Early planning activity in frontal and parietal cortex in a simplified task

Chess Stetson[1,2] and **Richard A. Andersen**[1,2]

[1]*Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, California; and* [2]*Program in Computation and Neural Systems, California Institute of Technology, Pasadena, California*

Stetson C, Andersen RA. Early planning activity in frontal and parietal cortex in a simplified task. *J Neurophysiol* 113: 3915–3922, 2015. First published March 11, 2015; doi:10.1152/jn.00104.2014.— Cortical planning activity has traditionally been probed with visual targets. However, external sensory signals might obscure early correlates of internally generated plans. We devised a nonspatial decision-making task in which the monkey is encouraged to decide randomly whether to reach or saccade in the absence of sensory stimuli. Neurons in frontal and parietal planning areas (in and around the arcuate and intraparietal sulci) showed responses predictive of the monkey's upcoming movement at early stages during the planning process. Neurons predicted the animal's future movements several seconds beforehand, sometimes before the trial even began. These data cast new light on the role of the cerebral cortex in the action planning process, when the animal is free to decide on his own actions in the absence of extraneous sensory cues.

decisions; motor; parietal; planning; target

NEURAL FIRING IN THE CORTEX can reflect an animal's intended movement even before he makes it. This kind of predictive activity can be found in multiple areas around the brain, including the intraparietal sulcus (IPS; Gnadt and Andersen 1988), and areas around the arcuate sulcus (AS) such as the dorsal premotor area and frontal eye fields (Andersen et al. 1990; Chafee and Goldman-Rakic 1998; Cisek and Kalaska 2005; Tanne-Gariepy et al. 2002). However, the exact meaning of predictive activity is still being debated. Spatial attention tasks strongly engage these areas (Bisley and Goldberg 2003), and some have suggested this persistent activity in the IPS represents sensory accumulation (Gold and Shadlen 2000). However, it is also modulated by expectation (Janssen and Shadlen 2005) and expected value (Musallam et al. 2004; Platt and Glimcher 1999; Sugrue et al. 2004). Persistent activity also predicts explicitly nonspatial aspects of movement such as the effector (e.g., hand or eye) the monkey plans to use (Cui and Andersen 2007; Scherberger et al. 2005; Snyder et al. 1997).

As neuroscience has transitioned from studying the brain purely as an input-output machine toward understanding it as a system that generates behavior on its own, evidence has begun to emerge that persistent activity can represent several kinds of internally generated decisions (Curtis and Lee 2010). Benjamin Libet famously showed that early neural signals can predict when a movement will be made (Libet et al. 1983; in monkeys, Maimon and Assad 2006). Frontoparietal neurons can reveal where the animal intends to move (Raposo et al. 2014). However, the same neurons can also reveal how the animal intends to move (e.g., to reach vs. saccade; Cui and Andersen 2007).

The decision of how is independent of where. However, out of convention, nonspatial decisions are most often probed by presenting a spatial target to the animal (Cui and Andersen 2007). Visual targets are strong drivers of activity in the frontoparietal network, and in the context of a task where the animal expects a visual target, the neural resources required to prepare for the target presentation might be diverted away from an encoding of how the monkey plans to move. We reasoned that correlates of the monkey's decision might emerge at an earlier time if we removed the conventional visual cue, giving the neural representation of the future action more time to evolve in the cortex.

To detect early correlates of planning activity in the cortex, we employed a task in which the animal decides between two effectors (a reach or a saccade) in the absence of extraneous visual stimuli (Fig. 1, *A* and *B*). Importantly, the task design strives to make the monkey's actions unpredictable. If a movement plan is predictable (from an earlier movement or stimulus), then it cannot be characterized as a free choice or decision. However, near-random behavior, without any predictable correlation with earlier stimuli or choices, can be a proxy for free choice in an experimental context. To facilitate random, unpredictable behavior, we used the "matching pennies" paradigm (Barraclough et al. 2004). The matching pennies task records the monkey's previous choices and makes a prediction about the monkey's choice on each trial. The monkey is only rewarded if his choice was not predicted by the computer. By encouraging near-random behavior (i.e., free choice), any detectable neural correlate of the upcoming decision will be unlikely to represent a previous behavioral predictor.

Single and multiunit neural data were recorded in IPS and around AS in two animals (Fig. 1*C*). Neural correlates of the animal's upcoming movement emerged as far back as the beginning of the trial, before the animal even adopts fixation. We found that these neural correlates of the upcoming plan occur significantly earlier in time than those in a task with a visual target. These results add to the growing evidence that frontal and parietal regions do more than plan a movement appropriate to a sensory stimulus. Rather, they seem to be involved in generating movement plans in their own right.

METHODS

Two adult male rhesus monkeys (*Macaca mulatta*, *R* and *L*) participated in this study. All surgical and animal care procedures were conducted in accordance with National Institutes of Health guidelines and were approved by the California Institute of Technology Animal Care and Use Committee.

Task. In the task, monkeys were required to choose either a reach or saccade to a known target after holding a fixation spot for 1–2 s

without any additional stimuli. The fixation spot [a light-emitting diode (LED) housed inside a button] was placed centrally in the animal's field of view, at comfortable arm's length. The fixation spot appeared, after which the monkey acquired fixation by touching the spot and looking at it with his eyes. This fixation lasted for 1–2 s, after which the fixation spot disappeared. Disappearance of the fixation spot cued the monkey to make a movement, either a reach or saccade to a peripheral target. Peripheral targets were also LED buttons placed 4 in. away from the fixation spot to the right or left. For the first few trials of each trial block, a peripheral target was illuminated during fixation so that the monkey understood where to reach/saccade for the remaining trials (these initial trials were excluded from analysis).

The task discouraged the animals from creating a predictable sequence of choices by tracking the previous 5 choices (and also the previous 4, 3, 2, and 1), picking the most probable choice following all previous such sequences, and rewarding the monkey on a given trial only if he made the improbable choice (i.e., matching pennies; see Barraclough et al. 2004). Our algorithm did not explicitly discourage strategies featuring interactions between actions and reward (e.g., "win/stay, lose/switch"). However, it was possible to post hoc filter these trials so as to remove from the analysis any trials with predictable strategy formation.

Similar to the real-time matching pennies task, the post hoc trial filtering calculated the probability of a given reach or saccade choice given the trial history. The last five trials were considered and compared with all previous possible versions of the same five-trial sequence. All such five-trial sequences followed by a reach were counted and compared with those followed by saccades, and if one set was significantly larger (according to the binomial distribution at $P < 0.001$, $P < 0.01$, $P < 0.1$; see Fig. 4), the more frequent action was predicted. If the predicted action matched the actual action, that trial was excluded. The primary difference in post hoc filtering for self-cued strategies based on reward was that rather than counting sequences of five reaches and saccades, the algorithm counted sequences of five rewarded-reaches/nonrewarded-saccades vs. rewarded-saccades/nonrewarded-reaches. These strategic trials could then be filtered out, as could any trials in which the animals may have engaged in predictable sequential behavior despite the reward structure (such as a small number of segments when the monkey repeatedly chose saccade despite not being rewarded for it). After successful completion of a reach or saccade and an additional 200-ms wait, the visual target finally appeared to confirm the monkey's movement was registered and preventing his movements from drifting in space over the course of the task. Two hundred milliseconds after holding either hand or eye position on the confirmatory target, the animal was rewarded with 50% probability (assuming his reaches and saccades had been randomly chosen up until that point). This time of reward delivery or withholding is referred to as "previous trial end" and was followed by the reappearance of the fixation spot.

The average number of trials per recording day was 650. The visual task shown in Fig. 3 worked in a similar fashion, but after fixating for 500-1,500 ms, an intervening visual stimulus (a lit LED) was presented for 500 ms on either the left or right of the fixation spot followed by an additional 500- to 1,500-ms memory period during which the animal continued to fixate. The location of the stimulus cued the location for the movement but not the effector, which was still the monkey's own choice. The average number of trials per recording day was 260 for the visual experiment.

Recording techniques. IPS recordings were made near the bottom of the sulcus, and AS recordings were centered on the posterior and medial side of the AS, anterior and lateral to the dorsal premotor cortex but somewhat posterior and medial to the frontal eye fields. Some AS sites elicited eye or pinna movements in response to 20-μA stimulation, although to a lesser extent than expected from the frontal eye fields. Recordings were made using multielectrode drives (Thomas Recording; amplifiers and storage software by Plexon). Recorded neurons were not preselected by their visual tuning prop-

erties. Rather, on a given channel, the first neuron that could be held in isolation for the 30–60 min preceding the start of the experiment was chosen and maintained in isolation. Well-isolated single neurons were prioritized for recording, although in many cases neuronal recordings might represent multiunit activity from two or more neurons that could not be disambiguated. Many multiunit recordings had significant predictive power, such as those in Fig. 2, *A–C*. All neurons were isolated either in real-time via box selection (Plexon) or offline via clustering in principal components space. Qualitatively, multiunit signals were similar to single-unit signals during recording with audible single spikes rather than a high-frequency hash. Single-unit data were differentiated from multiunit data by both exceeding a waveform signal-to-noise ratio of 4 and having a peak in their interspike interval distribution of >2 ms.

Analysis techniques. Action potentials were convolved with a 50-ms width Gaussian, and time courses were averaged across all trials in which the monkey made either a reach or a saccade to produce the traces in Fig. 2. Spike counts were binned every 100 ms, aggregated according to whether they preceded a reach or saccade, and then compared via an unpaired *t*-test to produce significance values for Figs. 3–5. Note that no other procedure was required to determine the spatial tuning of individual neurons: they were simply assessed based on reach/saccade selectivity in each hemispace. The same procedure of comparing spike trains was repeated but randomized across reaches and saccades to produce the white baseline curves in Fig. 3. To produce the asterisks in Figs. 3–5, the earliest time at which a neuronal signal encoded the upcoming movement with a significance value of $P < 0.05$ was recorded, registered with respect to the nearest behavioral event (e.g., FIX ON, FIX AQ., and MVMT. START), and binned into nonoverlapping time segments (the segment boundaries are signified by white spaces in Figs. 3–5) centered on each of these events. Medians for times of first significance were then calculated across these concatenated distributions. The distributions of first-significance times across neurons were likewise inspected. There was a wide peak (>1 time bin) at or near the median of the distribution in all cases of this analysis except Fig. 5, *monkey R*.

All results shown in Figs. 3–5 are combined over two monkeys. However, these results are dominated by data from *monkey L*. In the data we were able to collect from *monkey R*, an insignificant number of IPS neurons showed tuning. This animal performed the task using the arm on the same side as the recording site (i.e., ipsilateral to the recording chamber and recorded hemisphere), which might also partially account for the weak results. However, enough neurons from the AS region were tuned throughout the planning period for us to justify a repeat of the experiment in *monkey L*, this time with a more complete set of neurons and using the contralateral arm to the recorded hemisphere. Except for Fig. 5, all data are shown from both animals combined. However, all conclusions from the combined data shown and discussed in Figs. 3–5 are true for *monkey L* in isolation. The conclusions that are significant in *monkey R* in isolation are the appearance of decision correlates in the early stages of fixation before the target would normally even be presented (Fig. 3, *A and B*), the existence of a significant number of prefixation decision correlates in the AS region (Fig. 3A; 7.8%; $P < 0.025$) according to the chance distribution determined by shuffling, i.e., the white traces in Fig. 3, and the stronger decision-making representation in AS compared with IPS (on average, 10% of AS neurons show predictive activity during the period between fixation onset and the cue to move, $P < 0.025$ according to the shuffled distribution; see Fig. 5).

RESULTS

The monkeys chose an approximately equal number of reaches and saccades over the course of the task and across task days (Fig. 1*B*). Occasional trial segments in which the monkey behaved nonrandomly (e.g., if a monkey repeatedly chose saccades for 20 trials despite getting no reward for them after

A

B

C

D

% REACH / REWARD

100

50

0 TRIALS 100

100

50

0 EXPERIMENTAL DAYS

E

AS
N=231

N=144 IPS
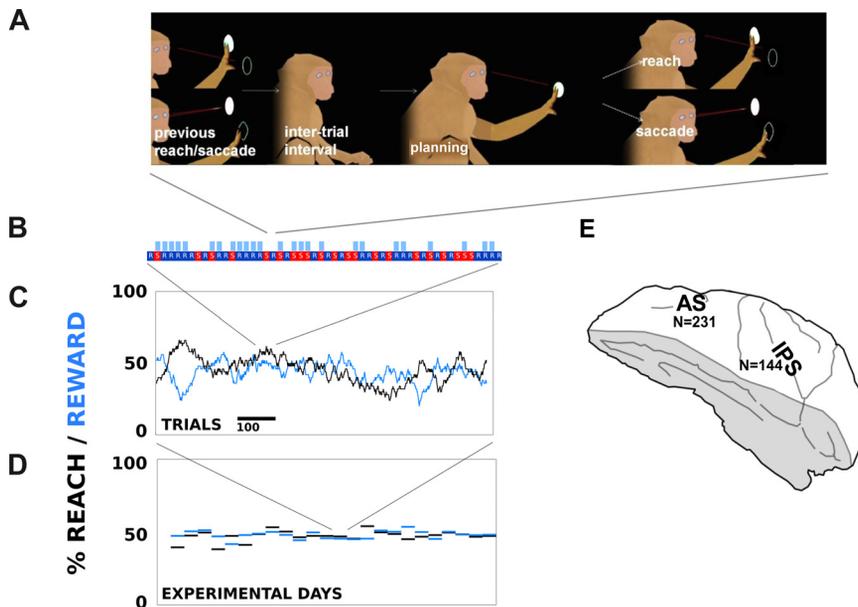
Fig. 1. *A*: after the conclusion of the previous trial, animals experienced a short intertrial interval (0.5–1 s) and then regained fixation (0.2–0.5 s), initiating the planning period of 1–2 s. When the fixation light disappeared, the monkey made either a reach or saccade and held the movement (0.2 s) until a light appeared on the target to indicate completion of the trial (0.2 s). Afterward, reward was delivered with ~50% probability, and the task restarted. *B*: a segment of *N* trials. Reward delivery is signified by blue boxes to the top right of each trial. R, reach; S, saccade. *C* and *D*: average reward and reach probabilities over a single day (*C*) and day by day (*D*). *E*: recorded regions. AS, arcuate sulcus; IPS, intraparietal sulcus.

the 3rd or 4th trial) were removed via post hoc filtering (see METHODS).

A total of 108 single units were recorded in the IPS site of *monkey L*, 61 single units in AS of *monkey L*, 16 single units in the IPS of *monkey R*, and 20 single units in the AS of *monkey R*. Additionally, 152 multiunit signals were recorded in the IPS of *monkey L*, 243 multiunits in the AS of *monkey L*, 99 multiunits in the IPS of *monkey R*, and 69 multiunits in the AS of *monkey R*. Neurons were classified as single-unit if they satisfied both a signal-to-noise threshold and a condition on their interspike interval distributions. Example neural signals are shown in Fig. 2. Figure 2*A* shows a multiunit recording that begins to differentiate around the time the fixation point appears and maintains increased firing throughout the planning period in advance of a reach, although the difference disappears by the time of the actual movement. Figure 2, *B* and *C*, shows multiunit signals that begin to differentiate between a planned saccade and reach as far back as the end of the previous trial (a reward or withheld reward after target acquisition; see METHODS) and continue to fire more throughout the planning period, all the way up until the saccade is actually made. Other neurons (data not shown) exhibited more exotic behavior, such as a parietal neuron that responded transiently to onset of the fixation spot, firing more if the animal was planning a saccade, and yet stopping its firing as the planning period progressed. Overall, 6% of signals significantly encoded the animal's upcoming plan before but not during the fixation period (corrected for multiple comparisons), and 4% significantly encoded the animal's upcoming plan during the fixation period but not during the movement period (as in Fig. 2*C*). Otherwise, the encoding of the upcoming movement by the neural signals progressed in a generally increasing pattern over time.

These examples indicate a variety of neuronal responses that reflect the monkey's ultimate movement choice at different epochs throughout the course of the trial. How does the ensemble activity develop? Each neuronal signal in Fig. 2 differentiates significantly between a planned reach and saccade at different times throughout the planning process. For each time bin, Fig. 3*A* aggregates the instances that a given neural signal differentiated between planned reach and saccade at a level of $P < 0.05$ and shows them plotted as a percentage of all neurons. The overall predictive signal grows as the previous trial concludes, and by the time the fixation point appears, 20% of all neural signals significantly differentiate the future movement. The asterisk in Fig. 3*A* represents the time that neuronal signals 1st differentiate the movement plan (median, across the neuronal population, of times when each neural signal 1st significantly encodes the upcoming movement). This median 1st-response time appears just after the fixation spot turns on and before the monkey even acquires fixation to start the planning period. The overall percentage of neuronal signals encoding the decision remains sustained throughout the planning period, increasing slightly around the time of the GO cue and increasing more when the animal actually moves. These results suggest that in the absence of extraneous stimuli appearing during fixation, neural firing correlates with the upcoming movement decision as early as the end of the previous trial.

In contrast, in a similar experiment with the same fixation spot, but also a visual stimulus that appeared during fixation in one hemispace or the other (Fig. 3*B*), neuronal encoding of the plan does not appear until after the visual stimulus. Also, the median appearance time is not until well into the memory period, the period of continued fixation after the target disappears. The rules of the task are such that the monkey must move to the presented target, although whether he reaches or saccades toward it is still his own decision. As the reach and saccade decision is independent of the direction of the visual stimulus, in principle the decision could be detected in advance of the visual stimulus. However, the data show that the cortical signals do not predict the monkey's choice until after the stimulus disappears. The level of significance rides slightly above chance (8%, or 3% above chance between the end of the previous trial and the cue to move in the current trial compared with 16% above chance in the no-visual-stimulus case shown in Fig. 3*A*). This trace level of predictive power may mirror earlier-than-expected decision-making signals hinted at by ear-
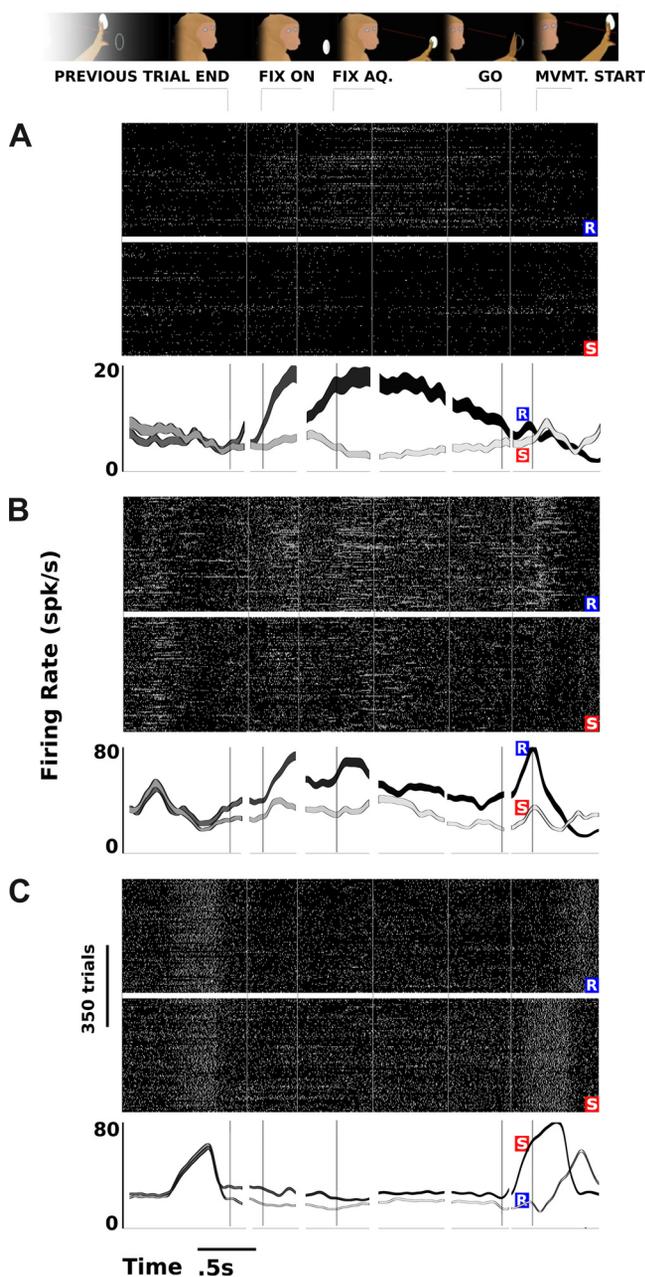
Fig. 2. Example neuronal signals. Rastergrams (above) show action potentials in time and across trials. Rastergrams show spiking events in 3-ms bins (1 pixel per bin, 9-pixel-high vertical lines; best viewed at 100 or 50% resolution), which are then convolved with a 50-ms Gaussian to produce the averaged time courses (below). Neuronal signals differentiated upcoming movements at a variety of task phases preceding the movement. $P < 0.05$, 2-tailed t-test between neural firing time courses in reach and saccade conditions, corrected for multiple comparisons (see METHODS). Each time segment corresponds to a different phase of the task. For each neuronal signal, the black trace corresponds to the action eliciting the strongest firing during the movement period. Signal traces fade from gray, at the left, to signify the evolution of the movement plan. Rastergrams and activity traces are labeled by the corresponding effector used (saccade or reach). Recording *A* comes from the IPS region (representing a single unit), and *B* and *C* come from the AS region (*monkey L*). Bars around time course traces are standard error of the mean. Numbers of trials for *A/B/C* reach/saccade are 350/350, 235/249, and 256/304. FIX ON, FIX AQ., GO, and MVMT. START are behavioral events. spk/s, Spikes per second.

lier studies (see Cui and Andersen 2007) and could represent a suppressed version of the same phenomenon shown in Fig. 3*A*. The numbers of neural signals of each isolation type recorded for this task are as follows: *monkey L*, single unit, 72 IPS, 87 AS; *monkey R*, single unit, 259 IPS, 218 AS; *monkey L*, multiunit, 71 IPS, 43 AS; *monkey R*, multiunit, 248 IPS, 217 AS. Note also that the differences in absolute level of detected selectivity are attributable to the smaller average number of trials in this task. To correct for the possibility that the median values were biased by this signal-to-noise difference, the median in Fig. 3*A* is calculated by leaving out enough neural signals with a highly significant *P* value that the peak percentage of significant neuronal signals at the time of movement start is just less (<30%) than the peak for the visual task. These data show that in a comparable task, in which an extra cue is added that affects the movement direction but not the effector decision, the effector decision is nevertheless obscured until after the cue has disappeared.

The dotted and dashed lines in Fig. 3 show the same data, broken out by isolation type (single vs. multiunit). Overall, the frequency of significant predictive power is largely the same. However, the multiunit data have slightly more predictive power than single units in the target-absent task (Fig. 3*A*), possibly indicating that neurons that can be maintained in isolation for long periods of time (with large waveforms and consistent firing rates across task phases) tend to report on animal's decisions with less frequency. These data suggest that with respect to decision correlates, there are minimal differences between single- and multiunit neural records.

The need to maximize trial number generally necessitated that the experiments in Fig. 3 be done separately. However, we did record a small number of neurons on target-present and target-absent tasks. There were 50 neuronal signals (11 single-unit) recorded from the AS of *monkey R* over 10 recording days with an average of twice as many trials for the former. This particular set of neuronal signals encoded the animal's upcoming effector decision more weakly (10% of the time in the target-absent task but only at chance level during the target-present task, 4%, likely attributable to the lower trial number). However, these neuronal signals were highly influenced by the visual target in the target-present task, firing an average of 70% over baseline in response to target presentation. These data indicate that the same neurons that respond to visual targets can also encode an animal's effector decisions in the absence of visual targets.

These results show that in the absence of a visual stimulus, neurons in the IPS and AS may encode the upcoming effector plan at times as early as the conclusion of the previous trial. Why not even earlier? It seems most likely that the events at the end of the previous trial, delivery or withholding of a reward, act as a cue for the monkey to begin making his subsequent decision. In the same way as the visual stimulus acted as a cue (Fig. 3*B*), the removal of which revealed earlier decision-making (Fig. 3*A*), we now seek to find a post hoc way to remove the cueing effect of delivery/withholding of reward.

Theoretically, the matching pennies task discourages the monkey from deciding on a preformed sequence of movements (Barraclough et al. 2004), such that in the best case, the monkey's choices are rewarded at random with 50% probability. However, even though the monkey does not execute a sequence, he could use the random value of this reward to cue his subsequent action (i.e., for a rewarded reach, produce a
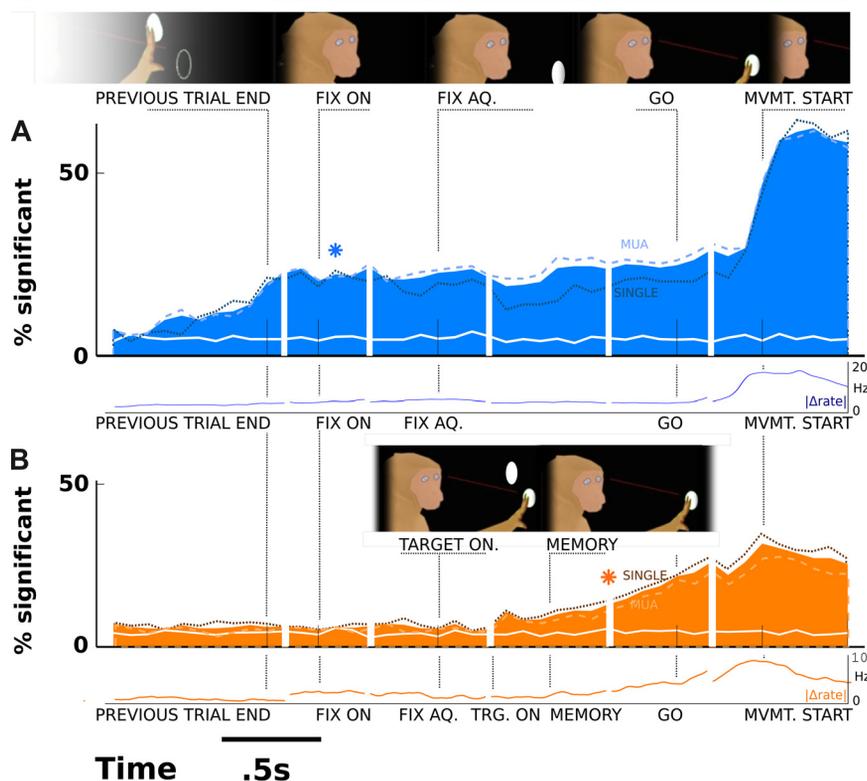
Fig. 3. Ensemble time course in the simple choice task compared with a task with an intervening visual stimulus. *A*: the percentage of neuronal signals in the population with significant tuning over time (bins are 50 ms) in the simple task. The white line corresponds to the same data but randomly shuffled across trials, simulating the results expected by chance. *Median of the times when significant tuning 1st appears, calculated across neurons (corrected; see METHODS). Dotted lines represent the same analysis but for single-unit data only; dashed lines represent multiunit activity (MUA; differentiated by waveform signal-to-noise and interspike interval; see METHODS). *B*: in a similar task, where the only addition is an intervening visual stimulus, significant tuning does not begin until a much later time. These differences in the overall time courses may suggest that the addition of a visual stimulus interferes with the capacity for IPS and AS neurons to encode movement plans at an early stage even when the movement plan is an effector decision, completely independent of the visual stimulus.

subsequent reach; otherwise saccade; and vice versa). Barraclough et al. (2004) created a task to discourage this kind of win-stay, lose-switch strategy. Instead, we choose to use a technique of post hoc filtering out such trials (see METHODS). In this way, we can isolate those trials that do not conform to a reward-cued strategy, the less predictable decisions, and compare them with the filtered-out trials, which will represent highly strategy-based decisions. In Fig. 4, we choose a significance threshold of $P = 0.1$, filtering out 44% of trials (the strategy-based trials) and leaving 56% remaining (the more random trials). The filtered and remaining trials now represent, in effect, two separate subtasks, which can have greater or lesser and earlier or later correlations of spiking with decisions.

The results of filtering are a general reduction of the proportion of neurons encoding the plan in the nonstrategy case (filtered) compared with strategy-only and also a relatively
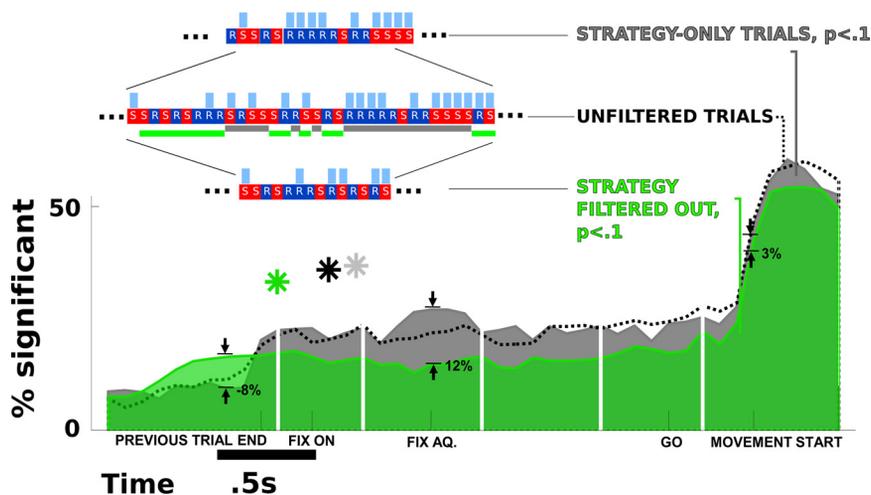
Fig. 4. Planning vs. decision-making. Whereas the task prevented the monkey from being rewarded for predetermined action sequences, he could still execute a strategy such as "win/stay, lose/switch" (see main text). By filtering out trials predictable by such a strategy (*inset*), only unpredictable (green) are preserved. This kind of trial filtering can reveal the difference between the neural encoding of a strategic plan (the filtered-out or strategy-only trials) and less-predictable decisions (green, trials filtered at $P < 0.1$). At the time that fixation is acquired, filtering out strategic trials reduces the number of predictable trials by 12% of neurons (54% of the unfiltered population, black, at that time point). However, by the time of the movement onset, the reduction due to filtering has reduced to 3% (6% of the significance level of unfiltered population). Nonetheless, trial filtering may also reveal a small amount of very early predictive capacity, as indicated by the higher level of predictive activity (8%, 66% of the unfiltered significance level) in the filtered case, before the end of the previous trial, which acts as a cue for the strategic decision (asterisks are median time of 1st tuning across neurons as in Fig. 3; gray, strategy-only; black, unfiltered; green, strategy filtered-out).

greater amount of movement-period encoding compared with planning-period encoding in the sense that the reduction due to filtering is less during the movement period than during the planning period. Figure 4 shows that after filtering, the percentage of neural signals significantly encoding the plan by the time of the fixation onset is lower by 12% of neuronal signals (54% of the unfiltered magnitude) compared with the strategy-only trials but only lower by 3% of neuronal signals (6% of the unfiltered case) during the movement itself. Note that features like the slight bump in significance around fixation acquisition in the strategy-only trials are possible because some neurons behave much more clearly in the isolated trials than overall. These results suggest that behaviors that are predictable by a strategy are also more strongly encoded by neural activity.

On the other hand, the effect of filtering promotes a slightly increased proportion of predictability very early in the trial. For trials without strategy, neural signals become more likely by 8% of neuronal signals (66% of the unfiltered percentage) to encode the upcoming movement decision. The median times of first significant effector encoding also reflect this time difference (note, these were not corrected for overall significance compared with the target experiment, as in Fig. 3). Therefore, it may be that nonstrategic trials, which do not use the reward delivery/withholding as a cue, have even earlier decision-making activity.

As reward is known to influence neurons in both IPS and AS, we sought to determine whether reward had a modulatory effect on the ability to decode reaches vs. saccades. As reward can enhance neural activity, a reward could result in greater subsequent predictive power in the spiking signal. We repeated the analysis from Fig. 3 but further separating trials by whether they followed rewarded or unrewarded trials. Between the end of the previous trial and the cue to move in the current trial, the average percentage of neuronal signals significantly predicting

the upcoming movement was 23% following an unrewarded trial and 20% following a rewarded trial (compared with 21% in the original reward-agnostic case shown in Fig. 3). Repeating the same technique of differentiating by previous reward but now applied only to the nonstrategic trials described in Fig. 4, we find a larger separation, although in an unexpected direction. Following reward during sequences of nonstrategic trials, only 14% of neuronal signals significantly predicted the upcoming movement as opposed to 19% following no-reward. Thus there is no evidence to suggest that the delivery of reward enhances the power of the decision-related signal. On the contrary, reward may weaken it.

In Fig. 5, the data are partitioned by recorded brain region and monkey. Whereas *monkey R*'s data are significantly weaker due to experimental constraints (see METHODS), both monkeys have qualitatively similar differences between cortical areas. In *monkey R*, AS neurons significantly encode the upcoming decision, but IPS neurons do not (apart from a small, very early bump). In *monkey L*, IPS neuronal signals have significant encoding power, but AS neuronal signals are still more likely to encode the movement decision as the time of the movement approaches. However, at very early times, the two populations are nearly equally likely to encode the decision, and the median time at which the decision is first encoded is the same across IPS and AS. These data suggest that although there are differences in decision correlates between AS and IPS, with AS arguably more predictive of imminent movements, both areas show early planning activity related to an upcoming decision.

DISCUSSION

To probe the extent to which neurons predict how an animal will move, we recorded extracellularly in the AS and IPS during a task without extraneous stimuli and examined neuro-

Fig. 5. Comparison across cortical areas (AS, gray; IPS, blue). The *top* image shows results for *monkey L*, and the *bottom* for *monkey R*, where data were recorded on the ipsilateral side to the hand being used. In both monkeys, a greater frequency of predictive activity appears in the AS site as the monkey prepares for the GO cue (33% of neurons in AS compared with 19% in IPS). However, in *monkey L* (top), the early component of the predictive activity is more similar between areas (12.5% in AS and 12% in IPS by the time of fixation acquisition), as is the median time of earliest detectable tuning (asterisks, calculated as in Figs. 3 and 4). *Monkey R* has significant predictive activity in the AS site only. Asterisks representing median times of 1st significant tuning are presented in parentheses for *monkey R* because they did not satisfy the condition of a wide peak in the distribution of 1st responses (see METHODS).
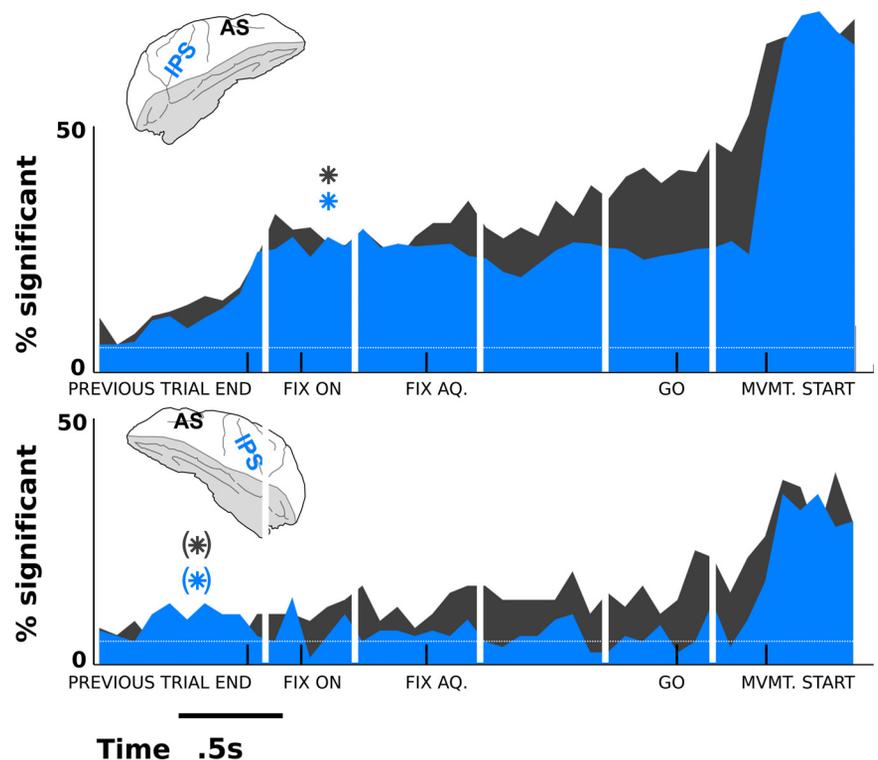
nal responses leading up to a reach or saccade the monkey had chosen to make. We found that neurons in cortical areas AS and IPS reveal a monkey's movement decision several seconds in advance, as early as the end of the previous trial. In comparison, during a task with a visual target stimulus, neural correlates do not appear until after the target disappears. In the targetless task, the cue for decision-making to start seems to backtrack to the previous reward, but when reward-influenced decisions are filtered out, neural correlates of the decision become detectable even earlier. Whereas the neural correlates of the decision are somewhat stronger in AS as the movement time approaches, the predictive capacity of AS and IPS were similar during the prefixation phase. These results indicate that when behavioral cues unrelated to the decision are removed, neural correlates of an upcoming movement decision can be detected further back in time in movement planning areas AS and IPS.

Why does the plan not appear earlier in the context of a task with a visual stimulus? Effector choice is modulated by target direction in many neurons (Snyder et al. 1997), and these areas also encode expectation (Janssen and Shadlen 2005). It could be that the combination of these two factors obscures the plan before the directional uncertainty is resolved. Alternately, the monkey may postpone the decision itself until after uncertainty about the target is resolved. At a minimum, it should be clear that decision correlates in the no-target task are unrelated to spatial attention (Bisley and Goldberg 2003), as the decision is nonspatial in nature.

There remains some question about whether these data indicate storage of a movement plan or something more akin to an evolving decision (Gold and Shadlen 2000). We decided to use a simple version of matching pennies task ("Type 1" from Barraclough et al. 2004). This version discourages predictable action sequences but not strategies in which the animal uses randomness inherent in the task (reward delivery/withholding) to cue his own future behavior. The post hoc filtering of strategic trials in Fig. 4 suggests that the self-cued decisions (gray trace) have strong neural correlates. Compared with just making a random decision when the movement is cued, it may seem energetically unfavorable that the monkey should *1*) store the previous action and *2*) its resultant reward and then *3*) combine them to generate an instruction for himself that he would *4*) store for seconds in persistent electrical activity until he next makes an action. However, as the present data seem to support the latter explanation, they may point to a need to reassess the role of randomness in decision-making theory.

On the other hand, a residual decision correlate remains (Fig. 4, green) after filtering out strategic trials. This signal begins at an earlier time than the reward-cued correlate (gray). Future experimentation will tell whether there is any limit on how far back in time neural activity could predict a decision in the absence of other stimuli.

The existence of early decision-making activity in parietal and frontal planning areas adds to an emerging awareness that these areas do more than encode working memory or prepare actions (Janssen and Shadlen 2005) that may have originated elsewhere. The earlier we find evidence of decision correlates in these areas, the more likely it seems that they themselves are involved in generating the decision. Moreover, because the decision correlates appear before fixation but persist throughout, the encoding of the decision may be independent of the role of these areas in satisfying immediate behavioral necessities like achieving and maintaining fixation. In other words, the decision or plan can persist in parallel to the other functions of these brain areas in the moment, and whereas on an individual basis many neurons are strongly modulated by the task phase (Fig. 2), in aggregate there is qualitatively little difference between the pre- and postfixation levels of decision encoding (Fig. 3*A*). Thus even as the transient requirements of successive behavioral states change and evolve, these cortical regions encode a parallel representation of the future movement plan the animal has decided to execute.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

AUTHOR CONTRIBUTIONS

C.S. and R.A.A. conception and design of research; C.S. performed experiments; C.S. analyzed data; C.S. and R.A.A. interpreted results of experiments; C.S. prepared figures; C.S. drafted manuscript; C.S. and R.A.A. edited and revised manuscript; C.S. and R.A.A. approved final version of manuscript.

REFERENCES

Andersen RA, Asanuma C, Essick G, Siegel RM. Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol* 296: 65–113, 1990.

Badre D, Hoffman J, Cooney JW, D'Esposito M. Hierarchical cognitive control deficits following damage to the human frontal lobe. *Nat Neurosci* 12: 515–522, 2009.

Barraclough DJ, Conroy ML, Lee D. Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci* 7: 404–410, 2004.

Bisley JW, Goldberg ME. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299: 81–86, 2003.

Chafee MV, Goldman-Rakic PS. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol* 79: 2919–2940, 1998.

Cisek P, Kalaska JF. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45: 801–814, 2005.

Cui H, Andersen RA. Posterior parietal cortex encodes autonomously selected motor plans. *Neuron* 56: 552–559, 2007.

Curtis CE, Lee D. Beyond working memory: the role of persistent activity in decision making. *Trends Cogn Sci* 14: 216–222, 2010.

Desmurget M, Reilly KT, Richard N, Szathmari A, Mottolese C, Sirigu A. Movement intention after parietal cortex stimulation in humans. *Science* 324: 811–813, 2009.

Gnadt JW, Andersen RA. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70: 216–220, 1988.

Gold JI, Shadlen MN. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.

Janssen P, Shadlen MN. A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8: 234–241, 2005.

Libet B, Gleason CA, Wright EW, Pearl DK. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain* 106: 623–642, 1983.

Maimon G, Assad JA. Parietal area 5 and the initiation of self-timed movements versus simple reactions. *J Neurosci* 26: 2487–2498, 2006.

Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA. Cognitive control signals for neural prosthetics. *Science* 305: 258–262, 2004.

Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. *Nature* 400: 233–238, 1999.

Raposo D, Kaufman MT, Churchland AK. A category-free neural population supports evolving demands during decision-making. *Nat Neurosci* 17: 1784–1792, 2014.

Scherberger H, Jarvis MR, Andersen RA. Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron* 46: 347–354, 2005.

Snyder LH, Batista AP, Andersen RA. Coding of intention in the posterior parietal cortex. *Nature* 386: 167–170, 1997.

Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science* 304: 1782–1787, 2004.

Tanne-Gariepy J, Rouiller EM, Boussaoud D. Parietal inputs to dorsal versus ventral premotor areas in the macaque monkey: evidence for largely segregated visuomotor pathways. *Exp Brain Res* 145: 91–103, 2002.