, the high-current μS experiments (Figure 3A) reassure us that the monkey is in fact able to use PDW to report decreased confidence when it occurs. Of course, we do not know what monkeys experience subjectively when we stimulate the brain, nor can we interrogate the subjective feeling of certainty itself. That said, any neuroscientific investigation is unlikely to furnish this level of explanation. What seems certain is that a quantitative reconciliation of choice, RT, and confidence will provide a basis for extending the neurobiology of decision making to more complex situations in which confidence itself plays a critical role.

EXPERIMENTAL PROCEDURES

Behavioral Task

Two adult male rhesus monkeys (*Macaca mulatta*) were trained to perform a direction discrimination task with postdecision wagering (PDW), as described previously (Kiani and Shadlen, 2009). The task (Figure 1A) was to determine the net direction of motion in a circular patch of dynamic random dots. Motion could be in one of two directions separated by 180°, and difficulty was controlled by varying both the viewing duration (truncated exponential distribution, mean = 270 ms, range = 60–880 ms) and the percentage of coherently moving dots (motion coherence: 0%, 3.2%, 6.4%, 12.8%, 25.6%, or 51.2%). After acquiring central fixation, two direction-choice targets appeared on opposite sides of the fixation point (9°–12° eccentricity), followed by the random dot motion display. After motion offset, the monkey maintained fixation through a variable delay period (range = 500–1,000 ms), during which a third target (the sure-bet target [T_S]) appeared on a random half of trials.

Importantly, the monkeys could not predict whether T_s would appear until at least 500 ms after stimulus offset, strongly encouraging them to complete a direction decision on all trials. T_s differed in color and size from the direction-choice targets and was positioned at an angle perpendicular to the motion axis at 6°–8° eccentricity.

After the delay period, the fixation point disappeared, cueing the monkey to make a saccadic eye movement to one of the targets. When given the opportunity, the monkey could choose T_s and receive a guaranteed reward (drop of water or juice) or waive T_s and make the higher-stakes direction choice. Correct direction choices yielded a larger liquid reward than T_s choices, while errors resulted in a 5–6 s timeout. The ratio of T_s reward size to direction-choice reward size (0.75–0.82 for monkey I, 0.64–0.72 for monkey D) was chosen to encourage the animals to choose T_s approximately 50% of the time at the weakest motion strengths. The ratio was not adjusted during the course of an experiment.

Surgery and Neurophysiological Methods

All procedures were in accordance with National Institutes of Health guidelines and approved by the Institutional Animal Care and Use Committees at the University of Washington and Columbia University. Animals were implanted with a head post and recording chamber using aseptic surgical methods. Electrical microstimulation and multiunit recordings were made with tungsten electrodes (Alpha Omega, impedance = 0.5–2 MΩ measured at 1 kHz). Areas MT (n = 32 sites) and MST (n = 31) were identified using structural MRI scans and standard physiological criteria, as well as histological analysis in one animal. Stimulation sites were chosen based on strong direction selectivity and consistent tuning (across ∼200 μm of cortex) for the direction, speed, and size of the motion stimulus (see Supplemental Experimental Procedures for details).

Once we encountered an acceptable site, we positioned the electrode tip near its center and began the discrimination task. On a random half of trials, including both T_s-present and T_s-absent trials, microstimulation was delivered through the recording electrode using a Grass S88 stimulator with two PSIU6 optical isolation units (Grass Technologies). Stimulation trains consisted of square-wave, biphasic pulses with the following parameters: 5, 7.5, or 10 μA; 200, 250, or 333 Hz, 0.4 ms negative and 0.4 ms positive phase (negative phase leading). Within these ranges, no systematic effects of pulse amplitude or frequency were detected. The pulses began 40 ms after motion onset and stopped 40 ms after motion offset to account for visual response latency. For eight sites in one monkey, following the standard block of trials, an additional block was collected in which the amplitude of pulses was increased to 75 μA ("high-current") while all other parameters remained the same. Note that our pulse duration (0.4 ms) was longer than in previous studies by a factor of 1.33 (Ditterich et al., 2003) or 2 (Murasugi et al., 1993; Salzman et al., 1992).

We performed a set of control experiments with the same task design and stimuli, except that electrical microstimulation was replaced with an offset to the motion coherence assigned by the computer on a given trial (dubbed "added-signal" trials in Salzman et al., 1992). The coherence offset (Δcoh) was fixed for a given session and varied from 5%–40% coh across sessions (see Figures S1B and S1C).

Behavioral Data Analysis

We fit the direction choices to the logistic regression model given by:

$$P_{pref} = \left\{1 + e^{-Q}\right\}^{-1}, \quad Q = \beta_0 + \beta_1 I_E + \beta_2 C \qquad \text{(Equation 1)}$$

where P_{pref} is the probability of a preferred-direction choice, C is signed motion coherence, I_E is an indicator variable for μS (1 or 0 for trials with/without μS), β_0 is the overall bias, β_1 estimates the effect of μS on the direction choice, and β_1/β_2 expresses this in units of motion coherence. Fitting was performed by the method of maximum likelihood (binomial error), with SEs of the parameters obtained from the inverted Hessian matrix. SEs were used to compute t statistics and thereby evaluate the null hypothesis (e.g., $\beta_1 = 0$). Effects of μS on choice were similar between MT and MST (two-sample K-S test, p = 0.39); thus, we pooled the data from the two areas for all analyses.

To quantify the change in sensitivity associated with μS, we fit the logistic model given by:

$$P_{pref} = \left\{1 + e^{-Q}\right\}^{-1}, \quad Q = \beta_0 + \beta_1 I_E + \beta_2 C + \beta_3 I_E C \qquad \text{(Equation 2)}$$

where β_3 captures the effect on sensitivity. Similarly, the difference in sensitivity with and without T_s present (Figure 2B, bottom) was examined by replacing I_E in Equation 2 with an indicator term for T_s.

For some analyses, we fit the probability of sure-bet choices as function of signed coherence with a Gaussian function (Supplemental Experimental Procedures, Equations S1 and S2). Note that the smooth curves in Figures 2B, 3, S2A, and S3 were generated from the bounded accumulation models (see below), not logistic regression or Gaussian fitting. For additional methods and results related to the Gaussian fits, see Supplemental Experimental Procedures.

Model Fits and Predictions

Here we provide an intuitive overview of the model (Figure 4) and fits displayed in Figures 1B–1D, 2B, and S3B. Variables constituting degrees of freedom are identified by bold font and listed in Table 1. For mathematical details, see Supplemental Experimental Procedures.

We explain choice and PDW using a simplified one-dimensional diffusion process (Gold and Shadlen, 2007; Kiani and Shadlen, 2009; Link, 1992; Ratcliff and Rouder, 1998; Smith and Vickers, 1988) in which noisy evidence favoring either direction (and against the other) is accumulated for its display duration or until the process reaches an upper or lower bound, ±**B**. The bounds would explain reaction time in other contexts, whereas here they render decision times shorter than the display duration on some trials, and they affect the predicted accuracy on these trials (Kiani et al., 2008). The accumulation has a drift and a diffusion component. The latter is the accumulation of independent random numbers at each time step. The drift is a line with slope (drift rate) proportional to the motion coherence ($\boldsymbol{\kappa} C$), where the sign of C indicates direction. The sign of the accumulated evidence, termed the decision variable (x), determines the choice. Confidence, in turn, is the log odds that such a choice would be correct. It is a function of both x and time (i.e., the stimulus duration or the time that the accumulation reached a bound). The time dependence arises because the reliability of the evidence (motion strength) varies unpredictably across trials and is not explicitly known by the observer (Drugowitsch et al., 2014; Kiani and Shadlen, 2009). When T_s is offered, we assert that the monkey exercises or waives this option based on a criterion, $\boldsymbol{\theta}$, applied to the log odds of being correct (Figures 4B–4D). The model generates the expected probability of each option by propagating and integrating the probability density of the decision variable within different regions of this space, as partitioned by θ and the bounds (Figures 4C and 4D).

On trials with μS, we assume that the drift rate is offset by $\boldsymbol{\delta_C}$, equivalent to a change in the motion coherence, and allow for the possibility that the diffusion noise is also affected (offset by $\boldsymbol{\delta_{\sigma^2}}$). On all trials, the drift rate includes an additional offset term $\boldsymbol{\gamma}$ to account for the compensatory bias that arises in microstimulation experiments (Salzman et al., 1992).

We employed a simple parameterization and tiered fitting strategy designed to minimize the number of degrees of freedom of the model (see Supplemental Experimental Procedures). We pursued this strategy to guard against over fitting and to support intuitions about the neural mechanisms. The model furnishes the smooth curves in the analyses of the main data set (low current μS) and the Δcoh control (Figures 2B and S3B, respectively; see also Figure S4), as well as the dashed curves in Figures 1B–1D. A more elaborate model, also described in the Supplemental Experimental Procedures, was required to explain the effects of high current stimulation (Figures 3 and S2).

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, four figures, and one table and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2014.07.011.

AUTHOR CONTRIBUTIONS

M.N.S, C.R.F., and R.K. conceived and designed the experiments and performed the analyses. C.R.F. and R.K. collected the data. All authors wrote the paper.

ACKNOWLEDGMENTS

This work was supported by the Howard Hughes Medical Institute, National Eye Institute grant EY11378, NEI Core Grant P30EY01730, and National Center for Research Resources grant RR00166. We thank members of the lab for discussions, G. Horwitz for histology, and A. Boulet, K. Ahl, and K. Morrisroe for technical assistance.

Accepted: July 7, 2014
Published: August 7, 2014

REFERENCES

Afraz, S.-R., Kiani, R., and Esteky, H. (2006). Microstimulation of inferotemporal cortex influences face categorization. Nature *442*, 692–695.

Bach, D.R., and Dolan, R.J. (2012). Knowing how much you don’t know: a neural organization of uncertainty estimates. Nat. Rev. Neurosci. *13*, 572–586.

Bartlett, J.R., and Doty, R.W. (1980). An exploration of the ability of macaques to detect microstimulation of striate cortex. Acta Neurobiol. Exp. (Warsz.) *40*, 713–727.

Barttfeld, P., Wicker, B., McAleer, P., Belin, P., Cojan, Y., Graziano, M., Leiguarda, R., and Sigman, M. (2013). Distinct patterns of functional brain connectivity correlate with objective performance and subjective beliefs. Proc. Natl. Acad. Sci. USA *110*, 11577–11582.

Bisley, J.W., Zaksas, D., and Pasternak, T. (2001). Microstimulation of cortical area MT affects performance on a visual working memory task. J. Neurophysiol. *85*, 187–196.

DeAngelis, G.C., Cumming, B.G., and Newsome, W.T. (1998). Cortical area MT and the perception of stereoscopic depth. Nature *394*, 677–680.

Del Cul, A., Dehaene, S., Reyes, P., Bravo, E., and Slachevsky, A. (2009). Causal role of prefrontal cortex in the threshold for access to consciousness. Brain *132*, 2531–2540.

Ditterich, J., Mazurek, M.E., and Shadlen, M.N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. Nat. Neurosci. *6*, 891–898.

Doty, R.W. (1965). Conditioned reflexes elicited by electrical stimulation of the brain in macaques. J. Neurophysiol. *28*, 623–640.

Drugowitsch, J., Moreno-Bote, R., and Pouget, A. (2014). Relation between Belief and Performance in Perceptual Decision Making. PLoS ONE *9*, e96511.

Foote, A.L., and Crystal, J.D. (2007). Metacognition in the rat. Curr. Biol. *17*, 551–555.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Gu, Y., Deangelis, G.C., and Angelaki, D.E. (2012). Causal links between dorsal medial superior temporal area neurons and multisensory heading perception. J. Neurosci. *32*, 2299–2313.

Hampton, R.R. (2001). Rhesus monkeys know when they remember. Proc. Natl. Acad. Sci. USA *98*, 5359–5362.

Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat. Neurosci. *9*, 682–689.

Histed, M.H., Bonin, V., and Reid, R.C. (2009). Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. Neuron *63*, 508–522.

Histed, M.H., Ni, A.M., and Maunsell, J.H.R. (2013). Insights into cortical mechanisms of behavior from microstimulation experiments. Prog. Neurobiol. *103*, 115–130.

Kahneman, D., Slovic, P., and Tversky, A. (1982). Judgment under Uncertainty: Heuristics and Biases. (Cambridge: Cambridge University Press).

Kepecs, A., Uchida, N., Zariwala, H.A., and Mainen, Z.F. (2008). Neural correlates, computation and behavioural impact of decision confidence. Nature *455*, 227–231.

Kiani, R., and Shadlen, M.N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. Science *324*, 759–764.

Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J. Neurosci. *28*, 3017–3029.

Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., and Miyamoto, A. (2013). Responses of pulvinar neurons reflect a subject’s confidence in visual categorization. Nat. Neurosci. *16*, 749–755.

Lau, H.C., and Passingham, R.E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. Proc. Natl. Acad. Sci. USA *103*, 18763–18768.

Link, S.W. (1992). The Wave Theory of Difference and Similarity. (Hillsdale: Lawrence Erlbaum Associates).

Logothetis, N.K., Augath, M., Murayama, Y., Rauch, A., Sultan, F., Goense, J., Oeltermann, A., and Merkle, H. (2010). The effects of electrical microstimulation on cortical signal propagation. Nat. Neurosci. *13*, 1283–1291.

Middlebrooks, P.G., and Sommer, M.A. (2011). Metacognition in monkeys during an oculomotor task. J. Exp. Psychol. Learn. Mem. Cogn. *37*, 325–337.

Murasugi, C.M., Salzman, C.D., and Newsome, W.T. (1993). Microstimulation in visual area MT: effects of varying pulse amplitude and frequency. J. Neurosci. *13*, 1719–1729.

Murphey, D.K., and Maunsell, J.H.R. (2007). Behavioral detection of electrical microstimulation in different cortical visual areas. Curr. Biol. *17*, 862–867.

Palmer, J., Huk, A.C., and Shadlen, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. J. Vis. *5*, 376–404.

Rahnev, D.A., Maniscalco, B., Luber, B., Lau, H., and Lisanby, S.H. (2012). Direct injection of noise to the visual cortex decreases accuracy but increases decision confidence. J. Neurophysiol. *107*, 1556–1563.

Ratcliff, R., and Rouder, J.N. (1998). Modeling response times for two-choice decisions. Psychol. Sci. *9*, 347–356.

Romo, R., Hernández, A., Zainos, A., and Salinas, E. (1998). Somatosensory discrimination based on cortical microstimulation. Nature *392*, 387–390.

Rounis, E., Maniscalco, B., Rothwell, J.C., Passingham, R.E., and Lau, H. (2010). Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. Cogn Neurosci *1*, 165–175.

Salzman, C.D., Britten, K.H., and Newsome, W.T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. Nature *346*, 174–177.

Salzman, C.D., Murasugi, C.M., Britten, K.H., and Newsome, W.T. (1992). Microstimulation in visual area MT: effects on direction discrimination performance. J. Neurosci. *12*, 2331–2355.

Smith, P.L., and Vickers, D. (1988). The accumulator model of 2-choice discrimination. J. Math. Psychol. *32*, 135–168.

Smith, J.D., Beran, M.J., Couchman, J.J., and Coutinho, M.V.C. (2008). The comparative study of metacognition: Sharper paradigms, safer inferences. Psychon. Bull. Rev. *15*, 679–691.

Smith, J.D., Couchman, J.J., and Beran, M.J. (2012). The highs and lows of theoretical interpretation in animal-metacognition research. Philos. Trans. R. Soc. Lond. B Biol. Sci. *367*, 1297–1309.

Tolman, E.C. (1948). Cognitive maps in rats and men. Psychol. Rev. *55*, 189–208.

Vickers, D. (1979). Decision Processes in Visual Perception. (New York: Academic Press).