

Appendix A

Errington, S. P., & Schall, J. D. (2020). Express saccades during a countermanding task. *Journal of Neurophysiology*, *124*(2), 484-496.

Published July 15th 2020

RESEARCH ARTICLE | *Higher Neural Functions and Behavior*

Express saccades during a countermanding task

Steven P. Errington and Jeffrey D. Schall

Department of Psychology, Vanderbilt Vision Research Center, Center for Integrative & Cognitive Neuroscience, Vanderbilt Brain Institute, Vanderbilt University, Nashville, Tennessee

Submitted 18 June 2020; accepted in final form 13 July 2020

Errington SP, Schall JD. Express saccades during a countermanding task. *J Neurophysiol* 124: 484–496, 2020. First published July 15, 2020; doi:10.1152/jn.00365.2020.—Express saccades are unusually short latency, visually guided saccadic eye movements. They are most commonly observed when the fixation spot disappears at a consistent, short interval before a target spot appears at a repeated location. The saccade countermanding task includes no fixation-target gap, variable target presentation times, and the requirement to withhold saccades on some trials. These testing conditions should discourage production of express saccades. However, two macaque monkeys performing the saccade countermanding task produced consistent, multimodal distributions of saccadic latencies. These distributions consisted of a longer mode extending from 200 ms to as much as 600 ms after target presentation and another consistently less than 100 ms after target presentation. Simulations revealed that, by varying express saccade production, monkeys could earn more reward. If express saccades were not rewarded, they were rarely produced. The distinct mechanisms producing express and longer saccade latencies were revealed further by the influence of regularities in the duration of the fixation interval preceding target presentation on saccade latency. Temporal expectancy systematically affected the latencies of regular but not of express saccades. This study highlights that cognitive control can integrate information across trials and strategically elicit intermittent very short latency saccades to acquire more reward.

NEW & NOTEWORTHY A serendipitous discovery that macaque monkeys produce express saccades under conditions that should discourage them reveals how cognitive control can adapt behavior to maximize reward.

cognitive control; foreperiod; reward; temporal predictability

INTRODUCTION

Saccade latency is a manifestation of visual, motor, and cognitive processes (Carpenter 1988). Saccade latency and dynamics are influenced profoundly by reward expectancy and value. Saccades to stimuli associated with higher reward typically are more accurate and have faster peak velocities and shorter latencies relative to unrewarded stimuli (Milstein and Dorris 2007; Takikawa et al. 2002; Vullings and Madelain 2018, 2019). Of course, reward contingencies are integral in learning behaviors, and previous work has highlighted that production of express saccades is learned over time (Bibi and Edelman 2009; Fischer et al. 1984; Jóhannesson et al. 2018; McPeck and Schiller 1994; Paré and Munoz 1996).

Saccade latency is also influenced by the reliability of the timing of events. Behavioral testing typically includes an

interval between the presentation of a warning stimulus and the presentation of a target stimulus, known as the foreperiod. The passage of time within a foreperiod can convey information about when to expect the target. Sampling foreperiods from a uniform rectangular distribution results in an aging distribution with the conditional probability of the target appearing increasing as the foreperiod elapses. Conversely, nonaging foreperiods have an exponentially declining probability of elapsing, resulting in a constant conditional probability. Response times, including saccade latencies, are typically quicker following predictable, longer foreperiods (Ameqrane et al. 2014; Correa and Nobre 2008; Drazin 1961; Näätänen 1970; Niemi and Näätänen 1981; Thomaschke et al. 2011). Saccade latencies can become extremely short following a brief (~200 ms) foreperiod predictably coupled with removal of the visual fixation stimulus (Saslow 1967). In this gap paradigm, many saccades are initiated with latencies of ~80 ms in macaques and ~100 ms in humans (Boch and Fischer 1986; Boch et al. 1984; Fischer and Boch 1983; Schiller et al. 2004). Because these latencies approach the lower limits imposed by conduction delays in the oculomotor system, these are known as express saccades.

The saccade countermanding task has been widely used to investigate response inhibition (Cabel et al. 2000; Colonius et al. 2001; Godlove and Schall 2016; Hanes and Carpenter 1999; Hanes and Schall 1995; Kornyló et al. 2003; Morein-Zamir and Kingstone 2006; Thakkar et al. 2011, 2015; Walton and Gandhi 2006; Wattiez et al. 2016). In this task, subjects fixate a central spot and following a variable amount of time make a saccade to a target stimulus at one of two fixed, spatially separated locations (no-stop trials) presented simultaneously with the disappearance of the fixation dot. On a minority of trials, the fixation point reappeared instructing the monkeys to cancel their planned saccade to the peripheral target (stop-signal trials). Even though express saccades have been reported under overlap conditions (Amatya et al. 2011; Boch and Fischer 1986; Knox and Wolohan 2015), multiple aspects of the stop-signal task should discourage production of express saccades. First, because successful performance depends on balancing inhibition and initiation, saccade latencies are typically slower than those observed in other response tasks (Verbruggen and Logan 2009). Indeed, saccade latency increases with the fraction of stop-signal trials (Emeric et al. 2007). Second, whereas express saccade production is facilitated by consistent spatiotemporal target presentation, target location and timing are randomized in this task design. Finally, express

Correspondence: J. D. Schall (jeffrey.d.schall@vanderbilt.edu).

saccades are most common in gap task conditions that encourage release of fixation. In contrast, successful saccade countermanding performance encourages stricter control over visual fixation.

Here we report a serendipitous observation of frequent express saccades produced by monkeys performing a saccade countermanding task. To explore why monkeys produce express saccades in this task, we simulated performance according to the Logan and Cowan (1984) race model. We found that producing a fraction of express saccades can increase reward rate. To verify that express saccade production was motivated by reward contingencies, reward for producing express saccades was eliminated. When a minimum saccade latency was enforced for reward, monkeys stopped producing express saccades. We also quantified how saccade latency and express saccade production was affected by the temporal predictability of target presentation. When target presentation could be anticipated, regular but not express saccade latencies decreased. When monkeys performed the countermanding task consistent with the assumptions of the Logan and Cowan (1984) race model, they produced more express saccades than when they performed the task in violation of the assumption. Being incidental findings, their interpretation must be cautious, but they indicate differences in the mechanisms responsible for regular and express saccades, suggest informative future experimental designs, and highlight the range of operation of cognitive control.

METHODOLOGY

Animal Care

Data were collected from three male bonnet macaques (*Macaca radiata*, 6.9 to 8.8 kg) and one female rhesus macaque (*Macaca mulatta*, 6.0 kg). Animal care exceeded policies set forth by the USDA and Public Health Service Policy on Humane Care and Use of Laboratory Animals and all procedures were carried out with supervision and approval from the Vanderbilt Institutional Animal Care and Use Committee. Titanium head posts were surgically implanted to facilitate head restraint during eye tracking. Surgical methods have previously been described in detail (Godlove et al. 2011).

Data Acquisition

Experiments were carried out in darkened, sound-attenuated rooms. During testing, monkeys were seated comfortably 43 to 47 cm from a CRT monitor (~48 × 38°, 70 Hz) in enclosed polycarbonate and stainless-steel primate chairs and head restrained using surgically implanted head posts. Stimulus presentation, task contingencies related to eye position, and delivery of liquid reinforcement were all under computer control in hard real time (TEMPO, Reflective Computing, Olympia, WA). With the exception of the 70-Hz screen refresh rate, task timing was controlled at 500 Hz. Stimulus sizes and eccentricities were calculated automatically by the stimulus presentation program based on subject to screen distance to allow for increased precision between primate chairs and recording room setups. Stimuli were presented using computer-controlled raster graphics (TEMPO Videosync 640 × 400 pixel resolution, Reflective Computing, Olympia, WA). Stimuli had a luminance of 3.56 cd/m² (fixation point and stop-signals) or 2 cd/m² (targets) on a 0.01 cd/m² background.

Behavioral Task

Behavior and electrophysiological signals were recorded during the countermanding (i.e., stop-signal) task (Fig. 1). Additional details

about the behavioral training regime and task have been described previously (Hanes et al. 1998; Hanes and Schall 1995). Data from *monkey F* was recorded from the first countermanding session. Data from *monkey H* was recorded after he was well trained on the task. After observations made in an initial 110 sessions, 31 additional sessions were recorded from *monkey F* to compare the effect of reward contingencies on express saccades. In these sessions, saccade latencies were monitored online. In the first 13 sessions, correct express saccades (i.e., express saccades on no-stop trials) were rewarded. In the following 18 sessions, reward for express saccades was eliminated. Results from this data set were compared against data recorded several years later from *monkeys Eu* and *X*. Data from these monkeys were collected after both were well trained on the task and neither was rewarded for generating express saccades at any point in their training.

Trials were initiated when monkeys fixated a centrally presented square that subtended 0.34° of visual angle. After a variable foreperiod, the center of the central fixation point was extinguished, leaving a white outline. A target subtending 3° of visual angle simultaneously appeared at 10° to the left or right of the fixation. For two monkeys (*F* and *H*) foreperiods were randomly sampled from a uniform distribution. For two other monkeys (*Eu* and *X*) foreperiods were randomly sampled from an approximately nonaging, exponentially decaying distribution. On no-stop trials (Fig. 1, top), no further visual stimuli were presented. Monkeys were required to make a saccade to the target and hold fixation to obtain reward. Correct trials were rewarded with several drops of juice. On a proportion of trials, the center of the fixation point was reilluminated after a variable delay providing a “stop signal” that instructed the monkeys to cancel their impending eye movements and maintain central fixation (Fig. 1, bottom). The average proportion of stop trials varied across monkeys for incidental reasons: *monkey F*: 37%, *monkey H*: 40%, *monkey Eu*: 52%, and *monkey X*: 50%.

On stop-signal trials, two trial outcomes were then possible. If monkeys successfully withheld the eye movement and maintained fixation for a period of time (typically 500 ms for *monkey F* and *H*, and 1,500 ms for *monkey Eu* and *X*), they obtained fluid reward. These trials were designated as “canceled.” If monkeys failed to inhibit the movement, no reward was given, and the trial was termed “noncanceled.” If a saccade was initiated before the stop signal was scheduled to appear, the trial was classified as noncanceled based on the logic of the Logan race model (Logan and Cowan 1984). No time outs were imposed for a noncanceled error and reward volume was consistent across trial types.

The stop-signal delay (SSD) or time between target and stop-signal presentation determines the probability with which movements can be successfully countermanded (Logan and Cowan 1984). An initial set of SSDs was selected for each recording session based on the experimenter’s knowledge of the animal’s past performance. Although varied from session to session, these SSDs typically ranged from 43 up to 443 ms, in steps of 16 ms. SSD was then manipulated using either an adaptive staircasing algorithm that adjusted stopping difficulty based on accuracy, or by randomly selecting one of the defined SSDs. In the staircasing design, when subjects failed to inhibit responses, the SSD was decreased by a random step of 1, 2, or 3 stop-signals increasing the likelihood of success on the next stop trial. Similarly, when subjects were successful in inhibiting the eye movement, the next SSD was increased by a random step of 1, 2, or 3, decreasing the future probability of success. This procedure was used to ensure that subjects failed to inhibit action on ~50% of stop trials overall. Plots showing the probability of responding at each SSD (inhibition functions) were constructed and monitored online to ensure adequate performance.

The median session length for *monkey F* was 51 min [interquartile range (IQR): 27 to 77 min], and for *monkey H* was 236 min (119 to 366 min). In later data recorded from *monkey Eu* and *X*,

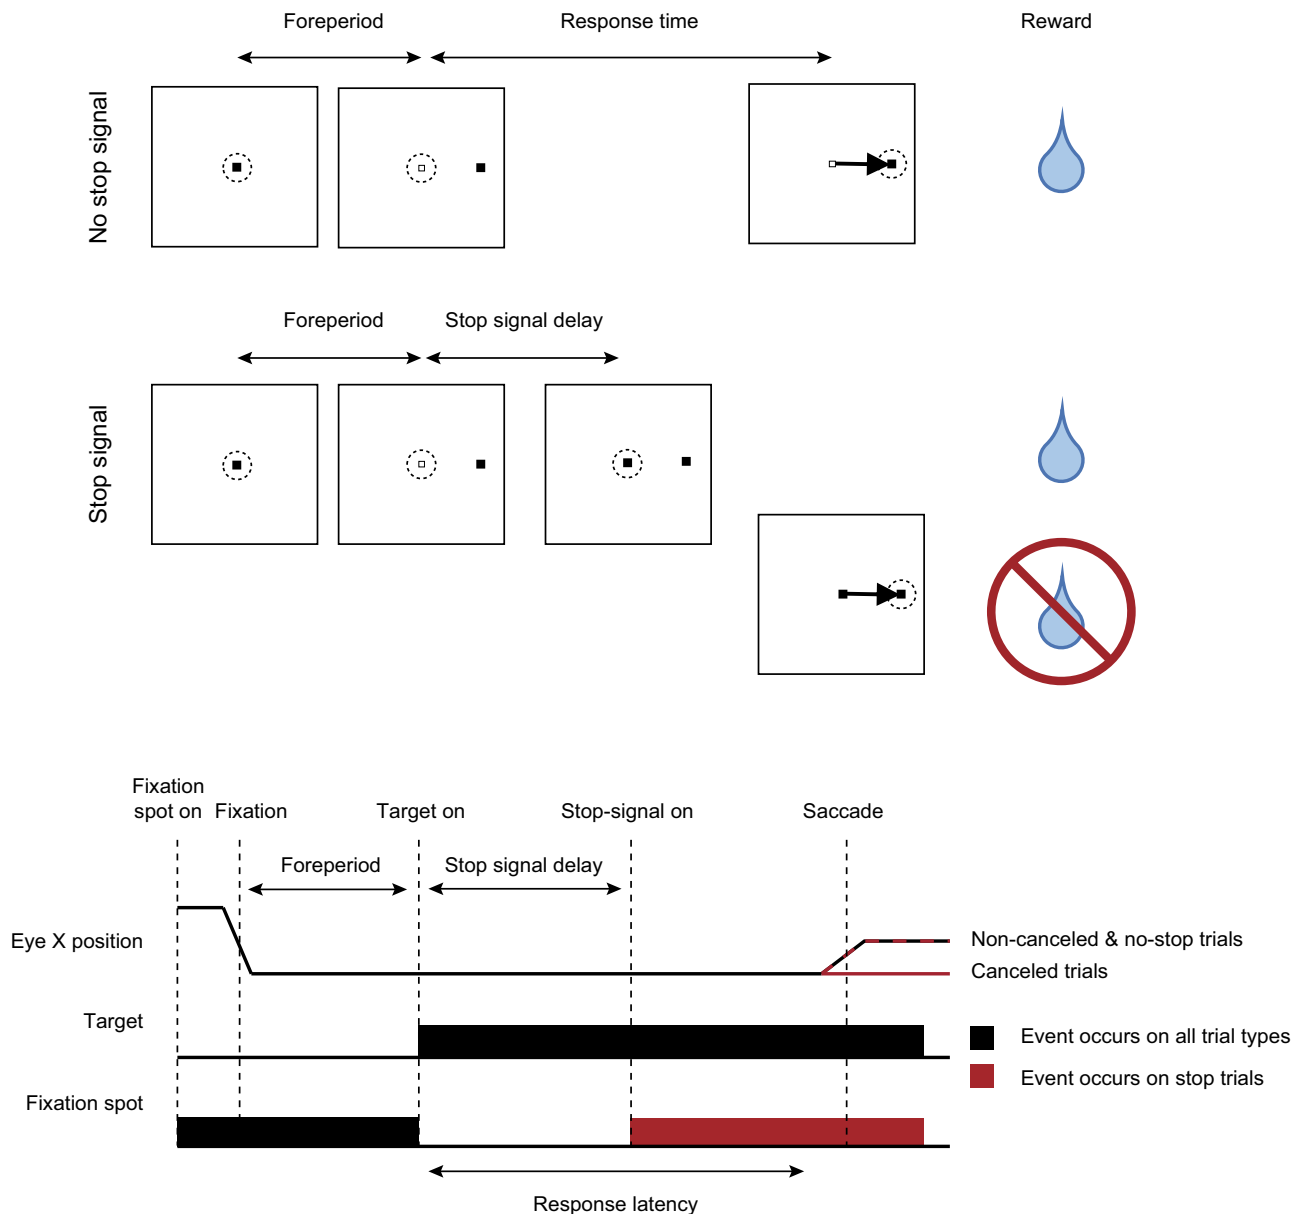


Fig. 1. Saccade countermanding task. Monkeys initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target was presented simultaneously at one of two possible locations. On no-stop-signal trials, monkeys were required to shift gaze to the target, whereupon after a variable period of time, fluid reward was provided. On stop-signal trials (~40% of trials), after the target appeared, the center of the fixation point was reilluminated after a variable stop-signal delay, which instructed the monkey to cancel the saccade. After holding fixation for a period of time, the monkey received feedback and reward. In a staircase procedure, stop-signal delay was adjusted such that monkeys successfully canceled the saccade in ~50% of trials. In the remaining trials, monkeys made noncanceled errors, in which no reward was delivered. In nonstaircased trials, stop-signal delays were randomly selected from a predetermined set of stop-signal delays.

trial length was fixed. For these monkeys, session lengths ranged from 87 to 184 min (modal session time: 169 min) for *monkey Eu*, and from 62 to 193 min (modal session time: 120 min) for *monkey X*. For *monkey F*, the intertrial interval was fixed (~1,000 ms), resulting in a variable trial length. However, for *monkey H*, the intertrial interval was varied. In all cases, trial length would be extended if a time-out (~500 ms for *monkey F* and *H*, ~5,000 ms for *monkey Eu* and *X*) was issued if the monkey aborted the trial (i.e., not maintaining fixation on a target, failing to make a saccade before a given deadline). For all sessions, time-outs were not issued for noncanceled trials; instead, the temporal progression of the trial would instead mirror that of a no-stop trial, but with reward omitted.

Eye Tracking

Eye position data were acquired, calibrated, and streamed to the Plexon computer using the EyeLink 1,000 infrared eye-tracking system (SR Research, Ontario, Canada). This system has an advertised resolution of 0.01°. Online, gaze was monitored using digital fixation windows. The size of these windows was determined by the experimenter based on the quality of the eye tracker calibration. Typically, subjects were allowed 1° of stimulus fixation error online. For final trial classification and analysis, saccade initiation and termination were detected offline using a custom algorithm implemented in the MATLAB programming environment (MathWorks, Natick, MA). Saccade starting and ending times were defined as periods when instantaneous velocity was elevated above 30°s⁻¹. The eye tracking

Table 1. Saccade latencies across trial types and data sets

Saccade type	Trial type	Aging Foreperiod		Nonaging Foreperiod	
		Monkey F (n = 84)	Monkey H (n = 157)	Monkey Eu (n = 12)	Monkey X (n = 17)
All	No-stop	230.7 ± 50.8	265.0 ± 50.6	311.9 ± 14.2	265.8 ± 18.9
	Noncanceled	238.1 ± 52.8	272.5 ± 66.5	282.3 ± 25.7	232.4 ± 12.9
Regular	No-stop	252.5 ± 45.7	297.8 ± 46.8	312.6 ± 14.4	266.3 ± 19.1
	Noncanceled	249.9 ± 46.5	301.8 ± 53.8	283.8 ± 25.6	233.3 ± 13.1
Express	No-stop	72.1 ± 6.7	84.0 ± 3.4	68.5 ± 2.5	85.5 ± 14.6
	Noncanceled	71.0 ± 5.5	84.1 ± 4.1	71.2 ± 8.4	85.3 ± 13.1

Values are average ± SD.

procedures reliably detected saccades > 0.2° in amplitude. Express saccades were classified as saccades to a target with a latency less than or equal to than 100 ms. No monkey produced enough anticipatory saccades to confound any of the analyses.

RESULTS

Monkeys Produce Express Saccades in a Countermanding Task

We retrospectively examined 425 sessions of saccade countermanding obtained from two monkeys (*monkey F*: 110 sessions; *monkey H*: 315 sessions). Collectively these monkeys completed 264,422 trials (*monkey F*: 77,438; *monkey H*: 186,984). While performing this task, monkeys *F* and *H* produced saccades with very short latencies which led to bimodal or multimodal latency distributions. Evident in both single behavioral sessions and across sessions, saccades with latencies ≤ 100 ms comprised a separate mode in the saccade latency distributions. These observations are congruent with previous descriptions of express saccades in macaques and were prominent in both monkeys.

Average saccade latencies across all data sets, split by trial type, are displayed in Table 1. On average express saccades comprised 6% of the saccade latencies in a given session, ranging between 0% and 66% of the saccade latencies within a given session. Across 110 sessions, *monkey F* generated saccades in 38,982 trials (32,992 no-stop, 5,990 noncanceled). The response distribution was clearly multimodal (Fig. 2A, top). Median saccade latency across these trials was 248 ms (IQR: 157–318 ms). Express saccades were elicited in 4,707 (12.0%) saccade trials and were significantly more common in one of the two target directions [78.73% to 21.27%, one-sample *t* test, $t(106) = 18.369$, $P < 0.001$]. The median express saccade latency was 87 ms (IQR: 82–92 ms). *Monkey H* generated saccades in 90,347 trials (73,750 no-stop, 16,597 noncanceled). Again, the response time distribution was multimodal (Fig. 2A, bottom) with a median saccade latency of 328 ms (IQR: 256–420 ms). Express saccades were elicited in 4,061 (4.5%) saccade trials and were significantly more common in one direction [44.05% to 55.95%, one-sample *t* test, $t(297) = -3.628$, $P < 0.001$]. The median express saccade latency was 92 ms (IQR: 91.93 ms – 96 ms). Findings for both monkeys are consistent when looking at the single session level (Fig. 2B).

Express Saccade Production and Race Model Violations

Performance of stop-signal countermanding tasks has been explained comprehensively by a race model in which trial

outcomes are dictated by the finishing time of stochastically independent GO and STOP processes (Logan and Cowan 1984). The race model is based on the assumption that the finishing time of noncanceled responses ($RT_{\text{noncanceled}}$) is consistently faster than the finishing time of responses when there is no stop signal ($RT_{\text{no stop}}$). Saccade latencies and express saccade proportions, split by trial types and monkeys, are presented in Table 2, for sessions with and without violations of the race model. In these archived data, we noted a number of sessions in which the response time (RT) relationship that justifies the application of the race model was violated. We were surprised to uncover a systematic pattern of variation of express saccade production across sessions when this relationship was or was not violated. We should note that previous publications about countermanding responses based on data from these monkeys utilized sessions in which the race model assumption was respected.

In these archived data, we observed an unexpected pattern of express saccade production related to satisfying the stochastic independence assumption of the race model. *Monkey F* violated the $RT_{\text{noncanceled}} < RT_{\text{no stop}}$ relationship in 50/84 sessions (59.5%). On no-stop trials, the proportion of express saccades when the relationship was respected was not different from that when the relationship was violated [$t(82) = 1.185$, $P = 0.239$, two-tailed]. However, during noncanceled trials, the proportion of express saccades when the relationship was respected was significantly greater than that when the relationship was violated [$t(82) = -2.495$, $P = 0.015$, two-tailed].

Monkey H showed a similar pattern of behavior and violated the $RT_{\text{noncanceled}} < RT_{\text{no stop}}$ relationship in 83/157 sessions (52.9%). During no-stop trials, the proportion of express saccades when the relationship was respected was not different from that when the relationship was violated [$t(155) = -1.393$, $P = 0.165$]. However, on noncanceled trials, the proportion of express saccades when the relationship was respected was significantly greater than that when the relationship was violated [$t(155) = -3.223$, $P = 0.002$]. Hence, when monkeys *F* and *H* performed the countermanding task consistent with the race model assumption of stochastic independence, they produced more express saccades than when they performed the task in violation of the assumption.

Monkey Eu had 3 sessions in which $RT_{\text{noncanceled}} > RT_{\text{no stop}}$, but the proportion of express saccades did not differ across sessions or trial types. *Monkey X* had no sessions in which $RT_{\text{noncanceled}} > RT_{\text{no stop}}$.

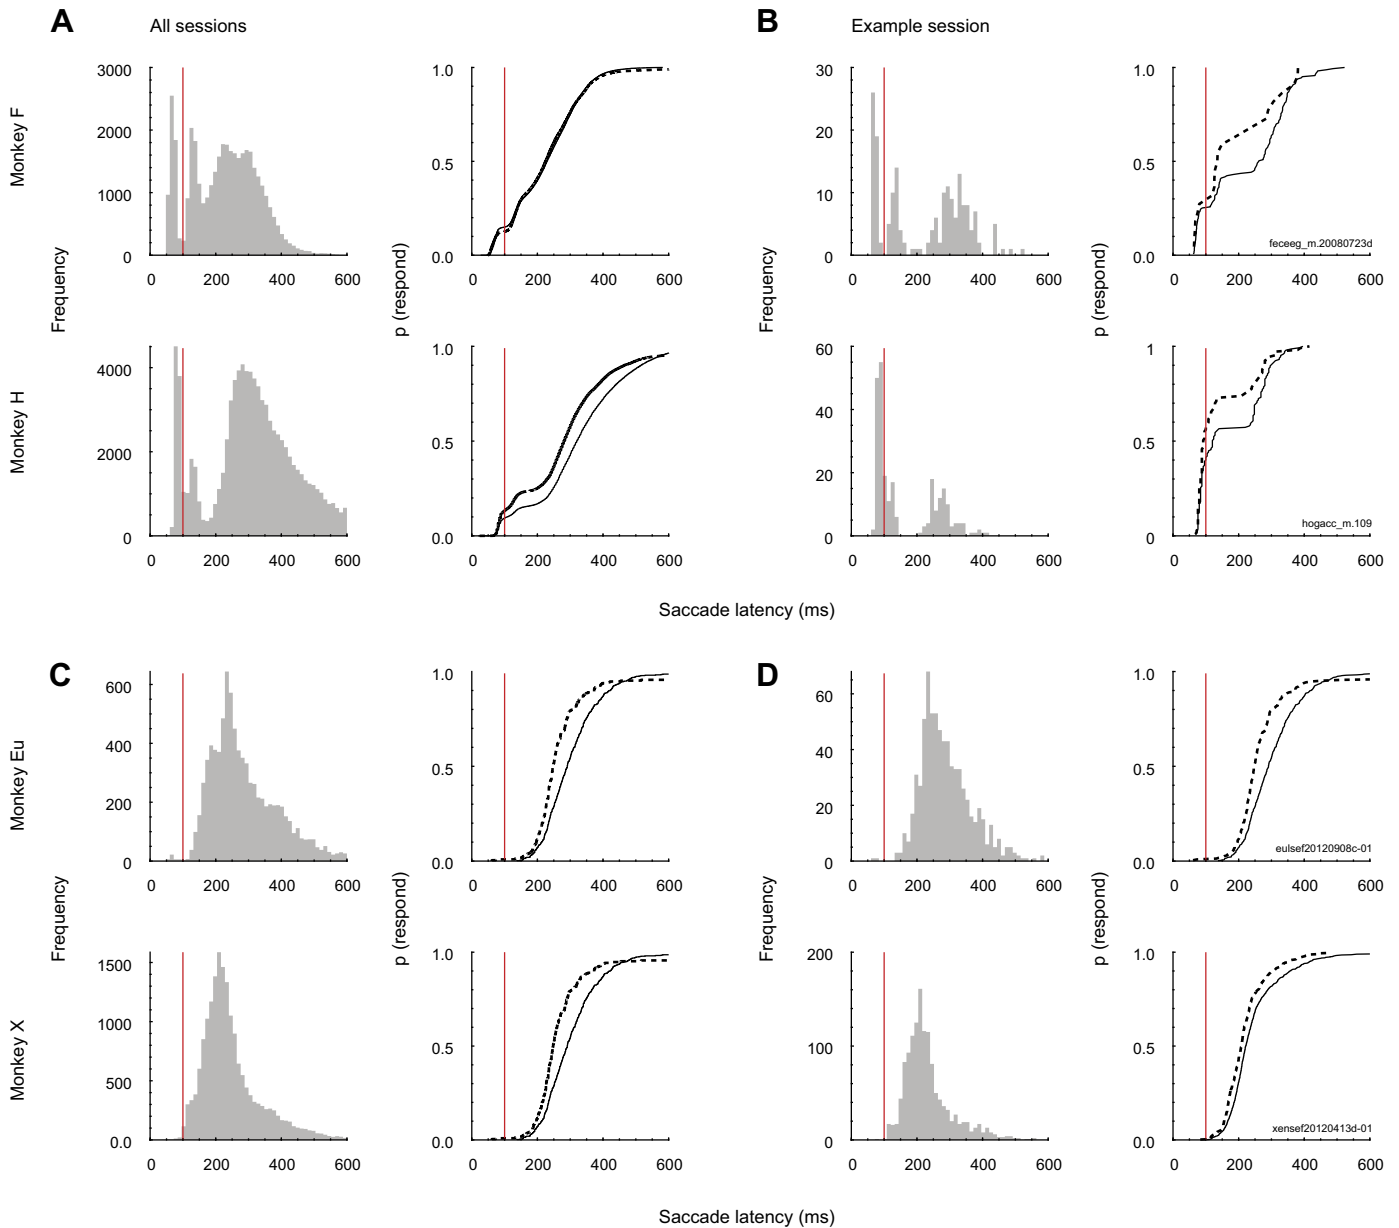


Fig. 2. Distributions of saccade latencies during saccade countermanding. Histograms (*left*) represent saccade latencies on noncanceled and no-stop trials combined, Cumulative distribution functions (*right*) of saccade latencies are presented for noncanceled (dashed lines) and no-stop trials (solid lines) separately. *A*: saccade latencies in example sessions for *monkeys F* and *H*. *B*: saccade latencies collapsed across all sessions for *monkeys F* and *H*. *C*: saccade latencies in example sessions for *monkeys Eu* and *X*. *D*: saccade latencies collapsed across all sessions for *monkeys Eu* and *X*. Red vertical lines represent the 100-ms saccade latency mark. Saccades with latencies before this value were classified as express saccades.

Simulated Express Saccade Production Influences Behavioral Outcomes

Given the apparent inconsistency between performance of the countermanding task and production of express saccades, we sought to understand why monkeys would initiate express saccades. Appreciating that monkeys are motivated to earn fluid reward, we explored whether express saccade production could be advantageous. To do so, we simulated countermanding performance with production of different fractions of express saccades. Unlike previous work (Boucher et al. 2007; Logan et al. 2015), we did not fit parameters based on observed measures of performance. Rather, we simply quantified the amount of reward earned when we simulated countermanding

performance while varying the fraction of express saccades produced as well as other parameters of the task. Using this approach, we found increasing express saccade production allowed more trials to be initiated under typical experimental parameters. Furthermore, a small increase in reward rate could be attained by including a small proportion of express saccades.

Saccade latencies were simulated using a linear ballistic accumulator (LBA) to instantiate the independent race underlying countermanding performance (Fig. 3A). The saccade GO process was simulated with accumulators having two distributions of rates. One accumulator had slower median accumulation rates, producing longer saccade latencies. The second had

Table 2. Saccade latencies and express saccades proportions across trial types and monkeys, split by race model violations

RT Relationship	Saccade Type	Aging				Nonaging			
		Monkey F (n=84)		Monkey H (n=157)		Monkey Eu (n=12)		Monkey X (n=17)	
		RT, ms	P (Express)	RT, ms	P (Express)	RT, ms	P (Express)	RT, ms	P (Express)
		<i>n</i> = 50 sessions		<i>n</i> = 83 sessions		<i>n</i> = 3 sessions		<i>n</i> = 0 sessions	
RT _{noncanceled} >	No-stop	225.8 ± 49.8	0.139 ± 0.119	261.4 ± 41.5	0.151 ± 0.084	303.7 ± 1.6			
RT _{no stop}	Noncanceled	249.8 ± 50.3	0.042 ± 0.075	301.2 ± 51.9	0.110 ± 0.114	316.6 ± 17.7	0.003 ± 0.005		
		<i>n</i> = 34 sessions		<i>n</i> = 74 sessions		<i>n</i> = 9 sessions		<i>n</i> = 17 sessions	
RT _{noncanceled} <	No-stop	237.8 ± 52.2	0.105 ± 0.138	267.0 ± 59.3	0.172 ± 0.100	314.7 ± 15.5	0.004 ± 0.004	265.8 ± 18.9	0.003 (0.003)
RT _{no stop}	Noncanceled	221.0 ± 52.4	0.112 ± 0.177	240.3 ± 66.7	0.195 ± 0.210	270.9 ± 15.6	0.009 ± 0.009	232.4 ± 12.9	0.006 (0.006)

Values are average ± SD. P(Express), probability of express saccade; RT, response time; RT_{noncanceled}, response times in noncanceled trials; RT_{no stop}, response times in no-stop trials.

faster median accumulator rates, producing express saccades. The fraction of trials governed by the faster accumulator was varied systematically. The STOP process was simulated with another accumulator. On simulated trials with a stop signal, if the STOP accumulator finished before the GO accumulator, then reward was earned for canceling the saccade, and if the STOP accumulator finished after the GO accumulator, then no reward was earned. On simulated trials with no stop-signal, reward was earned for saccades. Model parameters were not fit to performance measures. Instead, we simply explored qualitatively how reward rate varied as a function of various parameters. First, we varied the parameters of GO accumulators to generate different proportions of express saccades (0%, 10%, 50%, 90%; Fig. 3B). Second, we varied whether express saccades were rewarded or unrewarded. Third, we varied whether the trial length was fixed or varied. Finally, we varied the proportion of stop-signal trials. We examined how timing and experimental parameters similar to those used during data collection influenced average reward and trial rate.

Express saccade production increases the number of trials available. Under a fixed intertrial interval, trial length can vary as a function of response time. Under this paradigm, a faster saccade latency may lead to a shorter trial length and thus provide the monkey with more trials per minute to gain reward. With 40% stop trials, simulated data demonstrated that an increase in the proportion of express saccades led to an increase in trial rate (number of trials per minute) (Fig. 3C, squares). However, when trial lengths are fixed, trial rates only slightly decrease and subjects can only initiate around 11 trials per minute (Fig. 3C, circles).

Given that there is no penalty for making express saccades in no-stop trials, but there are time-out penalties in stop-trials,

an increase in stop trial proportion will increase the error rate and trial length and reduce overall trial rate. As such, we examined whether these values varied as a proportion of stop trials in a given session (Fig. 3E, top panels). Interestingly, if a session had a lower proportion of stop-trials (~10%), then producing more express saccades led to only a small increase in the number of trials available. However, if a session had a greater proportion of stop-trials (~70%), then producing more express saccades led to a greater increase in the number of trials available compared with lower stop-trial proportions. This is only true for variable trial lengths. If trial lengths were fixed, then express saccades decreased trial rates for all stop-signal proportions.

Express saccade production is not detrimental to reward rate. When trial length varied, express saccades were rewarded, and stop-signal delays adapted in a staircase procedure, we found systematic variation in reward rate with the proportion of express saccades in a session (Fig. 3D). At typical stop-trial proportions (~40%), producing a small proportion of express saccades, was not detrimental to performance and led to slightly increased reward rates (Fig. 3D). This relationship changed dependent on the proportion of stop-trials in a given session (Fig. 3E, bottom panels). If a session had a lower proportion of stop-trials (~10%), then producing more express saccades lead to a higher reward rate (Fig. 3E). However, if a session had a greater proportion of stop-trials (~60%), then producing more express saccades was detrimental to performance (Fig. 3E). This is only true for variable trial lengths. If trial lengths were fixed, then express saccades were detrimental to reward rates for all stop-signal proportions.

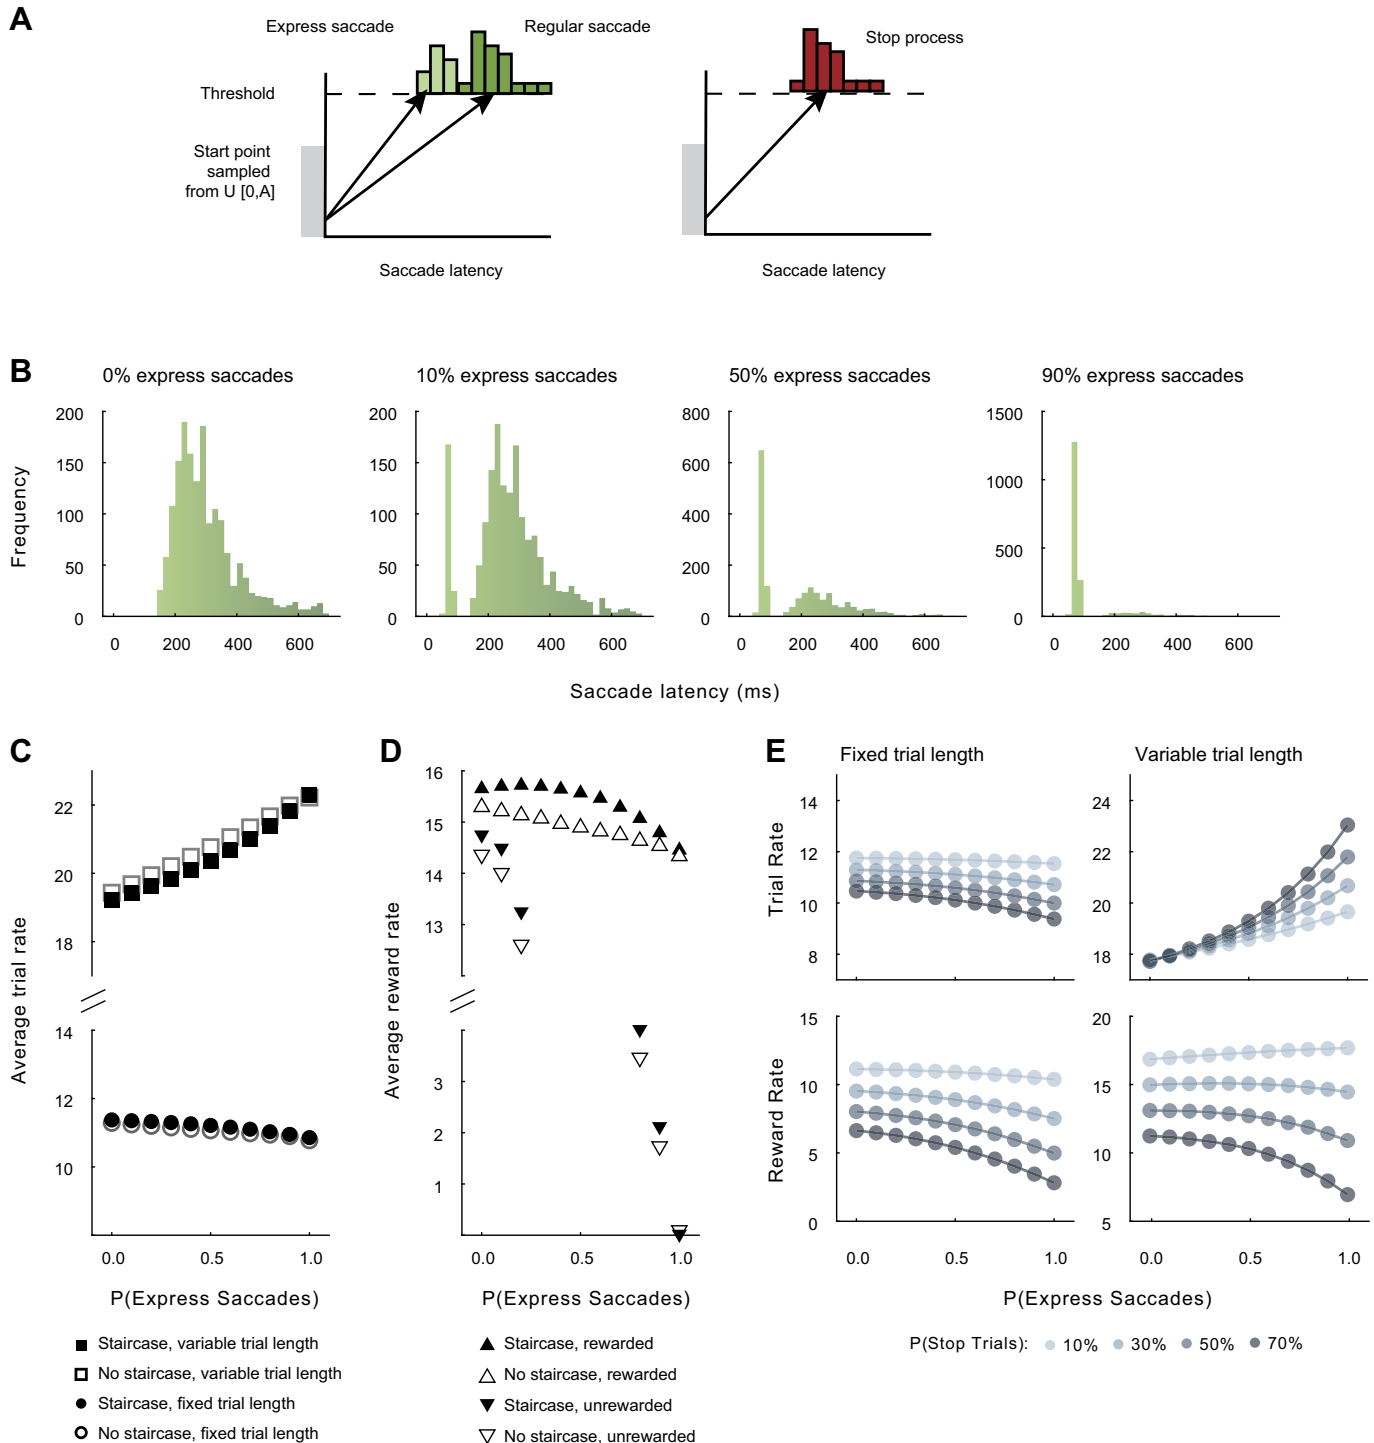
Fig. 3. Simulation of reward rate earned with variable fractions of express saccades. A: linear ballistic accumulators for GO process (left) and STOP process (right). Accumulation begins at a baseline level drawn from a uniform distribution (U) ranging from 0 to a maximum value of A. The GO process consisted of two mean rates, one producing express saccade latencies, and the other producing regular saccade latencies. Trial outcomes were specified by the finishing time of the fastest process according to the race model (Logan and Cowan 1984). On trials with no stop signal, reward was earned after the GO process produced either a regular or an express saccade. On trials with a stop signal, reward was earned only if the stop process linear ballistic accumulator (LBA) finished before both the regular and the express saccade LBAs. The following combinations of trial parameters were simulated: with or without staircase adjustment of stop-signal delay (SSD), with variable or fixed trial durations, with or without rewarding express saccade latencies. B: distributions of saccade latencies produced with indicated fractions of express saccades in a simulated session of 1,000 trials. C: average trial length over simulated sessions with (solid symbols) and without (open symbols) staircasing SSD and with variable (square) or fixed (circle) trial lengths. With variable but not fixed trial lengths, trial rate increased with the fraction of express saccades because more trials could be completed. Whether or not SSD was adjusted in a staircase made negligible difference. D: average reward rate over simulated sessions with (solid symbols) and without (open symbols) staircasing SSD and with express saccades rewarded (upright triangle) or not rewarded (inverted triangle). If express saccades were not rewarded, then reward rate decreased dramatically with increasing fraction of express saccades, whether or not SSD was adjusted in a staircase. If express saccades were rewarded, then reward rates were higher and decreased less with increasing fractions of express saccades. However, if SSD was adjusted in a staircase, then reward rate was maximal when ~10% express saccades were produced. E: variation of trial rate (top) and reward rate (bottom) as a function of proportion of express saccades [P(Express Saccades)] for fixed (left) and variable (right) trial durations. Progressively darker points plot values with progressively higher fraction of stop-signal trials. With lower fractions of stop-signal trials, the cost of express saccades on reward rate is reduced.

Experimental Express Saccade Production Varies with Reward Contingencies

The findings from the simulation demonstrated express saccade production can influence trial and reward rates under typical experimental conditions and parameters. Given this relationship, we then looked at these features within the data. In this experimental data, express saccade production increased from early to later training sessions and was mirrored by an increase in average reward rates. When express saccades were no longer rewarded, their prevalence significantly decreased.

Two other monkeys trained when express saccades were unrewarded produced very few express saccades across 10,000s of trials.

Express saccade production is learned through training. For monkey *F*, we tracked the progression of express saccade production from the initial training session onward. We observed the average proportion of express saccades significantly increased over time ($r = 0.49, P < 0.001$). In parallel, the average reward rate within each session increased significantly with experience ($r = 0.34, P < 0.001$). This pattern is also



clear when sessions were divided into five equal groups. From the earliest to latest training period, the proportion of express saccades increased significantly [$F(4,171) = 7.76$, $P < 0.001$, Fig. 4A], as did reward rate [$F(4,171) = 11.69$, $P < 0.001$, Fig. 4B].

Express saccade production is reduced when no longer rewarded. To causally test the association between reward rate and express saccade production, we recorded 31 additional behavioral sessions with *monkey F*. These sessions were recorded after the 425 sessions previously described and were not included in the previous analyses. In these sessions, we monitored saccade latencies online. We first recorded a block of sessions where correct express saccades were rewarded ($n = 13$ sessions). This was followed by a block of sessions in which express saccades were unrewarded ($n = 18$ sessions).

Sessions with both unrewarded and rewarded express saccades had multimodal saccade latency distributions (Fig. 4C).

When rewarded, express saccade trials comprised 22.6% of the latency distribution. However, when unrewarded, express saccades became much less common and a new mode appeared in the saccade latency distribution in the 200–250 ms range. In these sessions, the proportion of express saccades dropped considerably and comprised only 7.9% of the latency distribution, $\chi^2(1, N = 10,465) = 405.648$, $P < 0.001$ (Fig. 4D). Reverse to observations in the early training stages (when express saccades were rewarded), we found that the average proportion of express saccades significantly decreased with time from the reward manipulation ($r = -0.723$, $P < 0.001$, Fig. 4E). Interestingly, average reward rate exhibited no trends ($r = 0.183$, $P = 0.467$, Fig. 4F).

Monkeys trained with no reward contingencies don't produce express saccades. Further evidence that monkeys learn to produce express saccades through training was obtained by examining the saccade latency of two other monkeys who were

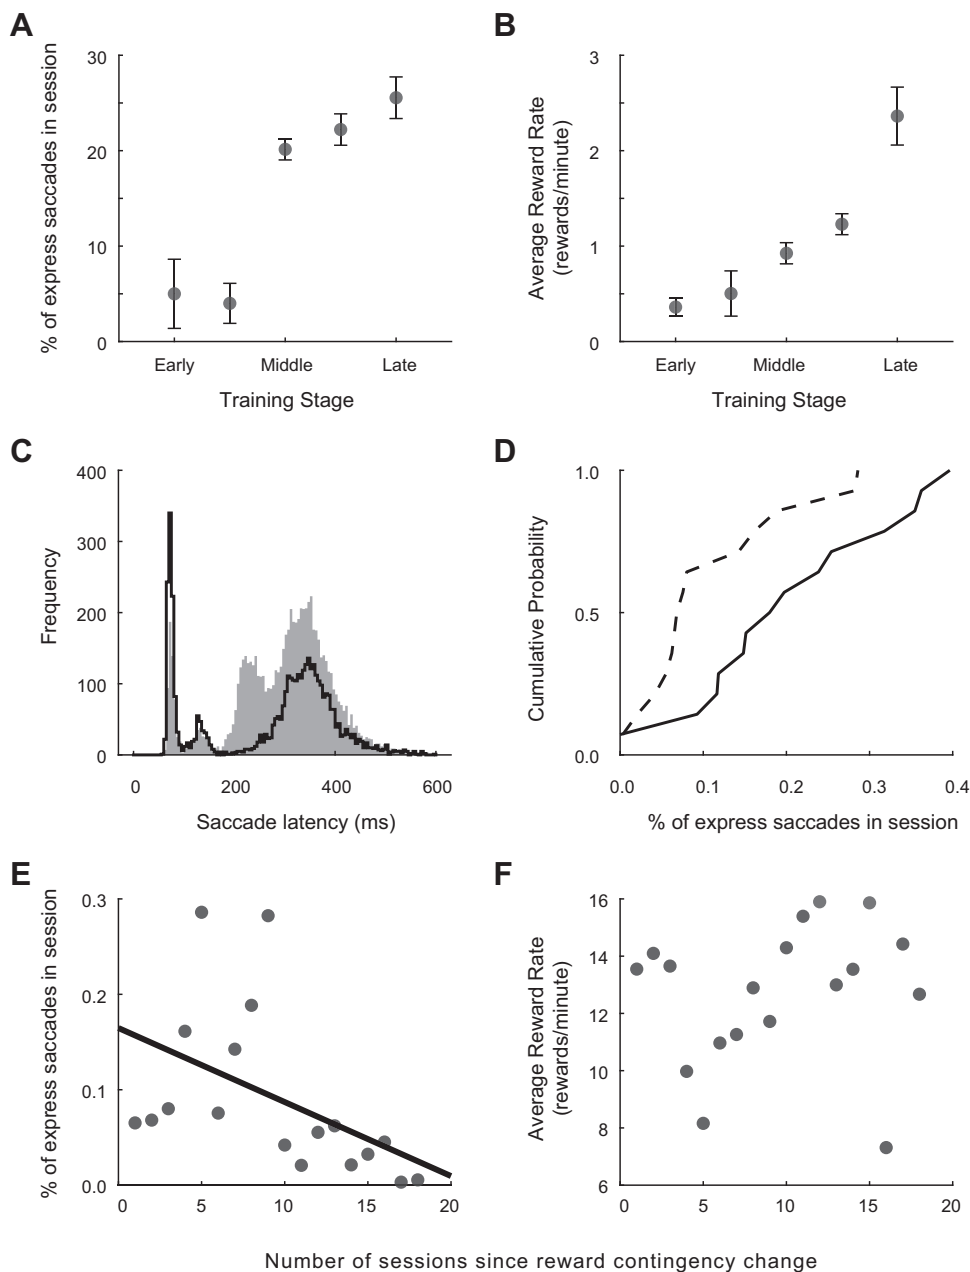


Fig. 4. Learning and unlearning express saccades. Variation of performance of *monkey F* as a function of training and task contingencies. *A*: average reward rate earned from early to late stages of training in 18 sessions when express saccades were rewarded just like regular saccades on no-stop trials. *B*: percentage of express saccades produced across corresponding stages of training. The monkey exploited contingencies of the task to earn more reward by producing a fraction of express saccades. *C*: distributions of saccade latencies when express saccades were rewarded (open symbols) and in subsequent sessions when only no-stop trials with regular saccade latencies were rewarded (solid symbols). *D*: cumulative probability of express saccade production in sessions when they were rewarded (solid line) and in later sessions when they were not (dashed line). *E*: the percentage of express saccades decreased across sessions following changing contingency to not reward express saccades. *F*: average reward rate did not decrease across corresponding sessions following the change of reward contingency.

trained with unrewarded express saccades. We examined an additional 29 sessions of saccade countermanding in two other monkeys (*monkey Eu*: 12 sessions; *monkey X*: 17 sessions) to compare training histories and reward contingencies. Collectively, these monkeys completed 33,816 trials (*monkey Eu*: 11,583; *monkey X*: 22,233). The saccade latency distributions were unimodal for both monkeys (Fig. 2, *C* and *D*, for all sessions and an example session respectively). Express saccades were elicited in only 107 trials, comprising only 0.40% of all saccade latencies. Across the 29 sessions, *monkey Eu* produced 42 (0.49%) express saccades and *monkey X* produced 65 (0.36%) express saccades.

Temporal Predictability Can Affect Express Saccade Production

Express saccades occur more often when the timing of a target presentation is predictable (Paré and Munoz 1996; Saslow 1967; Schiller et al. 2004). As such, we examined saccade latency as a function of the foreperiod between fixation at the central cue and target presentation. We started by looking at the distributions and respective survivor functions for foreperiods in our study. The survivor function is the probability that a target has not yet appeared by a given time. When foreperiods are sampled from a uniform distribution, the survivor function linearly decreases. Hence, as time passes, the proportion of the distribution from which the target onset time can be selected decreases. Although this function offers some insight to the temporal evolution of the distribution, temporal predictability is quantified by the hazard function, which is the conditional probability of an event occurring at a given time given that it has not yet occurred (Luce 1986; Nobre et al. 2007). Formally, the hazard rate is the ratio of the probability density of the event divided by the survivor function. In uniform distributions, the hazard rate for target presentation increases over time, resulting in an aging function. This aging function stipulates target appearance becomes predictable as the foreperiod progresses. Conversely, when sampled from exponentially decaying distributions, the hazard rate

is invariant over time. Under these conditions, the time of target onset is unpredictable. Given the retrospective nature of this study, we first examined the various patterns of foreperiod distributions experienced by monkeys within our data set. Although the pattern of foreperiods across some sessions for *monkey F* and *H* were variable, we identified a subset of sessions for each monkey with a reliable, uniform foreperiod distribution. Results for regular and express saccades separately are highlighted in Table 3.

In the majority of sessions for *monkey F* (84/110, 76%) a uniform distribution of foreperiods ranged from ~420 to ~630 ms (Fig. 5A). A linear decrease in the survivor function is associated with a continually increasing hazard rate of times when the foreperiod elapsed. In some of these sessions, saccade latency decreased with foreperiod duration (33/84 sessions, 39%). However, despite the decrease of saccade latencies at longer foreperiods, the proportion of express saccades did not vary over foreperiod duration [$F(2, 249) = 2.11, P = 0.123$].

We observed the opposite pattern of results from *monkey H*, who experienced a uniform distribution of foreperiods in half of the sessions (157/315 sessions, 50%) ranging over a much wider interval, ~760 to ~2,275 ms (Fig. 5B). No relation between saccade latency and foreperiod duration was found in most of these sessions (104/157, 66%), but in the remainder *monkey H* demonstrated increasing latencies with longer foreperiods (53/157, 34%). This slowing of saccade latencies was accompanied by reduced express saccades production [$F(15, 2,438) = 14.44, P < 10^{-5}$].

Monkey Eu and *monkey X* experienced a common pattern of foreperiod distribution ranging from ~600 to ~2,100 ms ($n = 29/29$ sessions). Compared with the foreperiod experienced by monkeys *F* and *H*, the survivor function of this distribution was exponentially decaying, resulting in a nonaging foreperiod (Fig. 5, *C* and *D*). As expected with nonaging foreperiods, none of *monkey Eu*'s sessions had significant changes in saccade latency or in the proportion of express saccades generated as a function of foreperiod. However, and somewhat unexpectedly,

Table 3. Number of significant correlations between saccade latency and foreperiod duration for each foreperiod distribution, each saccade type (all, regular, or express), and each trial type (all, no-stop, and noncanceled)

Saccade type	Trial type	Aging Foreperiod		Nonaging Foreperiod	
		<i>Monkey F</i>	<i>Monkey H</i>	<i>Monkey Eu</i>	<i>Monkey X</i>
All saccades	All	34/84 (33 negative)	54/157 (1 negative)	0/12 (0 negative)	9/17 (9 negative)
	No-stop	35/84 (34 negative)	51/157 (1 negative)	1/12 (0 negative)	7/17 (7 negative)
	Noncanceled	9/84 (9 negative)	31/157 (4 negative)	2/12 (1 negative)	5/17 (5 negative)
Regular saccades	All	27/84 (26 negative)	11/157 (6 negative)	0/12 (0 negative)	10/17 (10 negative)
	No-stop	30/84 (28 negative)	15/157 (9 negative)	1/12 (0 negative)	6/17 (6 negative)
	Noncanceled	11/84 (10 negative)	12/157 (4 negative)	2/12 (1 negative)	5/17 (5 negative)
Express saccades	All	4/84 (3 negative)	7/157 (5 negative)	0/12 (0 negative)	0/17 (0 negative)
	No-stop	4/84 (3 negative)	5/157 (3 negative)	0/12 (0 negative)	0/17 (0 negative)
	Noncanceled	1/84 (0 negative)	4/157 (4 negative)	1/12 (1 negative)	1/17 (1 negative)

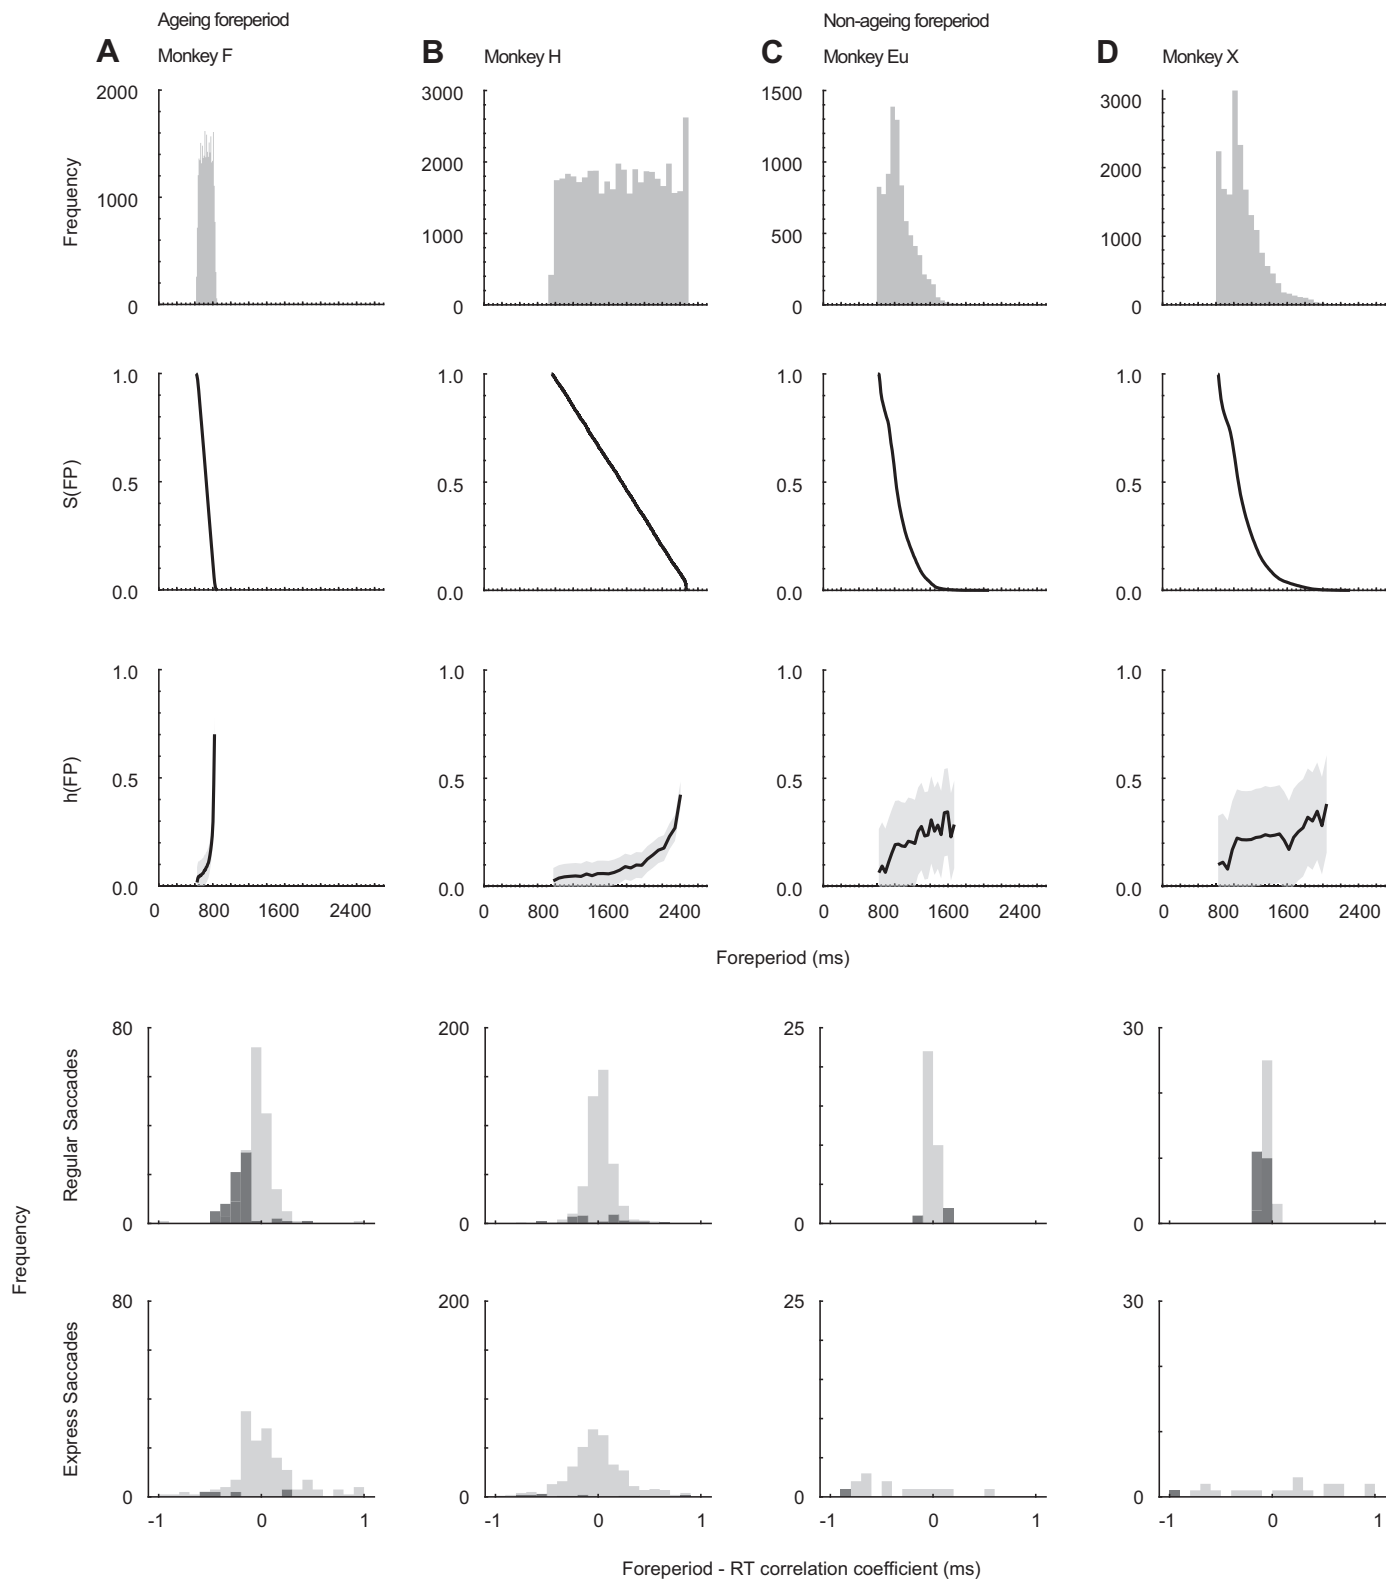


Fig. 5. Distributions of foreperiods and relation to saccade latency. Histograms (*1st row*), survivor distribution [$S(\text{FP})$; *2nd row*], and hazard rate [$h(\text{FP})$; *3rd row*] of foreperiods used for *monkey F* (*A*), *monkey H* (*B*), *monkey Eu* (*C*), and *monkey X* (*D*). *Monkeys F* and *H* experienced aging foreperiods, and *monkeys Eu* and *X* experienced approximately nonaging foreperiods. Correlation coefficients for the correlation between saccade latency and foreperiod are presented in the *bottom row*. Dark shades represent sessions in which the correlation coefficients were significant; light shades are nonsignificant sessions. Significant foreperiod effects were limited to regular saccades for *monkey F* and *Eu*. Foreperiod duration did not influence express saccade latency. RT, response time.

a foreperiod effect was observed in over half of the sessions for monkey X (9/17, 53%). Furthermore, although uncommonly produced, the proportion of express saccades did decrease as a function of foreperiod bin, $F(15, 206) = 2.26$, $P = 0.006$.

DISCUSSION

We observed multimodal distributions of saccade latencies in a saccade countermanding task. Furthermore, simulations revealed that trial and reward rate in this task can be increased by varying the proportion of express saccades made within a session, under certain conditions. We found that monkeys learned and exploited this association over time. We then manipulated the reward contingencies for express saccade production and found when monkeys made significantly lower proportion of express saccades when no longer rewarded for them. It is important to note the incidental nature of these findings, and that the features studied were not the main focus of the investigations. As such, our interpretations of these data are limited and future experimental designs should be employed to test them more thoroughly. However, this study has demonstrated for the first time how express saccade production can be accomplished through strategic modifications in saccade latencies, under the guidance of cognitive control.

Express Saccades Production in Atypical Conditions

Consistent with classic reports of express saccades in monkeys, the earliest modes in the resulting distributions peak below 100 ms and display very little variance (Boch and Fischer 1986; Boch et al. 1984; Fischer and Boch 1983; Schiller et al. 2004). Although express saccades have also been observed in “overlap” conditions, where the fixation spot remains on even while the target is illuminated (Amatya et al. 2011; Boch and Fischer 1986; Knox and Wolohan 2015), this finding in the saccade countermanding task was unexpected given the range of other features that are counterproductive for express saccade production. As such, although intended to be random and unpredictable, the production of express saccades under these conditions may suggest that monkeys have some knowledge about the underlying contingencies in our experiment and generated predictions to adapt saccade latencies accordingly.

Temporal Predictability and Saccade Latency

Monitoring the timing of events allows for temporal predictions of target onset to be generated, allowing for a movement to be planned and prepared ahead of time (Kolling and O'Reilly 2018; Petter et al. 2018). The use of hazard functions is motivated by the premise that participants have a representation of the elapsed time in the current situation based on knowledge of the experienced distribution of durations. If an interval elapsing before an event is randomly sampled from a uniform distribution (also known as an aging distribution), then participants will behave as if the event becomes more likely as time progresses. Thus, temporal prediction is critical in reaction time tasks in which foreperiods are a major determinant of response latency. Using distributions of foreperiods selected uniformly from a range of values, i.e., aging, manual response time studies have shown as foreperiod increases, response times become faster (Ameqrane et al. 2014; Correa and Nobre 2008; Drazin 1961; Näätänen 1970; Niemi and Näätänen 1981;

Thomaschke et al. 2011). Consistent with the aforementioned manual response studies, this finding has been replicated focusing on saccade latencies in humans (Findlay 1981) and monkeys (Schall 1988). We replicated this effect in the current study. Monkeys can make effective but idiosyncratic temporal predictions about the time a target may appear on screen.

Reward Maximization

The results demonstrate that monkeys adjusted saccade production to maximize earned rewards. By producing express saccades, monkeys could increase the number of trials available and thus the opportunity to gain a reward. Monkeys did not produce enough express saccades to be detrimental to reward rate.

If the total time of the task is fixed, a rational subject will trade off speed and accuracy to try and maximize the amount of reward they can gain in the given session. If more time is spent on one trial, there will be a greater chance of a correct response; however, longer trials reduce the total number of trials available and thus the total amount of reward available. This ultimately gives rise to a speed-accuracy trade-off which aims to maximize gains through flexible cognitive control of behavior.

Nonhuman primates are adept at optimizing reward rate by exploring different behavioral strategies and task parameters (Feng et al. 2009). Feng et al. developed a model allowing for the calculation of choice bias that yields the optimal harvesting of reward, given the animals' sensitivities to a visual stimulus. They found monkeys acquire over 98% of the possible maximum rewards, with shifts away from optimality erring in the direction of smaller penalties. In addition to this, monkeys have also shown to discover unexpected ways to exploit task contingencies. Lowe and Schall (2019) demonstrated this effect in a pro-/antisaccade visual search task where a vertical singleton cued that a prosaccade should be made toward it, and a horizontal singleton cued an antisaccade away from it. Close examination of the reward contingencies of the experiment found that shifting gaze toward a vertical stimulus was the correct outcome on 66% of trials. Exploiting this contingency, they found that one monkey produced more frequent and faster responses to vertical items than to any other item in the array, suggesting they adopted a strategy of searching for vertical items opposed to using the stimulus response rule provided by the singleton. Bichot et al. (1996) found a similar behavior in a color pop-out task: monkeys trained exclusively to find one color in an array persistently direct gaze to stimuli of that color, regardless of whether it is a distractor or target. These are just some of the examples of macaque monkeys creatively exploiting reward contingencies.

Conclusion

We report incidental findings revealing that express saccade production can be adjusted to increase the opportunities to gain reward. This observation suggests that the mechanisms of express saccade generation are sensitive to trial history and integrate information over multiple trials. This suggests that higher cortical areas that are involved in performance monitoring may contribute to the executive control of express saccade production (Dash et al. 2020; Donahue et al. 2013; Sajad et al. 2019; So and Stuphorn 2010; Stuphorn and Schall

2006; Stuphorn et al. 2000). Further research is needed to understand whether express saccade production can be controlled independently or whether express saccade production is accomplished as a corollary to strategic adjustments in overall RT.

ACKNOWLEDGMENTS

The authors thank K. A. Lowe and J. A. Westerberg, and Drs. T. Reppert, A. Sajad, and C. R. Subraveli for helpful discussions and comments on the manuscript.

Preprint is available at <https://doi.org/10.1101/2020.03.20.000760>.

GRANTS

This work was supported by R01-MH55806, R01-EY019882, and P30-EY08126 and by Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience (JS).

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.


AUTHOR CONTRIBUTIONS

S.P.E. analyzed data; S.P.E. and J.D.S. interpreted results of experiments; S.P.E. and J.D.S. prepared figures; S.P.E. and J.D.S. drafted manuscript; S.P.E. and J.D.S. edited and revised manuscript; S.P.E. and J.D.S. approved final version of manuscript; J.D.S. conceived and designed research; J.D.S. performed experiments.

REFERENCES

- Amatya N, Gong Q, Knox PC. Differing proportions of 'express saccade makers' in different human populations. *Exp Brain Res* 210: 117–129, 2011. doi:10.1007/s00221-011-2609-z.
- Ameqrane I, Pouget P, Wattiez N, Carpenter R, Missal M. Implicit and explicit timing in oculomotor control. *PLoS One* 9: e93958, 2014. doi:10.1371/journal.pone.0093958.
- Bibi R, Edelman JA. The influence of motor training on human express saccade production. *J Neurophysiol* 102: 3101–3110, 2009. doi:10.1152/jn.90710.2008.
- Bichot NP, Schall JD, Thompson KG. Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature* 381: 697–699, 1996. doi:10.1038/381697a0.
- Boch R, Fischer B. Further observations on the occurrence of express-saccades in the monkey. *Exp Brain Res* 63: 487–494, 1986. doi:10.1007/BF00237472.
- Boch R, Fischer B, Ramsperger E. Express-saccades of the monkey: reaction times versus intensity, size, duration, and eccentricity of their targets. *Exp Brain Res* 55: 223–231, 1984. doi:10.1007/BF00237273.
- Boucher L, Palmeri TJ, Logan GD, Schall JD. Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychol Rev* 114: 376–397, 2007. doi:10.1037/0033-295X.114.2.376.
- Cabel DW, Armstrong IT, Reingold E, Munoz DP. Control of saccade initiation in a countermanding task using visual and auditory stop signals. *Exp Brain Res* 133: 431–441, 2000. doi:10.1007/s002210000440.
- Carpenter RH. *Movements of the Eyes* (2nd ed.). London, UK: Pion Limited, 1988.
- Colonius H, Ozyurt J, Arndt PA. Countermanding saccades with auditory stop signals: testing the race model. *Vision Res* 41: 1951–1968, 2001. doi:10.1016/S0042-6989(01)00084-0.
- Correa A, Nobre AC. Neural modulation by regularity and passage of time. *J Neurophysiol* 100: 1649–1655, 2008. doi:10.1152/jn.90656.2008.
- Dash S, Peel TR, Lomber SG, Corneil BD. Impairment but not abolishment of express saccades after unilateral or bilateral inactivation of the frontal eye fields. *J Neurophysiol* 123: 1907–1919, 2020. doi:10.1152/jn.00191.2019.
- Donahue CH, Seo H, Lee D. Cortical signals for rewarded actions and strategic exploration. *Neuron* 80: 223–234, 2013. doi:10.1016/j.neuron.2013.07.040.
- Draizin DH. Effects of foreperiod, foreperiod variability, and probability of stimulus occurrence on simple reaction time. *J Exp Psychol* 62: 43–50, 1961. doi:10.1037/h0046860.
- Emeric EE, Brown JW, Boucher L, Carpenter RH, Hanes DP, Harris R, Logan GD, Mashru RN, Paré M, Pouget P, Stuphorn V, Taylor TL, Schall JD. Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vision Res* 47: 35–49, 2007. doi:10.1016/j.visres.2006.08.032.
- Feng S, Holmes P, Rorie A, Newsome WT. Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLOS Comput Biol* 5: e1000284, 2009. doi:10.1371/journal.pcbi.1000284.
- Findlay JM. Spatial and temporal factors in the predictive generation of saccadic eye movements. *Vision Res* 21: 347–354, 1981. doi:10.1016/0042-6989(81)90162-0.
- Fischer B, Boch R. Saccadic eye movements after extremely short reaction times in the monkey. *Brain Res* 260: 21–26, 1983. doi:10.1016/0006-8993(83)90760-6.
- Fischer B, Boch R, Ramsperger E. Express-saccades of the monkey: effect of daily training on probability of occurrence and reaction time. *Exp Brain Res* 55: 232–242, 1984. doi:10.1007/BF00237274.
- Godlove DC, Garr AK, Woodman GF, Schall JD. Measurement of the extraocular spike potential during saccade countermanding. *J Neurophysiol* 106: 104–114, 2011. doi:10.1152/jn.00896.2010.
- Godlove DC, Schall JD. Microsaccade production during saccade cancellation in a stop-signal task. *Vision Res* 118: 5–16, 2016. doi:10.1016/j.visres.2014.10.025.
- Hanes DP, Carpenter RH. Countermanding saccades in humans. *Vision Res* 39: 2777–2791, 1999. doi:10.1016/S0042-6989(99)00011-5.
- Hanes DP, Patterson WF II, Schall JD. Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J Neurophysiol* 79: 817–834, 1998. doi:10.1152/jn.1998.79.2.817.
- Hanes DP, Schall JD. Countermanding saccades in macaque. *Vis Neurosci* 12: 929–937, 1995. doi:10.1017/S095252380009482.
- Jóhannesson OI, Edelman JA, Sigurþórsson BD, Kristjánsson Á. Effects of saccade training on express saccade proportions, saccade latencies, and peak velocities: an investigation of nasal/temporal differences. *Exp Brain Res* 236: 1251–1262, 2018. doi:10.1007/s00221-018-5213-7.
- Knox PC, Wolohan FD. Temporal stability and the effects of training on saccade latency in "express saccade makers". *PLoS One* 10: e0120437, 2015. doi:10.1371/journal.pone.0120437.
- Kolling N, O'Reilly JX. State-change decisions and dorsomedial prefrontal cortex: the importance of time. *Curr Opin Behav Sci* 22: 152–160, 2018. doi:10.1016/j.cobeha.2018.06.017.
- Kornylo K, Dill N, Saenz M, Krauzlis RJ. Cancelling of pursuit and saccadic eye movements in humans and monkeys. *J Neurophysiol* 89: 2984–2999, 2003. doi:10.1152/jn.00859.2002.
- Logan GD, Cowan WB. On the ability to inhibit thought and action: a theory of an act of control. *Psychol Rev* 91: 295–327, 1984. doi:10.1037/0033-295X.91.3.295.
- Logan GD, Yamaguchi M, Schall JD, Palmeri TJ. Inhibitory control in mind and brain 2.0: blocked-input models of saccadic countermanding. *Psychol Rev* 122: 115–147, 2015. doi:10.1037/a0038893.
- Lowe KA, Schall JD. Sequential operations revealed by serendipitous feature selectivity in frontal eye field (Preprint). *bioRxiv* 683144, 2019.
- Luce RD. *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford, UK: Oxford University Press, 1986.
- McPeck RM, Schiller PH. The effects of visual scene composition on the latency of saccadic eye movements of the rhesus monkey. *Vision Res* 34: 2293–2305, 1994. doi:10.1016/0042-6989(94)90108-2.
- Milstein DM, Dorris MC. The influence of expected value on saccadic preparation. *J Neurosci* 27: 4810–4818, 2007. doi:10.1523/JNEUROSCI.0577-07.2007.
- Morein-Zamir S, Kingstone A. Fixation offset and stop signal intensity effects on saccadic countermanding: a crossmodal investigation. *Exp Brain Res* 175: 453–462, 2006. doi:10.1007/s00221-006-0564-x.
- Näätänen R. The diminishing time-uncertainty with the lapse of time after the warning signal in reaction-time experiments with varying fore-periods. *Acta Psychol (Amst)* 34: 399–419, 1970. doi:10.1016/0001-6918(70)90035-1.
- Niemi P, Näätänen R. Foreperiod and simple reaction time. *Psychol Bull* 89: 133–162, 1981. doi:10.1037/0033-2909.89.1.133.
- Nobre A, Correa A, Coull J. The hazards of time. *Curr Opin Neurobiol* 17: 465–470, 2007. doi:10.1016/j.conb.2007.07.006.
- Paré M, Munoz DP. Saccadic reaction time in the monkey: advanced preparation of oculomotor programs is primarily responsible for express

- saccade occurrence. *J Neurophysiol* 76: 3666–3681, 1996. doi:[10.1152/jn.1996.76.6.3666](https://doi.org/10.1152/jn.1996.76.6.3666).
- Petter EA, Gershman SJ, Meck WH.** Integrating models of interval timing and reinforcement learning. *Trends Cogn Sci* 22: 911–922, 2018. doi:[10.1016/j.tics.2018.08.004](https://doi.org/10.1016/j.tics.2018.08.004).
- Sajad A, Godlove DC, Schall JD.** Cortical microcircuitry of performance monitoring. *Nat Neurosci* 22: 265–274, 2019. doi:[10.1038/s41593-018-0309-8](https://doi.org/10.1038/s41593-018-0309-8).
- Saslow MG.** Latency for saccadic eye movement. *J Opt Soc Am* 57: 1030–1033, 1967. doi:[10.1364/JOSA.57.001030](https://doi.org/10.1364/JOSA.57.001030).
- Schall J.** Saccade latency and preparatory neuronal activity in the supplementary and frontal eye fields. *Soc Neurosci Abstr* 159, 1988.
- Schiller PH, Haushofer J, Kendall G.** An examination of the variables that affect express saccade generation. *Vis Neurosci* 21: 119–127, 2004. doi:[10.1017/S0952523804042038](https://doi.org/10.1017/S0952523804042038).
- So NY, Stuphorn V.** Supplementary eye field encodes option and action value for saccades with variable reward. *J Neurophysiol* 104: 2634–2653, 2010. doi:[10.1152/jn.00430.2010](https://doi.org/10.1152/jn.00430.2010).
- Stuphorn V, Schall JD.** Executive control of countermanding saccades by the supplementary eye field. *Nat Neurosci* 9: 925–931, 2006. doi:[10.1038/nn1714](https://doi.org/10.1038/nn1714).
- Stuphorn V, Taylor TL, Schall JD.** Performance monitoring by the supplementary eye field. *Nature* 408: 857–860, 2000. doi:[10.1038/35048576](https://doi.org/10.1038/35048576).
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O.** Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142: 284–291, 2002. doi:[10.1007/s00221-001-0928-1](https://doi.org/10.1007/s00221-001-0928-1).
- Thakkar KN, Schall JD, Boucher L, Logan GD, Park S.** Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biol Psychiatry* 69: 55–62, 2011. doi:[10.1016/j.biopsych.2010.08.016](https://doi.org/10.1016/j.biopsych.2010.08.016).
- Thakkar KN, Schall JD, Logan GD, Park S.** Cognitive control of gaze in bipolar disorder and schizophrenia. *Psychiatry Res* 225: 254–262, 2015. doi:[10.1016/j.psychres.2014.12.033](https://doi.org/10.1016/j.psychres.2014.12.033).
- Thomaschke R, Kiesel A, Hoffmann J.** Response specific temporal expectancy: evidence from a variable foreperiod paradigm. *Atten Percept Psychophys* 73: 2309–2322, 2011. doi:[10.3758/s13414-011-0179-6](https://doi.org/10.3758/s13414-011-0179-6).
- Verbruggen F, Logan GD.** Proactive adjustments of response strategies in the stop-signal paradigm. *J Exp Psychol Hum Percept Perform* 35: 835–854, 2009. doi:[10.1037/a0012726](https://doi.org/10.1037/a0012726).
- Vullings C, Madelain L.** Control of saccadic latency in a dynamic environment: allocation of saccades in time follows the matching law. *J Neurophysiol* 119: 413–421, 2018. doi:[10.1152/jn.00634.2017](https://doi.org/10.1152/jn.00634.2017).
- Vullings C, Madelain L.** Discriminative control of saccade latencies. *J Vis* 19: 16, 2019.
- Walton MM, Gandhi NJ.** Behavioral evaluation of movement cancellation. *J Neurophysiol* 96: 2011–2024, 2006. doi:[10.1152/jn.01323.2005](https://doi.org/10.1152/jn.01323.2005).
- Wattiez N, Poitou T, Rivaud-Péchoix S, Pouget P.** Evidence for spatial tuning of movement inhibition. *Exp Brain Res* 234: 1957–1966, 2016. doi:[10.1007/s00221-016-4594-8](https://doi.org/10.1007/s00221-016-4594-8).




Appendix B

Errington, S. P., Woodman, G. F., & Schall, J. D. (2020). Dissociation of medial frontal β -bursts and executive control. *Journal of Neuroscience*, *40*(48), 9272-9282.

Published November 25th 2020

Dissociation of Medial Frontal β -Bursts and Executive Control

Steven P. Errington, Geoffrey F. Woodman, and  Jeffrey D. Schall

Department of Psychology, Vanderbilt Vision Research Center, Center for Integrative and Cognitive Neuroscience, Vanderbilt Brain Institute, Vanderbilt University, Nashville, Tennessee 37240

The neural mechanisms of executive and motor control concern both basic researchers and clinicians. In human studies, preparation and cancellation of movements are accompanied by changes in the β -frequency band (15–29 Hz) of electroencephalogram (EEG). Previous studies with human participants performing stop signal (countermanding) tasks have described reduced frequency of transient β -bursts over sensorimotor cortical areas before movement initiation and increased β -bursting over medial frontal areas with movement cancellation. This modulation has been interpreted as contributing to the trial-by-trial control of behavior. We performed identical analyses of EEG recorded over the frontal lobe of macaque monkeys (one male, one female) performing a saccade countermanding task. While we replicate the occurrence and modulation of β -bursts associated with initiation and cancellation of saccades, we found that β -bursts occur too infrequently to account for the observed stopping behavior. We also found β -bursts were more common after errors, but their incidence was unrelated to response time (RT) adaptation. These results demonstrate the homology of this EEG signature between humans and macaques but raise questions about the current interpretation of β band functional significance.

Key words: countermanding; EEG; error monitoring; response inhibition; stop signal; stopping

Significance Statement

The finding of increased β -bursting over medial frontal cortex with movement cancellation in humans is difficult to reconcile with the finding of modulation too late to contribute to movement cancellation in medial frontal cortex of macaque monkeys. To obtain comparable measurement scales, we recorded electroencephalogram (EEG) over medial frontal cortex of macaques performing a stop signal (countermanding) task. We replicated the occurrence and modulation of β -bursts associated with the cancellation of movements, but we found that β -bursts occur too infrequently to account for observed stopping behavior. Unfortunately, this finding raises doubts whether β -bursts can be a causal mechanism of response inhibition, which impacts future applications in devices such as brain-machine interfaces.

Introduction

Response inhibition and performance monitoring are executive control functions supporting goal-directed behavior. The countermanding (stop-signal) task provides insights into these functions (Verbruggen and Logan, 2008; Verbruggen et al., 2019). Recently, researchers have described a higher incidence of β -bursts in electroencephalogram (EEG) recorded over medial

frontal cortex of humans during successful inhibition arising early enough to contribute to stopping the movement (Jana et al., 2020; Wessel, 2020). This result is difficult to reconcile with single-unit recordings within medial frontal cortex of macaques, which find neural signals modulating only after inhibition was achieved (Stuphorn et al., 2010) and most commonly after response inhibition errors (Stuphorn et al., 2000; Sajad et al., 2019).

The model of countermanding task performance as a race between GO and STOP processes (Logan and Cowan, 1984) offers clear criteria to attribute neural signals to movement initiation and inhibition when applied to electrophysiological (De Jong et al., 1995; Kok et al., 2004; Stahl and Gibbons, 2007; Godlove et al., 2011b; Reinhart et al., 2012; Swann et al., 2012; Wessel and Aron, 2015) and neurophysiological (Hanes et al., 1998; Paré and Hanes, 2003; Brown et al., 2008; Stuphorn et al., 2010; Schmidt et al., 2013; Brockett et al., 2020) measurements. To contribute to reactive response inhibition, a neural signal must be asserted before the STOP process finishes, and that

Received Aug. 6, 2020; revised Oct. 6, 2020; accepted Oct. 8, 2020.

Author contributions: J.D.S. and G.F.W. designed research; J.D.S. and G.F.W. performed research; S.P.E. and J.D.S. analyzed data; S.P.E., G.F.W., and J.D.S. wrote the paper.

This work was supported by the National Eye Institute Grant R01-EY019882, P30-EY08126 and by Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience. We thank B. Williams, R. Williams, M. Maddox, M. S. Schall, I. Haniff, S. Motomy, D. Richardson, L. Toy, and M. R. Feurtado for technical support. We also thank Dr. A. Bompas, Dr. K. A. Lowe, Dr. U. Rutishauser, Dr. A. Sajad, J. A. Westerberg, and Dr. T. Womelsdorf for useful conversations regarding the work.

The authors declare no competing financial interests.

Correspondence should be addressed to Jeffrey D. Schall at jeffrey.d.schall@vanderbilt.edu.

<https://doi.org/10.1523/JNEUROSCI.2072-20.2020>

Copyright © 2020 the authors

signal must scale in probability of occurrence proportional to the probability of canceling a response after a stop signal is presented. The spiking rate of movement-related neurons in cortical (Hanes and Schall, 1995; Murthy et al., 2009; Mirabella et al., 2011) and subcortical (Paré and Hanes, 2003; Schmidt et al., 2013) motor circuits satisfy these criteria. When preparing responses, the activity of movement-related neurons accumulates to a fixed threshold. Variation in the rate of this accumulation produces variation in response latency. Responses are withheld when the accumulation of activity is interrupted after the stop signal. The reliability of this relationship has been evaluated by comparing the psychometric inhibition function of the probability of responding despite the stop signal with a neurometric function of the probability of neural activity exceeding a threshold (Brown et al., 2008). The neurometric functions derived from movement neurons in FEF mirror the inhibition function.

Meanwhile, the spiking rate of neurons in medial frontal areas do not (Scangos and Stuphorn, 2010; Stuphorn et al., 2010). Instead, such signals are thought to contribute to executive control. As performance of the countermanding task includes failures to cancel the response on about half of all stop signal trials. This high failure rate also affords analysis of neural signals of error processing, a feature of executive control which can interact with proactive control.

To elucidate their origin, we have been establishing systematically the homology of visual and cognitive EEG signals in macaque monkeys and humans and locating contributing cortical areas (Woodman, 2012). We have established the homology of the error-related negativity (ERN) between macaques (Godlove et al., 2011b) and humans (Reinhart et al., 2012) and identified a medial frontal source for the ERN (Sajad et al., 2019). Here, we seek to establish the homology of β -bursts sampled in EEG recorded over medial frontal cortex of macaques and humans. If so, then subsequent investigation with invasive approaches in monkeys can offer mechanistic insights into the relationship between β -bursts and response inhibition.

We found the incidence of β -bursts observed in the EEG of macaque monkeys paralleled observations in humans. However, the probability of β -bursts lagged far behind the probability of canceling responses. Unexpectedly, β -bursts were most common after errors of inhibition, but their incidence was unrelated to adaptation in response times (RTs). We conclude that these β -bursts in macaque monkeys are homologous to those in humans, supporting further invasive investigation. It appears that β -bursts may serve as a rough index of executive control processes, but they lack any causal efficacy on the stopping response and thus much theoretical or practical utility.

Materials and Methods

Experimental model and subject details

All procedures were in accordance with the National Institutes of Health Guidelines, the American Association for Laboratory Animal Care *Guide for the Care and Use of Laboratory Animals* and approved by the Vanderbilt Institutional Animal Care and Use Committee in accordance with the United States Department of Agriculture and Public Health Service policies. Data were collected from one male bonnet macaque (Eu, *Macaca radiata*, 8.8 kg) and one female rhesus macaque (X, *Macaca mulatta*, 6.0 kg) performing a saccade countermanding task (Hanes and Schall, 1995; Godlove et al., 2014). Both animals were on a 12/12 h light/dark cycle and all experimental procedures were conducted in the daytime. Each monkey received

nutrient-rich, primate-specific food pellets twice a day. Fresh produce and other forms of environmental enrichment were given at least five times a week.

While human studies obtain data from more participants, for practical and regulatory reasons, data cannot be collected from as many monkeys. However, 25 years of investigating the stop signal task with macaque monkeys has shown that their performance matches in various nuanced details that of humans (Hanes and Schall, 1995) and the race model accounts for human and macaque performance equivalently (Boucher et al., 2007; Camalier et al., 2007). Investigations of this task by multiple nonhuman primate laboratories have found no differences beyond the incidental individual differences that are evident in human performance.

Surgical procedures

Surgical details have been described previously (Godlove et al., 2011a). Briefly, magnetic resonance images (MRIs) were acquired with a Philips Intera Achieva 3T scanner using SENSE Flex-S surface coils placed above or below the animal's head. T1-weighted gradient-echo structural images were obtained with a 3D turbo field echo anatomic sequence (TR = 8.729 ms; 130 slices, 0.70 mm thickness). These images were used to ensure Cilux recording chambers were placed in the correct area (Crist Instruments). Chambers were implanted normal to the cortex (monkey Eu: 17°; monkey X: 9°; relative to stereotaxic vertical) centered on midline: 30 mm (monkey Eu) and 28 mm (monkey X) anterior to the interaural line.

Data collection protocol

An identical daily recording protocol across monkeys and sessions was conducted. In each session, the monkey sat in an enclosed primate chair with their head restrained 45 cm from a CRT monitor (Dell P1130, background luminance of 0.10 cd/m²). The monitor had a refresh rate of 70 Hz, and the screen subtended 46° × 36° of the visual angle. Eye position data were collected at 1 kHz using an EyeLink 1000 infrared eye-tracking system (SR Research). All data were streamed to a single data acquisition system (MAP, Plexon). Time stamps of trial events were recorded at 500 Hz.

Macaque electroencephalography

The EEG was recorded from the cranial surface with an electrode located over medial frontal cortex. The electrode implants were constructed from Teflon-coated braided stainless-steel wire and solid-gold terminals. Implanted wires were cut to 8.5 cm, the wire ends exposed, and gold Amphenol pins were crimped to both ends. One end of the wires was inserted into a plastic connector, whereas the gold pin on the other end was ground down until 1 mm of the pin remained. During aseptic surgery, a ~1-mm hole was drilled into the surface of the skull (3–5 mm thick), allowing the terminal end of the electrode to be tightly inserted. The inserted gold pin was then covered with a small amount of acrylic cement. After the EEG electrode was implanted, the plastic connector was attached to exposed acrylic to allow access to the channels. Leads that were not embedded in the acrylic were covered by skin that was sutured back over the skull. This allowed for the EEG electrode to be minimally invasive once implanted. Unlike recordings from skull screws that extend to the dura mater through the skull, recordings from these electrodes approximate those used in human electrophysiological studies because the signals must propagate through the layers of brain, dura, and skull. Electrodes were referenced to linked ears using ear-clip electrodes (Electro-Cap International). The EEG from each electrode was amplified with a high-input impedance head stage (Plexon) and band-pass filtered between 0.7 and 170 Hz. All data were streamed to a data acquisition system (MAP, Plexon).

Saccade stop-signal (countermanding) task

The saccade stop-signal task used in this study has been widely used previously (Hanes and Schall, 1995; Hanes and Carpenter, 1999; Cabel et al., 2000; Colonius et al., 2001; Kornlyo et al., 2003; Morein-Zamir and Kingstone, 2006; Walton and Gandhi, 2006; Thakkar et al., 2011, 2015; Godlove and Schall, 2016; Wattiez et al., 2016; Verbruggen et al., 2019).

Recently, a set of guidelines has been proposed for designing and analyzing the stop-signal task to allow for valid comparisons to be made across studies (Verbruggen et al., 2019). Our study followed all of the recommendations but two. These adjustments were necessary to obtain sufficient neural data and to address issues arising because monkeys gain so much more experience with the task parameters relative to human participants. First, 40% of trials in our study were stop trials, compared with the recommended 25%. This higher value was used to achieve the necessary power to analyze neural data at the individual stop-signal level, but it did not introduce excessive slowing of responses (Emeric et al., 2007). Second, although we employed a staircase procedure, this stepped by one to three stop-signal delays (SSDs). We do this to prevent monkeys from anticipating the staircase (Nelson et al., 2010).

Briefly, trials were initiated when monkeys fixated a central point. Following a variable time period, the center of the fixation point was removed leaving an outline. At this point, a peripheral target was presented simultaneously on either the left or right hand of the screen. In this study, one target location was associated with a larger magnitude of fluid reward. The lower magnitude reward ranged from 0% to 50% of the higher magnitude reward amount. This incidence was adjusted to encourage the monkey to continue responding to both targets. The stimulus-response mapping of location-to-high reward changed across blocks of trials. Block length was adjusted to maintain performance at both targets, with the number of trials in each block determined by the number of correct trials performed. In most sessions, the block length was set at 10–30 correct trials. Erroneous responses led to repetitions of a target location, ensuring that monkeys did not neglect low-reward targets in favor of high-reward targets, a phenomenon demonstrated in previous implementations of asymmetrically rewarded tasks (Kawagoe et al., 1998).

On most of the trials, the monkey was required to make an eye movement to this target (no-stop trials). However, on a proportion of trials the center of the fixation point was re-illuminated (stop-signal trials); this stop signal appeared at a variable time after the target had appeared (SSD). An initial set of SSDs, separated by either 40 or 60 ms, was selected for each recording session. The delay was then manipulated through an adaptive staircasing procedure in which stopping difficulty was based on performance. When a subject failed to inhibit a response, the SSD was decreased by a random step to increase the likelihood of success on the next stop trial. Similarly, when subjects were successful in their inhibition, the SSD was increased to reduce the likelihood of success on the next stop trial. This procedure was employed to ensure that subjects failed to inhibit action on ~50% of all stop-signal trials. On no-stop trials, the monkey was rewarded for making a saccade to the target. On stop-signal trials, the monkey was rewarded for withholding the saccade and maintaining fixation on the fixation spot. Following a correct response, an auditory tone was sounded 600 ms later, and followed by a high or low fluid reward, depending on the stimulus-response mapping.

Experimental design and statistical analysis

Bayesian modeling of stop-signal performance

As performance on the stop-signal task can be considered as the outcome of a race between a GO and STOP process, then a stop-signal reaction time (SSRT) can be calculated (Logan and Cowan, 1984). This value can be considered as the latency of the inhibitory process that interrupts movement preparation.

SSRT was estimated using a Bayesian parametric approach (Matzke et al., 2013a,b). Compared with classical methods of calculating SSRT (i.e., integration-weighted method; Logan and Cowan, 1984), this approach allows for a distribution of SSRT to be derived by using the distribution of reaction times on no-stop trials, and by considering reaction times on non-canceled trials as a censored no-stop RT distribution. Furthermore, this model also allows for the estimation of the probability of trigger failures for a given session (Matzke et al., 2017). Individual parameters were estimated for each session. The priors were bounded uniform distributions [μ_{Go} , μ_{Stop} : $U(0.001, 1000)$; σ_{Go} , σ_{Stop} : $U(1, 500)$; τ_{Go} , τ_{Stop} : $U(1, 500)$; pTF: $U(0, 1)$]. The posterior distributions were estimated using Metropolis-within-Gibbs sampling and we ran multiple (3) chains. We ran the model for 5000 samples with a thinning of 5.

EEG processing and β -burst detection

For each session, raw data were extracted from the electrode. This signal was then bandpass filtered between 15 and 29 Hz. This signal was then epoched from –1000 to 2500 ms relative to multiple key events in a trial, including target onset, saccade, and stop-signal presentation. β -Burst detection was performed as previously described (Shin et al., 2017; Wessel, 2020). The description is adapted from therein. We then convolved the epoched signal for each trial with a complex Morlet wavelet of the form:

$$w(t, f) = A \exp\left(\frac{t^2}{2\sigma_t^2}\right) \exp(2i\pi ft),$$

with $\sigma = \frac{m}{2\pi f}$, $A = \frac{1}{\sigma} \sqrt{2\pi}$, and $m = 7$ (cycles) for each of the 15 evenly spaced frequencies spanning the β band (15–29 Hz). Time-frequency power estimates were extracted by calculating the squared magnitude of the complex wavelet-convolved data. Individual β -bursts were defined as local maxima in the trial-by-trial band time-frequency power matrix, for which the power exceeded a threshold of 6-times the median power of the entire time-frequency power matrix for the electrode. To compute the burst % across trials, we binary coded the time of the peak β -amplitude. A β -burst density function was generated by convolving the binary-coded array of β -burst activity with a Gaussian function of the form:

$$\beta - \text{bdf}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2},$$

where $\mu = 00.0$ ms, and $\sigma = 22.5$ ms. These values represent the time required for half a cycle at the median β -frequency.

Behavioral comparisons between monkeys

For each session, we extracted the mean response latencies on no-stop and non-canceled trials. After using the Bayesian approach described above, we extracted estimates of the SSRT mean and SD for each session. We also extracted estimates of the proportion of trigger failures for each session. Mean values between monkeys were compared using a one-way independent measure ANOVA. Greenhouse–Geisser corrections were applied when assumptions of sphericity were violated.

Comparing β -bursts during stopping

To examine how the β -bursts activity may vary dependent on trial type, the incidence of β -bursts observed during the stopping process were calculated for each session. The STOP process interval was defined as the time between SSD onset and SSRT. While stop trials by definition had a predefined SSD associated with them, no-stop trials were assigned an SSD value similar to that used in the most recent stop-trial. We compared this activity against a baseline period. This period was an equivalent interval of time ranging from –200 ms before the target onset, to –200 ms minus the mean SSRT on the given session. For each trial type and time period (baseline and stopping), we then calculated the proportion of trials in which at least one β -burst occurred. A two-way repeated measures ANOVA was conducted, with time window and trial type as factors, and the proportion of β -bursts as the dependent variable. This approach allowed us to determine whether β -bursts were more prevalent during particular trial types, and whether these events were clearly task-related activity. *Post hoc* tests were conducted if ANOVAs were statistically significant. To determine the time at which the incidence of β -bursts differentiated between non-canceled and canceled trials, we found the first time point at which the 95% confidence intervals (CIs) for the β -burst density functions no longer overlapped. Greenhouse–Geisser corrections were applied when assumptions of sphericity were violated.

Linking incidence of β -bursts to response inhibition

To examine how neural function may reflect changes in stopping behavior, we looked at how the incidence of β -bursts varied with the probability of inhibiting a movement. This analysis was limited to SSDs with 15

or more canceled trials. At each threshold, we subtracted the proportion of the β -bursts observed at a given SSD from the p (respond | stop-signal) at the same SSD. This difference between bursts and p (respond | stop-signal) was squared and values in the given session were summed, creating a sum of squared error between the two measures for each session and at each burst threshold. We then examined whether these values across sessions significantly differed from zero using a one-sample t test at each threshold. We performed this analysis on both raw measures of β -burst proportions, and normalized measures, where incidences were relative to the maximum proportion of β -bursts observed. Findings were the same across both approaches.

Linking β -burst incidence to error monitoring

Error-related activity was examined by comparing the incidence of bursts on non-canceled trials at the middle most SSD to latency matched no-stop trials. Previous work from our lab has highlighted error-related spiking activity from this dataset became most prominent in the 100- to 300-ms period following an erroneous saccade (Sajad et al., 2019). As such, we calculated the incidence of β -bursts that occurred during this period. Across sessions we compared the incidence of bursts observed in error trials against those observed in trials where a saccade was correctly executed using a one-way repeated measures ANOVA. Greenhouse–Geisser corrections were applied when assumptions of sphericity were violated.

Linking β -bursts to post-error RT adaptation

To examine how β -bursts may contribute to RT adaptations following errors or successful inhibition, we first quantified an index to capture the degree of slowing for each session. For post-error slowing, this was done by dividing the mean RT on no-stop trials following error trials by the mean RT on no-stop trials following no-stop trials for a given session. This value represents the proportional change in RT resultant from an error. We repeated this approach for post-stopping slowing, instead using the mean RT in no-stop trials following canceled trials as the numerator in this ratio in which the superscript identifies the current trial type, and the subscript identifies the preceding trial type:

$$\text{Post - error slowing index} = \frac{RT_{\text{Non-canceled}}^{\text{No-stop}}}{RT_{\text{No-stop}}^{\text{No-stop}}};$$

$$\text{Post - canceled slowing index} = \frac{RT_{\text{Canceled}}^{\text{No-stop}}}{RT_{\text{No-stop}}^{\text{No-stop}}}.$$

For each monkey, we examined how the incidence of β -bursts observed in the error monitoring period of a given session varied with the given sessions post-error index. To determine the association between post-error β -burst activity and post-error slowing, we fit a generalized linear model. From this we extracted R^2 values and determined whether the observed slope was significant.

We compared the effects of previous trial outcome (trial $n-1$) on β -burst activity observed in the baseline of the following trial (trial n). We identified no-stop trials which immediately followed canceled, non-canceled, and no-stop trials and calculated the proportion of these trials in which β -bursts occurred in the -400 - to -200 -ms period before the target appearing. We then compared whether the proportion of β -bursts during this baseline period differed between different trial types using a one-way repeated measures ANOVA.

Results

We acquired 33,816 trials across 29 sessions from two macaques (Eu: 11,583; X: 22,233) performing the saccade stop-signal (countermanding) task (Fig. 1A). Both monkeys exhibited typical sensitivity to the stop-signal. Summary measures of performance are in Table 1. First, response latencies on non-canceled (error) trials were faster than those on no-stop trials (Fig. 1B, left). Second, the probability of failing to cancel and executing an erroneous saccade was greater at longer SSDs (Fig. 1B, right). These

two observations validated the assumptions of the independent race model (Logan and Cowan, 1984), allowing us to estimate the SSRT, the time needed to cancel to partially prepared saccade.

Previous studies used SSRT for distinguishing whether neural signals can contribute directly to reactive control, so estimates of this duration must be accurate and precise. We calculated SSRT using a Bayesian parametric approach (Matzke et al., 2013a,b), which offers estimates of the SSRT for each session. The monkeys had indistinguishable mean SSRT (one-way independent measure ANOVA: $F_{(1,27)} = 0.108$, $p = 0.745$, $BF_{10} = 0.367$) and variance of SSRT (one-way independent measures ANOVA: $F_{(1,27)} = 0.819$, $p = 0.819$, $BF_{10} = 0.360$). This approach also quantified trigger failures when stopping was unsuccessful because the STOP process was not initialized. Trigger failures were significantly more common for monkey Eu relative to monkey X (one-way independent measures ANOVA: $F_{(1,27)} = 18.458$, $p < 0.001$, $BF_{10} = 114.778$).

β -Bursts and response inhibition

The monkeys' EEG was recorded with a lead placed on the cranial surface over the medial frontal cortex at location analogous to FCz in humans (Fig. 1C). At the individual trial level, β band activity was characterized by obvious, burst-like events, rather than by steady changes in modulations (Fig. 1D). As observed in human studies (Jana et al., 2020; Wessel, 2020), the overall prevalence of these bursts was low during both baseline (-400 to -200 ms pretarget, $\sim 12.6 \pm 3.8\%$ across all sessions) and task-relevant (0 – 200 ms post-target, $\sim 16.1 \pm 5.2\%$ across all sessions) periods.

Following previous studies (Jana et al., 2020; Wessel, 2020), to examine the relationship between β -burst activity and stopping behavior in the countermanding task, we first compared the prevalence of β -bursts across trial types (two-way repeated measures ANOVA with time window and trial type as factors, Greenhouse–Geisser corrected: $F_{(1.72,48.31)} = 9.816$, $p < 0.001$, $BF_{10} = 53.763$; Fig. 2A; Table 2). We found no significant changes in the incidence of β -bursts in a baseline period and during the stop process on non-canceled trials (Holm *post hoc* test, adjusted $p = 0.300$) or during an equivalent period of time when stopping would have occurred on no-stop trials (Holm *post hoc* test, adjusted $p > 0.999$). However, compared with a baseline period, β -bursts were significantly more common during the STOP process when a movement was successfully canceled (Holm *post hoc* test, adjusted $p = 0.019$). Furthermore, β -bursts were significantly more common during the STOP process on canceled compared with non-canceled trials (*post hoc* test, adjusted $p < 0.001$) but not compared with an equivalent period of time on no-stop trials (*post hoc* test, adjusted $p = 0.057$). This pattern of β -burst incidence replicates previous reports from human participants (Jana et al., 2020; Wessel, 2020).

Neurophysiological investigations quantify neural signals on a finer time scale using spike density functions (Hanes et al., 1998). To examine how changes in β -bursts occur over time, we derived β -burst density functions. First, we binary coded the time of the peak β -amplitude (Fig. 2B). A β -burst density function was determined by convolving this discretized array with a Gaussian function over time since the stop-signal (Fig. 2C). These plots reveal more information about the dynamics of β -burst production through trial time and across trial types. β -Burst frequency increases through the trial during saccade preparation. On no stop trials and non-canceled trials, β -burst

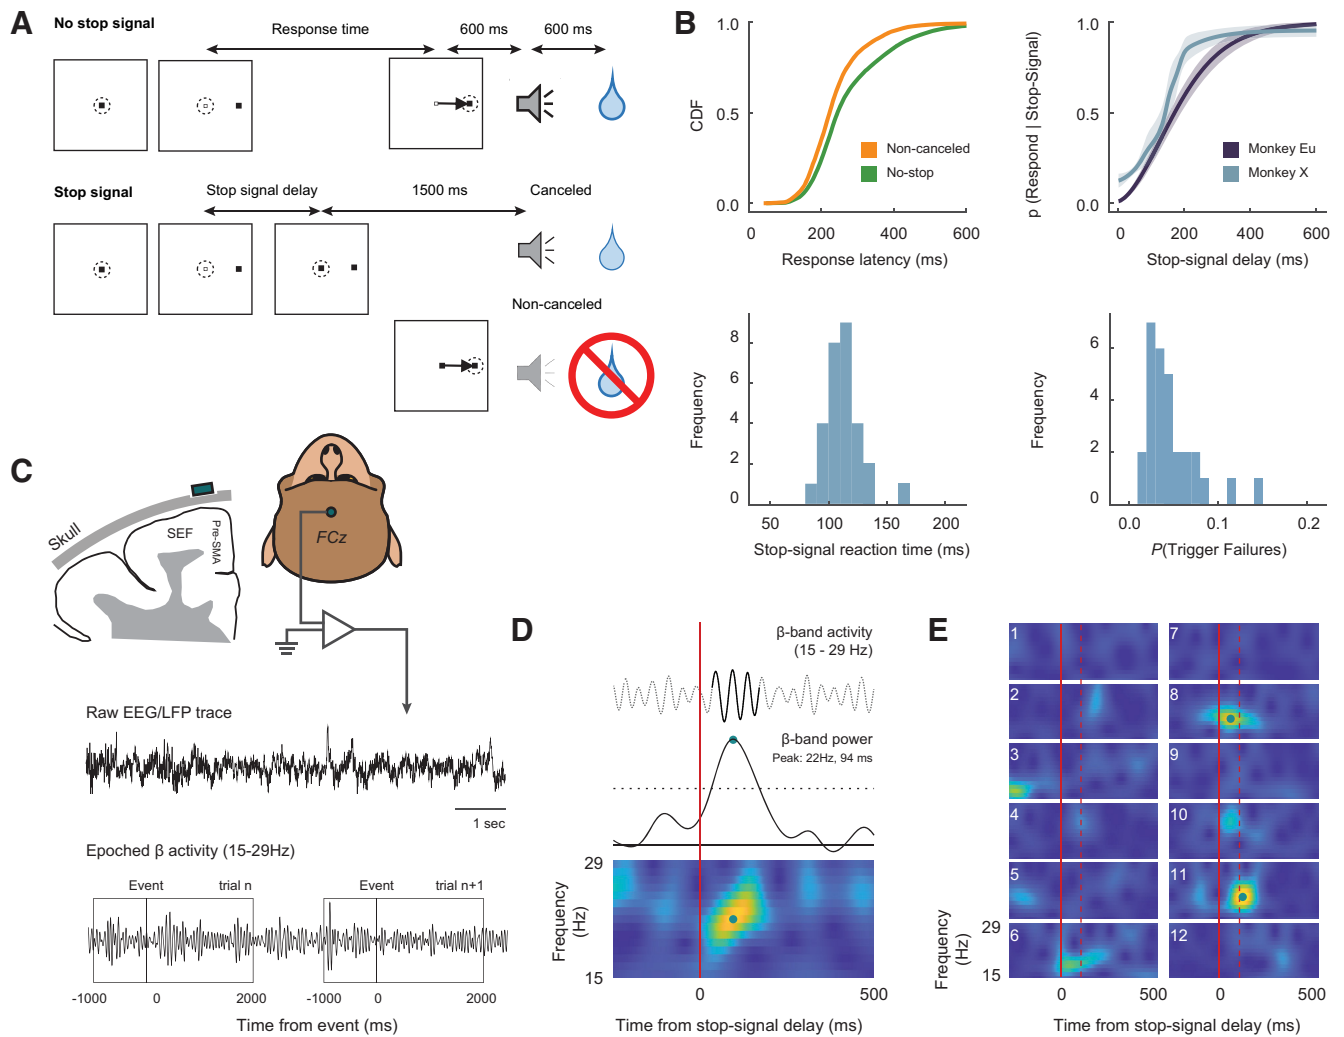


Figure 1. Experimental procedures. **A**, Saccade-countermanding task. Monkeys initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target was presented simultaneously at one of two possible locations. On no-stop-signal trials monkeys were required to shift gaze to the target, whereupon after 600 ± 0 ms a high-pitched auditory feedback tone was delivered, and 600-ms later fluid reward was provided. On stop-signal trials ($\sim 40\%$ of trials), after the target appeared the center of the fixation point was re-illuminated after a variable SSD, which instructed the monkey to cancel the saccade in which case the same high-pitched tone was presented after a 1500 ± 0 ms hold time followed, after 600 ± 0 ms by fluid reward. SSD was adjusted such that monkeys successfully canceled the saccade in $\sim 50\%$ of trials. In the remaining trials, monkeys made non-canceled errors which were followed after 600 ± 0 ms by a low-pitched tone, and no reward was delivered. Monkeys could not initiate trials earlier after errors. **B**, Countermanding behavior. Top left, Cumulative distribution function of response latencies on no-stop (green) and non-canceled (yellow) trials. Response latencies on non-canceled trials were faster than those on no-stop trials. Top right, Inhibition function plotting the probability of responding across SSDs. Weibull functions were fitted to data from each session. The mean of these Weibull functions across sessions and the corresponding 95% CI is plotted for each monkey (monkey Eu: purple; X: blue). Bottom left, Distribution of mean SSRTs across sessions. Bottom right, Distribution of the proportion of trigger failures across sessions. **C**, LFP processing. EEG was recorded with leads placed on the cranial surface over the medial frontal cortex at the location analogous to FCz in humans. The EEG lead was located over the supplementary eye field (SEF) and pre-supplementary motor area (pre-SMA). For each session, raw data were extracted. After bandpass filtering between 15 and 29 Hz, this signal was epoched from -1000 to 2500 ms relative to target presentation, saccade initiation, and stop-signal presentation. **D**, β -Burst processing. The epoched signal for each trial was convolved with a complex Morlet wavelet. Time-frequency power estimates were extracted by calculating the squared magnitude of the complex wavelet-convolved data. Individual β -bursts were defined as local maxima in the trial-by-trial band time-frequency power matrix, for which the power exceeded a threshold of six times the median power of the entire time-frequency power matrix for the electrode. An example burst is shown in the time-frequency plot at the bottom. **E**, Examples of β band time frequency in 12 randomly selected trials, aligned on SSD (solid red), with the corresponding SSRT (dashed red). These plots are indistinguishable from counterparts derived from human data.

Table 1. Stop-signal task performance (mean \pm SEM) for both monkeys across all sessions

	Monkey Eu ($n = 12$ sessions)	Monkey X ($n = 17$ sessions)
No-stop RT (ms)	313.4 ± 1.6	263.0 ± 1.0
Non-canceled RT (ms)	259.4 ± 2.0	229.6 ± 1.0
SSRT _{mean} (ms)	112.4 ± 6.4	114.3 ± 1.8
SSRT _{std} (ms)	30.5 ± 1.8	30.9 ± 1.0
p (trigger failures)	0.069 ± 0.010	0.032 ± 0.003

frequency decreases following saccade initiation. The decrease in non-canceled trials begins earlier because the latency of non-canceled saccades is systematically less than that of no-stop trials. However, after noncancelled errors the incidence of β -bursts observed increases markedly. This will be characterized further below.

The incidence of β -bursts differentiated between correctly inhibited or incorrectly executed stop trials on average across sessions 132 ms after a stop-signal appeared. However, unlike the

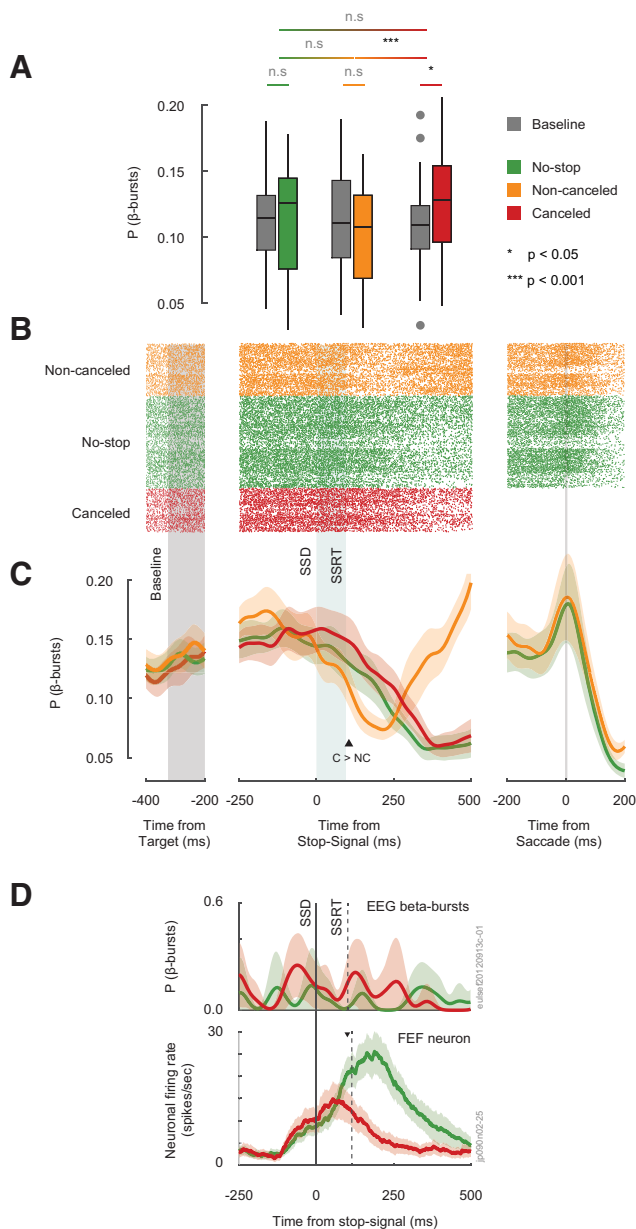


Figure 2. β -Bursts during stopping. **A**, Boxplots showing the incidence of β -bursts observed during the STOP process interval (from scheduled SSD to SSRT) during no-stop (green), non-canceled (orange), canceled (red), and during an equivalent period of time before target presentation (gray). β -Bursts are observed in ~10–15% of trials and are slightly but significantly more commonly observed when saccades are inhibited. **B**, Raster plot of β -bursts aligned on a pretarget baseline interval (left), stop signal (middle), and saccade initiation (right) across all sessions. Each tick-mark shows the time of peak β -amplitudes satisfying inclusion criteria on each trial. Rasters are shown for non-canceled, no-stop, and canceled trials. The rough equivalence of β -burst frequency across types of trials is evident, as is the elevation of β -burst rate at the end of non-canceled error trials. **C**, β -Burst density function derived from raster plots. β -Burst peak times were convolved with a Gaussian function. During the stopping period, β -bursts were slightly but significantly more common on canceled trials (red line) than on no-stop (green) or non-canceled trials (yellow). **D**, Comparing time course of β -burst (top) and single neuron discharges (bottom) on canceled and latency-matched no stop signal trials for a single session. At no time did the incidence of β -bursts on single session differentiate between movement initiation and inhibition. In contrast, as demonstrated previously, the discharge rate of an example FEF movement neuron sampled in one session shows a clear separation between trial types occurring before the STOP process concludes. The neuron was recorded from another monkey in a separate study performing a choice countermanding task (Middlebrooks et al., 2020) and is provided as an example to demonstrate the mechanistic differences between the signals.

Table 2. Percentage of trials (mean \pm SEM) with β -bursts during a baseline and stopping period, for all trial types

	No-stop	Non-canceled	Canceled
Baseline	11.6 \pm 0.6%	11.7 \pm 0.7%	11.2 \pm 0.7%
Stopping period	11.8 \pm 0.8%	10.5 \pm 0.7%	13.1 \pm 0.8%

discharge rates of neurons causally involved in movement initiation and inhibition (Middlebrooks et al., 2020; Fig. 2D, bottom panel), β -burst density functions were so small and noisy that no time distinguished canceled from no stop trials in individual sessions (Fig. 2D, upper panel). Notably, across sessions, β -burst incidence decreases after SSRT in canceled trials. If β -bursts are supposed to enforce response inhibition, this decay is curious. Because monkeys must sustain fixation for 1500 ms, response inhibition is sustained long after β -bursts cease. If stopping consists of multiple processes, one responsible for the interruption of the motor plan and another for maintenance of the inhibition goal (Logan and Cowan, 1984; Bompas et al., 2020), perhaps β -bursts initiate the stop process but do not maintain the inhibition.

The response inhibition function plots the fraction of non-canceled trials in which saccades are produced as a function of stop signal delay (Fig. 1B). The fraction of non-canceled trials is an increasing function of stop signal delay, because movements become less likely to cancel as movement preparation progresses. Using single neuron discharge rates, a neurometric function plots the probability of modulating within SSRT as a function of SSD. The neurometric function derived from the single neuron discharges of movement-related neurons parallels the inhibition function (Brown et al., 2008). Although this relationship has only been demonstrated for neural activity that initiates responses, the relationship with neural activity that instantiates the STOP process is just the inverse with greater likelihood of modulation when the probability of inhibition is highest.

We determined whether a neurometric function derived from β -bursts parallels the probability of inhibiting a prepared response. For each session, we measured the number of β -bursts observed on canceled trials at each SSD, 50 ms before SSRT, using the conventional threshold of $6\times$ median amplitude plus lower ($2\times$ median) and higher ($10\times$ median) thresholds. This β -burst neurometric function was compared with the probability of inhibiting a response at each SSD (Fig. 3A) quantitatively through the sum of their squared differences at each stop signal delay. Summed squared differences close to zero indicate similarity of the two measures. Across sessions, the distribution of summed squared differences using the $6\times$ median threshold was significantly different from zero (one-sample t test: $t_{(28)} = 6.70$, $p < 0.001$, $BF_{10} = 56,120.13$; Fig. 3B). This conclusion did not depend on β -burst measurement threshold, for even with the lowest threshold finding β -bursts in only 40% of canceled trials. No β -burst measurement threshold produced a neurometric function similar to the inhibition function (one-sample t test at each threshold, $p < 10^{-10}$ for all thresholds after corrections for multiple comparisons; Fig. 3B).

β -Bursts, error monitoring, and executive control

The stop-signal task is useful for exploring performance monitoring because, by design, errors occur in 50% of stop-signal trials (here, 40% of all trials). Using this task, previous work has demonstrated neural activity in supplementary eye field (SEF) that occurs following errors, the magnitude of which is predictive

of changes in response latencies in the following trial (Stuphorn et al., 2010; Sajad et al., 2019). Hence, we compared the incidence of β -bursts 100–300 ms after error and correct saccades, when spiking activity related to errors is maximal. β -Bursts were significantly more prevalent on error trials ($11.1 \pm 0.7\%$) compared with correct trials ($6.7 \pm 0.5\%$) during this period (one-way repeated measures ANOVA: $F_{(1,28)} = 55.103$, $p < 10^{-3}$, $BF_{10} = 9068.665$; Fig. 4A).

Behaviorally, RT on a trial varies according to the outcome of the previous trial (Emeric et al., 2007). Both monkeys produced longer RT in no-stop trials following erroneous non-canceled trials (one-way repeated measures ANOVA with previous trial type as factor, Greenhouse–Geisser corrected, $F_{(1.54,43.30)} = 226.341$, $p < 0.001$, $BF_{10} = 6.785e + 52$, Holm *post hoc* comparison between no-stop and non-canceled: $p < 0.001$, $BF_{10} = 2.273e + 7$). To examine whether post-saccade β -bursts influence post-error slowing, we calculated a post-error slowing index for each session by dividing the mean RT on no-stop trials following non-canceled trials, by the mean RT on no-stop trials following another no-stop trial (Fig. 4B). This value measures the relative change in RT on no-stop trials dependent on the previous trial type. We found that the incidence of β -bursts after errors bore no relation to RT on the following trial for either monkey (monkey Eu: $R^2 = 0.0333$, $p = 0.571$, $BF_{10} = 0.521$; monkey X: $R^2 = 0.0732$, $p = 0.294$, $BF_{10} = 0.624$; Fig. 4C).

Given the outcome-dependent adaption in response latency, we investigated whether the outcome on the previous trial, influenced β -burst occurrence during the baseline period on the following trial. We found no significant effect of previous trial outcome on the proportion of β -bursts observed in the baseline of a following no-stop trial (one-way repeated measures ANOVA: $F_{(1.47,41.04)} = 3.09$, $p = 0.071$, $BF_{10} = 1.22$).

Discussion

We found that the prevalence and timing of the β -bursts do not account for the likelihood of canceling a planned response. Mirroring findings with human participants, we found macaque monkeys exhibit small but consistent pulses of β activity recorded in EEG over medial-frontal cortex during the inhibition of prepared movements (Jana et al., 2020; Wessel, 2020). As previous findings, we observed β -bursts infrequently ($\sim 15\%$ of trials) and not uncommonly in trials in which a response was generated. If β -bursts cause response inhibition, then they should be more prevalent at earlier SSDs where inhibition is more successful. This relationship between neurometric and psychometric measures of response inhibition has been observed in the discharges of movement-related neurons in frontal eye fields (Brown et al., 2008). However, in the mechanistic, interactive race models of saccade countermanding, unlike the progressive activation of the GO unit over time and across stop signal delays, the action of the STOP unit is effectively all-or-none (Boucher et al., 2007; Logan et al., 2015). Thus, an effective mechanism of response inhibition must happen on every trial in which a response is inhibited. Collectively and unfortunately, these

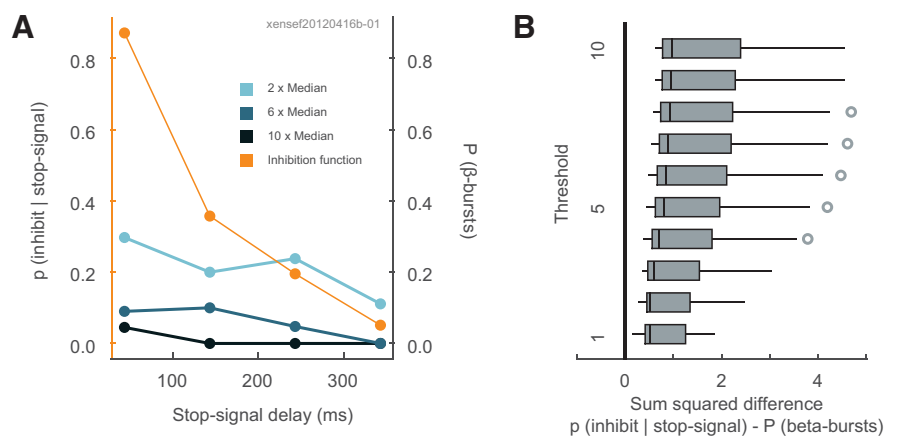


Figure 3. No relationship between β -bursts and response inhibition. **A**, Primary ordinate, average probability for a representative session of inhibiting on stop signal trials as a function of stop signal delay (orange line). Secondary ordinate, average probability of β -bursts determined with amplitudes exceeding the median threshold by 2 (light blue), 6 (blue), and 10 (dark blue) times in that session. The neurometric function derived from the probability of β -bursts does not correspond to the probability of inhibiting as a function of stop signal delay. **B**, Boxplots of the sum of squared differences between probabilities of canceling and of β -bursts across sessions determined with amplitudes exceeding from one to ten times the median threshold. Open circles indicate outlier values. From generous to severe measurement thresholds, the incidence of β -bursts did not account for the probability of cancellation with stop signal delay.

observations raise doubts about the proposal that β -bursts are a causal mechanism of response inhibition, and limit future applications in devices such as brain-machine interfaces.

Previous findings of β -bursts over frontal cortex have been interpreted as part of a larger framework proposing that the inferior frontal gyrus of the right hemisphere (rIFG) and the presupplementary motor area (pre-SMA) contribute to reactive control through the hyper-direct pathway (Aron and Poldrack, 2006; Aron et al., 2007). While the rIFG has been implicated in the attentional capture of the stop-signal and initiating the STOP process (Swann et al., 2012; Jana et al., 2020), the role of the pre-SMA is less clear. Human fMRI studies show greater activity in pre-SMA on stop trials with manual responses (Aron and Poldrack, 2006; Aron et al., 2007; Rae et al., 2015) and in supplementary eye field with saccades (Thakkar et al., 2014). Lesion studies also associate medial frontal areas with impaired stopping of limbs (Floden and Stuss, 2006; Nachev et al., 2007; Sumner et al., 2007) and eyes (Husain et al., 2003). This is mirrored in human electrophysiological evidence reporting stronger signals over pre-SMA during canceled trials (Swann et al., 2012). However, our observation that β -bursts are slightly more common during response inhibition is unexpected in macaques based on previous neurophysiological results. In single-unit recordings, the modulation of neurons in MFC occurs too late to contribute to the reactive control of movement and instead contributes to performance monitoring and the exertion of proactive control (Emeric et al., 2010; Stuphorn et al., 2010).

Such discrepancies have sparked debate about whether macaques are a useful model of executive control in humans (Cole et al., 2009; Schall and Emeric, 2010). The countermanding task has the advantage for comparisons between macaques and humans because across species it has been tested with the same response modalities, task designs, and measurement scales. Nevertheless, cross-species comparisons require caution because of differences in how the data are collected. First, like previous studies, we cannot localize β -bursts in scalp data to specific cortical areas, given the spatial resolution of EEG. To draw comparisons with human studies, we have to focus on literature that highlights the role of pre-SMA. These previous studies used tasks

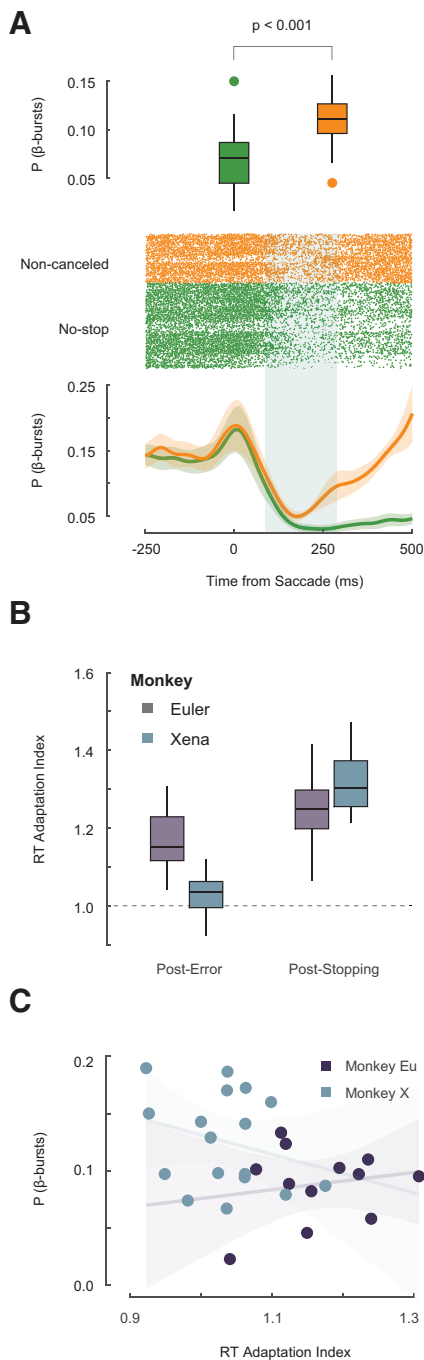


Figure 4. Relationship between β -bursts and performance monitoring. **A**, β -Bursts during error and correct responses. Top, Boxplot showing significantly greater incidence of β -bursts observed 100–300 ms following non-canceled (orange) compared with correct saccades (green). Middle, Raster plot aligned on saccade. Each tick-mark shows the time of one β -burst on a non-canceled (yellow) and no-stop (green) trials. Bottom, β -Burst density function. Following a saccade, the incidence of bursts on both error and correct trials decreases. This is followed by a pronounced increase in β -burst frequency \sim 150 ms after noncancelled error saccades. **B**, Boxplots of RT adaptation following non-canceled errors and successful cancellations across sessions for Eu (purple) and X (cyan). Values greater than one represent slowing. As observed previously, both monkeys tend to delay responses somewhat after errors and more after successful stopping. **C**, β -Bursts are unrelated to RT adaptation. The incidence of β -bursts observed after errors did not vary as a function of the post-error RT adaptation index across sessions for Eu (purple) or X (cyan). Non-significant regression lines (\pm 95% CI) include 0 slope.

which typically required reaching or manual movements. As our study used a saccade-countermanding task, we expect activity to be present in the supplementary eye field, considered to be the “eye area” of the pre-SMA. SEF in macaques is homologous to SEF in humans, as SMA/pre-SMA in macaques is homologous to SMA/pre-SMA in humans (Amiez and Petrides, 2009). In macaques, the pre-SMA is located in the cortex forming the medial wall above cingulate cortex. Thus, from a biophysical standpoint, it is unlikely that the electrical field generated in this area would be captured through the EEG we record. However, SEF is on the dorsomedial convexity and activity in such an area is evident in the overlying EEG (Sajad et al., 2019). Second, compared with experiments with humans, monkeys are well trained and well accustomed to the task, typically completing up to 10 times more trials than humans. This may be demonstrated through our observation that monkeys have a much smaller proportion of trigger failures (\sim 5%) compared with what has been previously observed in human (\sim 20%) studies (Skippen et al., 2019, 2020). However, although RTs and SSRTs are typically shorter in macaques than humans, the relative relationship between the two is similar between species. Indeed, our first investigation of the stop signal task with macaque monkeys established that their performance matches in various nuanced details that of humans (Hanes and Schall, 1995). Moreover, instantiations of the race model account for both human and macaque countermanding performance equivalently (Boucher et al., 2007; Camalier et al., 2007).

While we found β -bursts were more common during response inhibition, we also found they were more common following errors, suggesting a contribution to performance monitoring and executive control. This observation is consistent with patterns of spiking activity described in medial frontal cortex (Stuphorn et al., 2000; Sajad et al., 2019) and complementing previous observations that error, conflict, and reward monitoring are typically associated with θ (4–8 Hz) band modulation (Luu et al., 2004; Trujillo and Allen, 2007; Cohen et al., 2008; Cavanagh et al., 2009; Nigbur et al., 2011; Amarante et al., 2017). Over medial frontal cortex, β power is elevated when cognitive control was required and following negative feedback (Stoll et al., 2016). Specifically, in a change-stop-signal task, previous work has highlighted greater β activity on canceled compared with non-canceled trials in pre-SMA, following the completion of the STOP process (Jha et al., 2015). This has also been demonstrated in a combined Flanker stop-signal task, in which β activity increased following the commission of an error (Marco-Pallarés et al., 2008; Beyer et al., 2012). Interestingly, the magnitude of this β activity correlated with the θ -activity underlying the ERN (Marco-Pallarés et al., 2008).

While it has been previous argued that there is mixed evidence for a relationship between the error-related negativity and performance adjustments (Gehring et al., 1993; Gehring and Fencsik, 2001; Rodriguez-Fornells et al., 2002; Hajcak et al., 2003; Kerns et al., 2004; Holroyd et al., 2005; Ladouceur et al., 2007; West and Travers, 2008; Núñez Castellar et al., 2010; Godlove et al., 2011b; Reinhart et al., 2012; Fu et al., 2019), recent meta-analyses and large-scale studies have demonstrated a more robust relationship between single-trial ERN amplitude and post-error slowing. In a meta-analysis of mid-frontal θ , Cavanagh and Shackman (2015) showed systematic evidence that increases in the magnitude of the ERN were coupled with greater post-error slowing; the magnitude of this effect was over twice as large within-subjects

compared with between-subjects. Supporting this, a recent large-scale study of ~ 900 participants demonstrated that greater ERN amplitude was associated with greater post-error slowing within subjects, but not between (Fischer et al., 2016). In contrast to previous findings, we found incidence of β -bursts after errors did not predict post-error adjustments in RT. Previous observations have reported increased β activity in pre-SMA was associated with greater post-error slowing (Marco-Pallarés et al., 2008; Jha et al., 2015) but not with the latency of corrective responses (Marco-Pallarés et al., 2008).

Although we have observed these significant changes in β -burst activity during response inhibition and executive control, β -bursts were neither necessary or sufficient for stopping and were also commonly observed on trials with erroneous saccades executed. Furthermore, changes in stopping performance were not reflected through changes in the proportion of β -bursts observed. The overall low prevalence of β -bursts must be noted. This resulted in β -burst density functions which were small and noisy compared with single neuron spike density functions. Given that β -bursts are relatively uncommon and non-specific, we draw inferences only cautiously. Several alternative hypotheses can be considered. First, medial frontal β -bursts may not be specific to stopping. Indeed, we observed pronounced changes in β -burst occurrence after errors. Second, as argued previously (Jana et al., 2020; Wessel, 2020), EEG has a poor signal-to-noise ratio. This issue can be addressed in future work by measuring the occurrence of β -bursts in intracranial signals recorded in medial frontal areas. Finally, methods for detecting β -bursts could miss significant but subthreshold bursts. Perhaps, more sensitive methods can be developed.

To conclude, by replicating measurements of β -bursts in EEG of macaque monkeys, we establish an animal model of this phenomenon. First, we demonstrated β -bursts were more prevalent when movements were successfully inhibited. Second, we demonstrated a greater incidence of β -bursts over the medial frontal cortex when a movement was erroneously executed. However, in neither context were β -bursts frequent enough to account for behavior. Given the pulses of β -bursts at different stages during the task, it is uncertain whether they index different mechanisms or are produced by one mechanism at different times. In previous work, this uncertainty and apparent absence of causal efficacy has been explained as a consequence of poor signal-to-noise ratio of noninvasive EEG recordings. Establishing the same phenomena in a monkey model engenders confidence in proceeding with the systematic investigation of the neural mechanisms of β -burst generation in the cerebral cortex during countermanding or other tasks. Future studies which sample β -bursts across the layers of the cortex can provide insights into the mechanisms of their generation and possibly help elucidate the role of β -bursts during response inhibition and executive control.

References

- Amarante LM, Caetano MS, Laubach M (2017) Medial frontal theta is entrained to rewarded actions. *J Neurosci* 37:10757–10769.
- Amiez C, Petrides M (2009) Anatomical organization of the eye fields in the human and non-human primate frontal cortex. *Prog Neurobiol* 89:220–230.
- Aron AR, Poldrack RA (2006) Cortical and subcortical contributions to Stop signal response inhibition: role of the subthalamic nucleus. *J Neurosci* 26:2424–2433.
- Aron AR, Behrens TE, Smith S, Frank MJ, Poldrack RA (2007) Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *J Neurosci* 27:3743–3752.
- Beyer F, Münte TF, Fischer J, Krämer UM (2012) Neural aftereffects of errors in a stop-signal task. *Neuropsychologia* 50:3304–3312.
- Bompas A, Campbell AE, Sumner P (2020) Cognitive control and automatic interference in mind and brain: a unified model of saccadic inhibition and countermanding. *Psychol Rev* 127:524–561.
- Boucher L, Palmeri TJ, Logan GD, Schall JD (2007) Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychol Rev* 114:376–397.
- Brockett AT, Hricz NW, Tennyson SS, Bryden DW, Roesch MR (2020) Neural signals in red nucleus during reactive and proactive adjustments in behavior. *J Neurosci* 40:4715–4726.
- Brown JW, Hanes DP, Schall JD, Stuphorn V (2008) Relation of frontal eye field activity to saccade initiation during a countermanding task. *Exp Brain Res* 190:135–151.
- Cabel DW, Armstrong IT, Reingold E, Munoz DP (2000) Control of saccade initiation in a countermanding task using visual and auditory stop signals. *Exp Brain Res* 133:431–441.
- Camalier CR, Gotler A, Murthy A, Thompson KG, Logan GD, Palmeri TJ, Schall JD (2007) Dynamics of saccade target selection: race model analysis of double step and search step saccade production in human and macaque. *Vision Res* 47:2187–2211.
- Cavanagh JF, Shackman AJ (2015) Frontal midline theta reflects anxiety and cognitive control: meta-analytic evidence. *J Physiol Paris* 109:3–15.
- Cavanagh JF, Cohen MX, Allen JJ (2009) Prelude to and resolution of an error: EEG phase synchrony reveals cognitive control dynamics during action monitoring. *J Neurosci* 29:98–105.
- Cohen MX, Ridderinkhof KR, Haupt S, Elger CE, Fell J (2008) Medial frontal cortex and response conflict: evidence from human intracranial EEG and medial frontal cortex lesion. *Brain Res* 1238:127–142.
- Cole MW, Yeung N, Freiwald WA, Botvinick M (2009) Cingulate cortex: diverging data from humans and monkeys. *Trends Neurosci* 32:566–574.
- Colonius H, Özyurt J, Arndt PA (2001) Countermanding saccades with auditory stop signals: testing the race model. *Vision Res* 41:1951–1968.
- De Jong R, Coles MG, Logan GD (1995) Strategies and mechanisms in non-selective and selective inhibitory motor control. *J Exp Psychol Hum Percept Perform* 21:498–511.
- Emeric EE, Brown JW, Boucher L, Carpenter RH, Hanes DP, Harris R, Logan GD, Mashru RN, Paré M, Pouget P, Stuphorn V, Taylor TL, Schall JD (2007) Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vision Res* 47:35–49.
- Emeric EE, Leslie M, Pouget P, Schall JD (2010) Performance monitoring local field potentials in the medial frontal cortex of primates: supplementary eye field. *J Neurophysiol* 104:1523–1537.
- Fischer AG, Danielmeier C, Villringer A, Klein TA, Ullsperger M (2016) Gender influences on brain responses to errors and post-error adjustments. *Sci Rep* 6:24435.
- Floden D, Stuss DT (2006) Inhibitory control is slowed in patients with right superior medial frontal damage. *J Cogn Neurosci* 18:1843–1849.
- Fu Z, Wu DJ, Ross I, Chung JM, Mamelak AN, Adolphs R, Rutishauser U (2019) Single-neuron correlates of error monitoring and post-error adjustments in human medial frontal cortex. *Neuron* 101:165–177.e5.
- Gehring WJ, Fencsik DE (2001) Functions of the medial frontal cortex in the processing of conflict and errors. *J Neurosci* 21:9430–9437.
- Gehring WJ, Goss B, Coles MG, Meyer DE, Donchin E (1993) A neural system for error detection and compensation. *Psychol Sci* 4:385–390.
- Godlove DC, Schall JD (2016) Microsaccade production during saccade cancellation in a stop-signal task. *Vision Res* 118:5–16.
- Godlove DC, Garr AK, Woodman GF, Schall JD (2011a) Measurement of the extraocular spike potential during saccade countermanding. *J Neurophysiol* 106:104–114.
- Godlove DC, Emeric EE, Segovis CM, Young MS, Schall JD, Woodman GF (2011b) Event-related potentials elicited by errors during the stop-signal task. I. Macaque monkeys. *J Neurosci* 31:15640–15649.

- Godlove DC, Maier A, Woodman GF, Schall JD (2014) Microcircuitry of agranular frontal cortex: testing the generality of the canonical cortical microcircuit. *J Neurosci* 34:5355–5369.
- Hajcak G, McDonald N, Simons RF (2003) To err is autonomic: error-related brain potentials, ANS activity, and post-error compensatory behavior. *Psychophysiology* 40:895–903.
- Hanes DP, Schall JD (1995) Countermanding saccades in macaque. *Vis Neurosci* 12:929–937.
- Hanes DP, Carpenter RH (1999) Countermanding saccades in humans. *Vision Res* 39:2777–2791.
- Hanes DP, Patterson WF 2nd, Schall JD (1998) Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J Neurophysiol* 79:817–834.
- Holroyd CB, Yeung N, Coles MG, Cohen JD (2005) A mechanism for error detection in speeded response time tasks. *J Exp Psychol Gen* 134:163–191.
- Husain M, Parton A, Hodgson TL, Mort D, Rees G (2003) Self-control during response conflict by human supplementary eye field. *Nat Neurosci* 6:117–118.
- Jana S, Hannah R, Muralidharan V, Aron AR (2020) Temporal cascade of frontal, motor and muscle processes underlying human action-stopping. *Elife* 9.
- Jha A, Nachev P, Barnes G, Husain M, Brown P, Litvak V (2015) The frontal control of stopping. *Cereb Cortex* 25:4392–4406.
- Kawagoe R, Takikawa Y, Hikosaka O (1998) Expectation of reward modulates cognitive signals in the basal ganglia. *Nat Neurosci* 1:411–416.
- Kerns JG, Cohen JD, MacDonald AW 3rd, Cho RY, Stenger VA, Carter CS (2004) Anterior cingulate conflict monitoring and adjustments in control. *Science* 303:1023–1026.
- Kok A, Ramautar JR, De Ruiter MB, Band GP, Ridderinkhof KR (2004) ERP components associated with successful and unsuccessful stopping in a stop-signal task. *Psychophysiology* 41:9–20.
- Kornylow K, Dill N, Saenz M, Krauzlis RJ (2003) Cancelling of pursuit and saccadic eye movements in humans and monkeys. *J Neurophysiol* 89:2984–2999.
- Ladouceur CD, Dahl RE, Carter CS (2007) Development of action monitoring through adolescence into adulthood: ERP and source localization. *Dev Sci* 10:874–891.
- Logan GD, Cowan WB (1984) On the ability to inhibit thought and action – a theory of an act of control. *Psychol Rev* 91:295–327.
- Logan GD, Yamaguchi M, Schall JD, Palmeri TJ (2015) Inhibitory control in mind and brain 2.0: blocked-input models of saccadic countermanding. *Psychol Rev* 122:115–147.
- Luu P, Tucker DM, Makeig S (2004) Frontal midline theta and the error-related negativity: neurophysiological mechanisms of action regulation. *Clin Neurophysiol* 115:1821–1835.
- Marco-Pallarés J, Camara E, Münte TF, Rodríguez-Fornells A (2008) Neural mechanisms underlying adaptive actions after slips. *J Cogn Neurosci* 20:1595–1610.
- Matzke D, Love J, Heathcote A (2017) A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behav Res Methods* 49:267–281.
- Matzke D, Dolan CV, Logan GD, Brown SD, Wagenmakers EJ (2013a) Bayesian parametric estimation of stop-signal reaction time distributions. *J Exp Psychol Gen* 142:1047–1073.
- Matzke D, Love J, Wiecki TV, Brown SD, Logan GD, Wagenmakers EJ (2013b) Release the BEESTS: Bayesian estimation of ex-Gaussian STop-signal reaction time distributions. *Front Psychol* 4:918.
- Middlebrooks PG, Zandbelt BB, Logan GD, Palmeri TJ, Schall JD (2020) Countermanding perceptual decision-making. *iScience* 23:100777.
- Mirabella G, Pani P, Ferraina S (2011) Neural correlates of cognitive control of reaching movements in the dorsal premotor cortex of rhesus monkeys. *J Neurophysiol* 106:1454–1466.
- Morein-Zamir S, Kingstone A (2006) Fixation offset and stop signal intensity effects on saccadic countermanding: a crossmodal investigation. *Exp Brain Res* 175:453–462.
- Murthy A, Ray S, Shorter SM, Schall JD, Thompson KG (2009) Neural control of visual search by frontal eye field: effects of unexpected target displacement on visual selection and saccade preparation. *J Neurophysiol* 101:2485–2506.
- Nachev P, Wydell H, O'Neill K, Husain M, Kennard C (2007) The role of the pre-supplementary motor area in the control of action. *Neuroimage* 36:T155–T163.
- Nelson MJ, Boucher L, Logan GD, Palmeri TJ, Schall JD (2010) Nonindependent and nonstationary response times in stopping and stepping saccade tasks. *Atten Percept Psychophys* 72:1913–1929.
- Nigbur R, Ivanova G, Stürmer B (2011) Theta power as a marker for cognitive interference. *Clin Neurophysiol* 122:2185–2194.
- Núñez Castellar E, Kühn S, Fias W, Notebaert W (2010) Outcome expectancy and not accuracy determines posterror slowing: ERP support. *Cogn Affect Behav Neurosci* 10:270–278.
- Paré M, Hanes DP (2003) Controlled movement processing: superior colliculus activity associated with countermanded saccades. *J Neurosci* 23:6480–6489.
- Rae CL, Hughes LE, Anderson MC, Rowe JB (2015) The prefrontal cortex achieves inhibitory control by facilitating subcortical motor pathway connectivity. *J Neurosci* 35:786–794.
- Reinhart RM, Carlisle NB, Kang MS, Woodman GF (2012) Event-related potentials elicited by errors during the stop-signal task. II: Human effector-specific error responses. *J Neurophysiol* 107:2794–2807.
- Rodríguez-Fornells A, Kurzbuch AR, Münte TF (2002) Time course of error detection and correction in humans: neurophysiological evidence. *J Neurosci* 22:9990–9996.
- Sajad A, Godlove DC, Schall JD (2019) Cortical microcircuitry of performance monitoring. *Nat Neurosci* 22:265–274.
- Scangos KW, Stuphorn V (2010) Medial frontal cortex motivates but does not control movement initiation in the countermanding task. *J Neurosci* 30:1968–1982.
- Schall JD, Emeric EE (2010) Conflict in cingulate cortex function between humans and macaque monkeys: more apparent than real. *Brain Behav Evol* 75:237–238.
- Schmidt R, Leventhal DK, Mallet N, Chen F, Berke JD (2013) Canceling actions involves a race between basal ganglia pathways. *Nat Neurosci* 16:1118–1124.
- Shin H, Law R, Tsutsui S, Moore CI, Jones SR (2017) The rate of transient beta frequency events predicts behavior across tasks and species. *Elife* 6.
- Skippen P, Matzke D, Heathcote A, Fulham WR, Michie P, Karayanidis F (2019) Reliability of triggering inhibitory process is a better predictor of impulsivity than SSRT. *Acta Psychol (Amst)* 192:104–117.
- Skippen P, Fulham WR, Michie PT, Matzke D, Heathcote A, Karayanidis F (2020) Reconsidering electrophysiological markers of response inhibition in light of trigger failures in the stop-signal task. *Psychophysiology* e13619.
- Stahl J, Gibbons H (2007) Dynamics of response-conflict monitoring and individual differences in response control and behavioral control: an electrophysiological investigation using a stop-signal task. *Clin Neurophysiol* 118:581–596.
- Stoll FM, Wilson CRE, Faraut MCM, Vezoli J, Knoblauch K, Procyk E (2016) The effects of cognitive control and time on frontal beta oscillations. *Cereb Cortex* 26:1715–1732.
- Stuphorn V, Taylor TL, Schall JD (2000) Performance monitoring by the supplementary eye field. *Nature* 408:857–860.
- Stuphorn V, Brown JW, Schall JD (2010) Role of supplementary eye field in saccade initiation: executive, not direct, control. *J Neurophysiol* 103:801–816.
- Sumner P, Nachev P, Morris P, Peters AM, Jackson SR, Kennard C, Husain M (2007) Human medial frontal cortex mediates unconscious inhibition of voluntary action. *Neuron* 54:697–711.
- Swann NC, Cai W, Conner CR, Pieters TA, Claffey MP, George JS, Aron AR, Tandon N (2012) Roles for the pre-supplementary motor area and the right inferior frontal gyrus in stopping action: electrophysiological responses and functional and structural connectivity. *Neuroimage* 59:2860–2870.
- Thakkar KN, Schall JD, Boucher L, Logan GD, Park S (2011) Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biol Psychiatry* 69:55–62.
- Thakkar KN, van den Heiligenberg FM, Kahn RS, Neggers SF (2014) Frontal-subcortical circuits involved in reactive control and monitoring of gaze. *J Neurosci* 34:8918–8929.

- Thakkar KN, Schall JD, Logan GD, Park S (2015) Cognitive control of gaze in bipolar disorder and schizophrenia. *Psychiatry Res* 225:254–262.
- Trujillo LT, Allen JJ (2007) Theta EEG dynamics of the error-related negativity. *Clin Neurophysiol* 118:645–668.
- Verbruggen F, Logan GD (2008) Response inhibition in the stop-signal paradigm. *Trends Cogn Sci* 12:418–424.
- Verbruggen F, Aron AR, Band GP, Beste C, Bissett PG, Brockett AT, Brown JW, Chamberlain SR, Chambers CD, Colonius H, Colzato LS, Corneil BD, Coxon JP, Dupuis A, Eagle DM, Garavan H, Greenhouse I, Heathcote A, Huster RJ, Jahfari S, et al. (2019) A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *Elife* 8:e46323.
- Walton MM, Gandhi NJ (2006) Behavioral evaluation of movement cancellation. *J Neurophysiol* 96:2011–2024.
- Wattiez N, Poitou T, Rivaud-Péchoux S, Pouget P (2016) Evidence for spatial tuning of movement inhibition. *Exp Brain Res* 234:1957–1966.
- Wessel JR (2020) β -Bursts reveal the trial-to-trial dynamics of movement initiation and cancellation. *J Neurosci* 40:411–423.
- Wessel JR, Aron AR (2015) It's not too late: the onset of the frontocentral P3 indexes successful response inhibition in the stop-signal paradigm. *Psychophysiol* 52:472–480.
- West R, Travers S (2008) Tracking the temporal dynamics of updating cognitive control: an examination of error processing. *Cereb Cortex* 18:1112–1124.
- Woodman GF (2012) Homologues of human ERP components in nonhuman primates. In: *Oxford handbook of event-related potential components*, Ed 1, pp 611–626. New York: Oxford University Press.

Appendix C

Sajad, A.*, Errington, S.P.* & Schall, J.D (2022). Functional architecture of executive control and associated event-related potentials in macaques. *Nature Communications*, 13 (6270).


Published October 21st 2022

Functional architecture of executive control and associated event-related potentials in macaques

Received: 27 April 2021

Accepted: 7 October 2022

Published online: 21 October 2022

 Check for updatesAmirsaman Sajad^{1,3}, Steven P. Errington ^{1,3} & Jeffrey D. Schall ^{1,2} 

The medial frontal cortex (MFC) enables executive control by monitoring relevant information and using it to adapt behavior. In macaques performing a saccade countermanding (stop-signal) task, we simultaneously recorded electrical potentials over MFC and neural spiking across all layers of the supplementary eye field (SEF). We report the laminar organization of neurons enabling executive control by monitoring the conflict between incompatible responses, the timing of events, and sustaining goal maintenance. These neurons were a mix of narrow-spiking and broad-spiking found in all layers, but those predicting the duration of control and sustaining the task goal until the release of operant control were more commonly narrow-spiking neurons confined to layers 2 and 3 (L2/3). We complement these results with evidence for a monkey homolog of the N2/P3 event-related potential (ERP) complex associated with response inhibition. N2 polarization varied with error-likelihood and P3 polarization varied with the duration of expected control. The amplitude of the N2 and P3 were predicted by the spike rate of different classes of neurons located in L2/3 but not L5/6. These findings reveal features of the cortical microcircuitry supporting executive control and producing associated ERPs.

Effective control of behavior is necessary to override conflicting, habitual, or inappropriate responses, and to facilitate stopping, switching, and updating of task goals. Investigating features of executive control is afforded through the countermanding (stop-signal) task¹, during which macaque monkeys, like humans, exert response inhibition and adapt performance based on stimulus history, response outcomes, and the temporal structure of task events².

Converging evidence from imaging, electrophysiology, and lesion studies indicates that MFC, including the supplementary motor complex, is essential for executive control^{3–6}. In humans, noninvasive ERP measures derived from a negative–positive waveform over the medial frontal cortex, known as the N2/P3, have been used to test hypotheses about executive control function⁷. Executive control of gaze is mediated by the SEF—an agranular area on the dorsomedial convexity of the MFC. Here, neurons modulate too late to enable reactive control of eye

movements and instead contribute to proactive control⁸. Electrical microstimulation of SEF improves performance in the countermanding task by slowing response time (RT)⁹ by postponing the accumulation of pre-saccadic activity in gaze control structures¹⁰. SEF also supports working memory^{11,12} and signals surprising events¹³, event timing^{14,15}, response conflict¹⁶, and consequences¹⁷.

Whilst these executive control signals are well documented in MFC, how they arise is uncertain^{3–6}. Understanding neural spiking with laminar resolution is necessary to clarify circuit-level mechanisms because neurons in different layers have different biophysical properties and anatomical connections. Although such approaches have been integral in developing our understanding of processing in the early visual system¹⁸ the canonical cortical microcircuit derived from granular sensory areas¹⁹ does not explain agranular frontal areas like SEF^{20–24}.

¹Department of Psychology, Vanderbilt Vision Research Center, Center for Integrative & Cognitive Neuroscience, Vanderbilt University, Nashville, TN, USA.

²Department of Biology, Centre for Vision Research, Vision Science to Application, York University, Toronto, ON, Canada. ³These authors contributed equally: Amirsaman Sajad, Steven P. Errington. ✉ e-mail: schalljd@yorku.ca

To understand mechanistically how executive control signals are generated in the MFC, we sampled neural spiking activity in SEF beneath where the frontal N2/P3 is sampled. We previously described the laminar microcircuitry of performance monitoring signals in the SEF and their relationship to error-related negativity (ERN)¹⁷. Here we describe the laminar microcircuitry of signals that monitor events occurring during successful stopping performance. We describe three classes of neurons that signal response conflict, event timing, and maintenance of task goals. We also provide evidence that macaque monkeys produce the N2/P3 ERP associated with response inhibition, describe the task factors indexed by this ERP, and elucidate the neuron classes predicting the polarization.

Results

Countermanding performance and neural sampling

Neurophysiological and electrophysiological data were recorded from two macaque monkeys performing the saccade stop-signal task with

explicit feedback tone cues preceding the possible delivery of fluid reward (Fig. 1a)²⁵. Data collection and analysis were informed by the consensus guide for the stop-signal task¹. Briefly, monkeys earned fluid reward for shifting their gaze to a target on its appearance and for inhibiting the planned saccade when an infrequent stop-signal appeared. The delay of the stop-signal was varied experimentally to yield an equal probability of successful (canceled) or unsuccessful (noncanceled, NC) stop-signal trials. Two kinds of NC trials were distinguished (Fig. 1a): NC trials in which the gaze shifted after the stop-signal appeared (RT > SSD) were identified as errors (NC_{error}), whereas trials when gaze shifted before SSD (RT < SSD) were identified as premature noncanceled trials (NC_{premature}). Both trials resulted in no juice reward being delivered.

We acquired 33,816 trials across 29 sessions (Monkey Eu, male, 12 sessions, 11,583 trials; Monkey X, female, 17 sessions, 22,233 trials). Typical performance was produced by both monkeys. First, response times (RT) on noncanceled trials (mean ± SD Eu: 294 ± 179 ms; X:

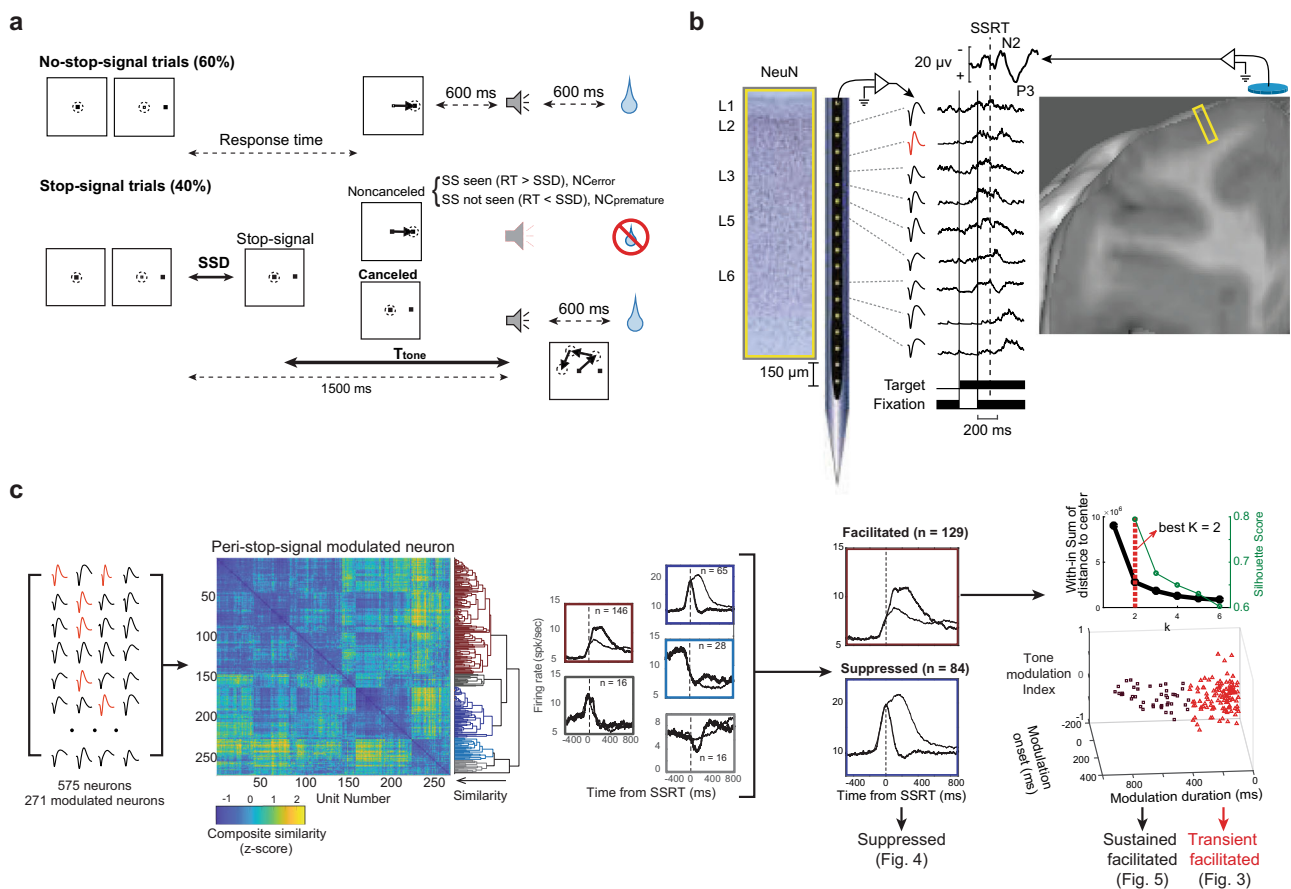


Fig. 1 | Experimental approach. **a** Saccade countermanding task. Monkeys earned fluid reward for shifting gaze (dashed circle) to a visual target unless a stop-signal appeared after a variable stop-signal delay (SSD) adjusted to achieve ~50% canceled trials. Successful no-stop-signal or canceled trial outcome was signaled by a high-pitched tone after T_{tone} preceding fluid delivery. Noncanceled errors were signaled by a low-pitched tone. Monkeys could shift gaze and blink after T_{tone} , which equaled 1500 ms - SSD. Noncanceled trials with RT > SSD were explicit errors. Noncanceled trials with RT < SSD were premature responses. Details in text.

b Neural sampling. Neural spiking was recorded across all layers of agranular SEF (NeuN stain) using Plexon U-probe. Neurons with both broad (black) and narrow (red) spikes were sampled. Spiking modulation was measured relative to the presentation of task events (thin solid, visual target; thick solid, stop-signal) and performance measures like SSRT (dashed vertical). Simultaneously, EEG was recorded from an electrode in the cranial surface over the MFC (10–20 location Fz). The yellow rectangle portrays the cortical area sampled in a T1 MR image. **c** Neuron

classification. Among 575 sampled neurons 271 were classified by their modulation during successful stopping. Their respective spike density functions (SDF) were submitted to an unsupervised consensus clustering pipeline²⁷ yielding the composite similarity matrix (z-score color map) and associated dendrogram. The color coding in this figure bears no relation to colors used in other figures. Consensus clustering yielded 5 clusters with 2 clusters containing >80% of the neurons. The different modulation patterns are evident in the average SDF aligned on SSRT of each cluster. Following manual curation, 129 facilitated and 84 suppressed neurons were analyzed further based on heterogeneity in modulation latency, duration, and pattern with the feedback tone. K-means clustering further divided facilitated neurons into sustained and transient classes (right panel). The inset in the right panel shows that $k = 2$ clusters were chosen based on the Elbow method and the Silhouette score. Therefore, three classes of neurons were analyzed in this study.

230 ± 83 ms) were systematically shorter than those on no-stop-signal trials (Eu: 313 ± 119 ms, X: 263 ± 112 ms; Mixed-effects linear regression (two-tailed) grouped by monkey, $t(27507) = -17.4$, $p < 0.001$). Second, the probability of noncanceled trials increased with stop-signal delay (SSD). These two observations validated the use of the independent race model²⁶ to estimate the stop-signal reaction time (SSRT), an estimate of the time needed to cancel a partially prepared saccade. SSRT measures the time of successful stopping. If the stop process does not inhibit saccade preparation before this time, then gaze will shift. Accordingly, neural modulation occurring before SSRT can contribute to reactive stopping but modulation occurring after SSRT cannot^{8,27}. SSRT across sessions did not differ between monkeys (Eu: 118 ± 23 ms, X: 103 ± 24 ms, independent groups t -test (two-tailed), $t(27) = -1.69$, $p = 0.103$).

Electroencephalogram (EEG) was recorded over the MFC with leads placed on the cranial surface beside the chamber while a linear electrode array (Plexon, 24 channels, 150 μm spacing) was inserted in SEF (Fig. 1b). SEF was localized by anatomical landmarks and intracortical electrical microstimulation²¹. Neural spiking was sampled from 575 single units (Eu: 244, X: 331) across five sites. Overall, 213 neurons (Eu: 105, X: 108) modulated around SSRT and revealed the functional signals reported in this study. The description of the laminar distribution of signals is based on 16 of the 29 sessions during which electrode arrays were oriented perpendicular to the cortical layers. In total, 119 neurons (Eu: 54; X: 65) were assigned to layers confidently (Supplementary Fig. 1b, Supplementary Table 1). Due to variability in the estimates and the indistinct border between L6 and white matter, units appearing beneath the average gray-matter estimate were assigned to L6.

Functional classification of neural activity related to successful stopping

Aspects of the behavioral and neural dataset analyzed here have been used to address other questions. The current results describe the activity on successfully canceled stop-signal trials. According to the race model²⁶ and validated neurophysiologically²⁷ noncanceled trials result from faster responses than canceled trials ($RT < SSD + SSRT$, Supplementary Fig. 2a) and therefore are associated with distinct neural mechanisms. Activity on noncanceled stop-signal trials was described previously¹⁷.

We classified neurons primarily based on the pattern of neural activation around SSRT (Fig. 1c, Supplementary Fig. 1c). A semi-supervised consensus cluster algorithm²⁸ followed by manual curation revealed populations of neurons that were facilitated (129) or suppressed (84) following successful response inhibition. Facilitated neurons were further classified into two distinct subpopulations based on other response profile properties (Fig. 1c).

To identify the putative functional roles of each pattern of modulation, we contrasted neural modulation on canceled trials with a subset of no-stop-signal trials with matching temporal dynamics, identified through a process of latency-matching. This process only included no-stop-signal trials with response times long enough to have been canceled had the stop-signal been presented ($RT \geq SSD + SSRT$; Supplementary Fig. 2a). We distinguished canceled trials by SSD to account for variations in modulation dynamics arising from differences in the timing of task events (Fig. 1a). We also examined the neural activity before the feedback tone which terminated operant control on behavior after successful stopping (Fig. 1a). Variations in neural activity across SSDs were tested against a variety of parameters derived from different theories of MFC function described below (Fig. 2b; Supplementary Table 2).

Performance of the stop-signal task is explained as the outcome of a race between stochastic GO and STOP processes²⁶, instantiated by specific interactions enabling the interruption of the GO process by a STOP process^{29,30} (Fig. 2a). A theory of medial frontal function posits

that it encodes the conflict between mutually incompatible processes³¹. Such conflict arises naturally as the mathematical product of the activation of the mutually incompatible GO and STOP units. The probability of noncanceled trials at each SSD, $p(\text{NC}|\text{SSD})$, served as a proxy for conflict because it is the outcome of the race between the conflicting processes. This was validated by simulations of the GO and STOP units in the interactive race model (Fig. 2a). Therefore, the conflict model predicts variations in neural activity as a function of $p(\text{NC}|\text{SSD})$.

Inspired by reinforcement learning models, we determined whether neural modulation after stop-signal appearance varied with the likelihood of error associated with an experienced SSD³². Note that a stop-signal appearing after RT cannot contribute to this association. The experienced SSD can be learned only in trials when the stop-signal was seen (SS_{seen}), which are canceled trials and noncanceled trials that are explicit errors (NC_{error}). Therefore, the error-likelihood model predicts that neural activity varies with the likelihood of error, which was operationalized by $p(\text{NC}_{\text{error}}|\text{SSD})/p(SS_{\text{seen}}|\text{SSD})$. Neural modulation scaling positively (negatively) with this quantity encodes the error (success) likelihood associated with each SSD. This quantity diverges from $p(\text{NC}|\text{SSD})$ at longer SSDs (Fig. 2b, top-left panel).

Finally, due to the temporal regularities of SSD and other events, we determined whether neural modulation signaled interval timing^{14,15,33,34}, moment-by-moment expectation for events^{35,36}, or the surprise associated with the violation of the expectations^{37,38} (Fig. 2b, Supplementary Table 2, Supplementary Fig. 3). Guided by previous research on time perception^{33,35,39}, we tested whether neural modulation varied with elapsed time in linear or logarithmic scales. Expectation was operationalized by hazard rate

$$h(t) = \frac{f(t)}{[1 - F(t)]} \quad (1)$$

where $f(t)$ is the probability density and $F(t)$ is the associated cumulative distribution. Surprise at the violation of this expectation was operationalized by Shannon's information,

$$s(t) = -\log_2[h(t)] \quad (2)$$

Different hazard rates and surprise quantities were computed based on different representations of temporal statistics plus the perceptual precision of SSD and the time until the feedback tone (T_{tone} ; Fig. 2b; Supplementary Fig. 3a). Hazard rate models predict that neural activity before an event varies with $h(t)$, and surprise models predict that neural activity after an event varies with $s(t)$. The various models resulting from these quantities were compared through mixed-effects model comparison with Bayesian Information Criteria (BIC).

Conflict monitoring neurons. The classification pipeline identified 75 neurons with transient facilitation on canceled trials compared to latency-matched no-stop-signal trials (Fig. 1c). This facilitation was not a visual response to the stop-signal because it did not occur on noncanceled trials (Supplementary Figs. 1e and 2b). It also cannot contribute to reactive response inhibition because for nearly all neurons (71/75) it arose after SSRT (Fig. 3a, Supplementary Fig. 4a). The facilitation started 90.4 ± 74 ms after SSRT with peak recruitment at -110 ms at which time $\sim 60\%$ of neurons were active (Fig. 3a).

Through model comparison, we assessed how the magnitude of this modulation varied with task and performance parameters. The facilitation was best described by the conflict model with higher activity associated with larger $p(\text{NC}|\text{SSD})$ (Fig. 3e; Supplementary Table 3, Mixed-effects linear regression (two-tailed) grouped by neuron, $t(212) = 4.24$, $p < 0.001$). Time-based and surprise models derived from the temporal statistics on canceled trials were candidates ($\Delta\text{BIC} < 2$). The error-likelihood ($\Delta\text{BIC} = 3.52$) and other surprise

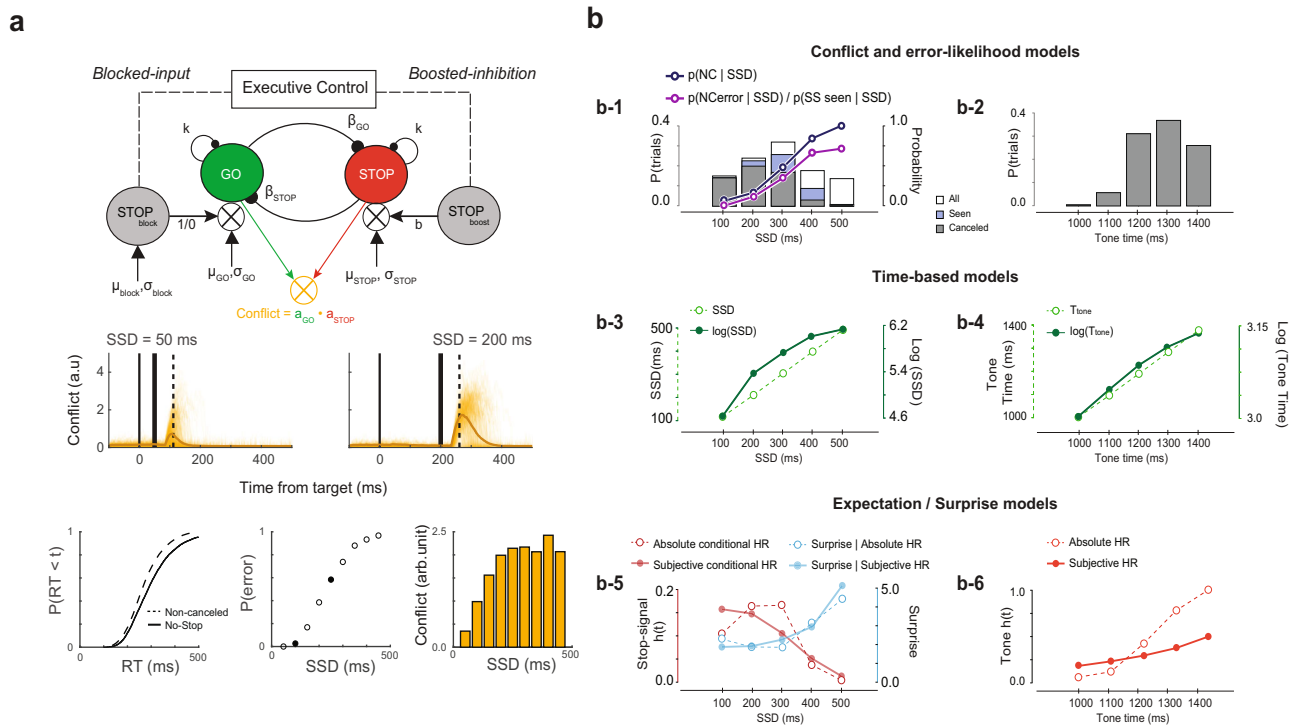


Fig. 2 | Models of neural modulation. **a** Interactive race model architecture. In interactive race architectures^{29,30} GO and STOP units are racing, competing response channels, so the product of their co-activation measures response conflict in this task. In simulations with parameters simulating observed RT and $p(\text{NC}|\text{SSD})$ values (lower panels) the instantaneous product of $a_{GO} \cdot a_{STOP}$ of individual (thin) and average (thick) increases around and peaks after SSRT (middle panels). The product of the activation of GO and STOP units producing countermmanding performance is proportional to $p(\text{NC}|\text{SSD})$ (lower right panel). Consequently, $p(\text{NC}|\text{SSD})$ is an effective proxy for the conflict measure. **b** Quantities related to different behavioral and task parameters that arise from the variability in SSD and T_{tone} are shown for a representative session. These quantities defined models with different predictions about variations in neural activity. Conflict and error-likelihood models (**b-1**) were based on the proportions of stop-signal trials (left ordinate) in which the saccade was successfully canceled (gray bars), those in which the stop-signal was seen because of RT > SSD (SS_{seen} ; blue bars) and all stop-signal trials (white bars).

The probability of NC trials as a function of SSD (right ordinate) indexed conflict with $p(\text{NC}|\text{SSD})$ and error-likelihood with $p(\text{NC}_{\text{error}}|\text{SSD})/p(\text{SS}_{\text{seen}}|\text{SSD})$, which diverge at long SSD. Time-based models were based on the duration of SSD (**b-1**) and T_{tone} (**b-2**) with t (left ordinate) and $\log(t)$ (right ordinate) values (**b-3**, **b-4**). Expectation and surprise models were based on hazard rates of SSD and T_{tone} derived from their respective distributions (**b-1** and **b-2**). Absolute and Subjective, incorporating imprecision in duration estimation, hazard rates (red) of SSD (**b-5**) and tone (**b-6**) and the surprise associated with stop-signal (b-5; blue) are shown. Other variants were also considered based on the underlying assumption about knowledge of task structure and imprecision in duration estimation (Supplementary Table 2, Supplementary Fig. 3a). Expectation of SS was quantified by hazard rate conditional on stop-signal trial probability (~40%). Surprise as a violation of expectations was quantified by the Shannon information derived from the hazard rate.

models earned weak support ($\Delta\text{BIC} = 5.7$) or were rejected ($\Delta\text{BIC} > 6$) (Fig. 3d; Supplementary Table 3). Although some surprise and time-based models explained the modulation, preference for the conflict model aligns with an earlier conjecture that these neurons signal conflict derived from the co-activation of GO and STOP unit activation¹⁶. Because this signal arises too late to influence response inhibition, we conjecture that it contributes to conflict monitoring.

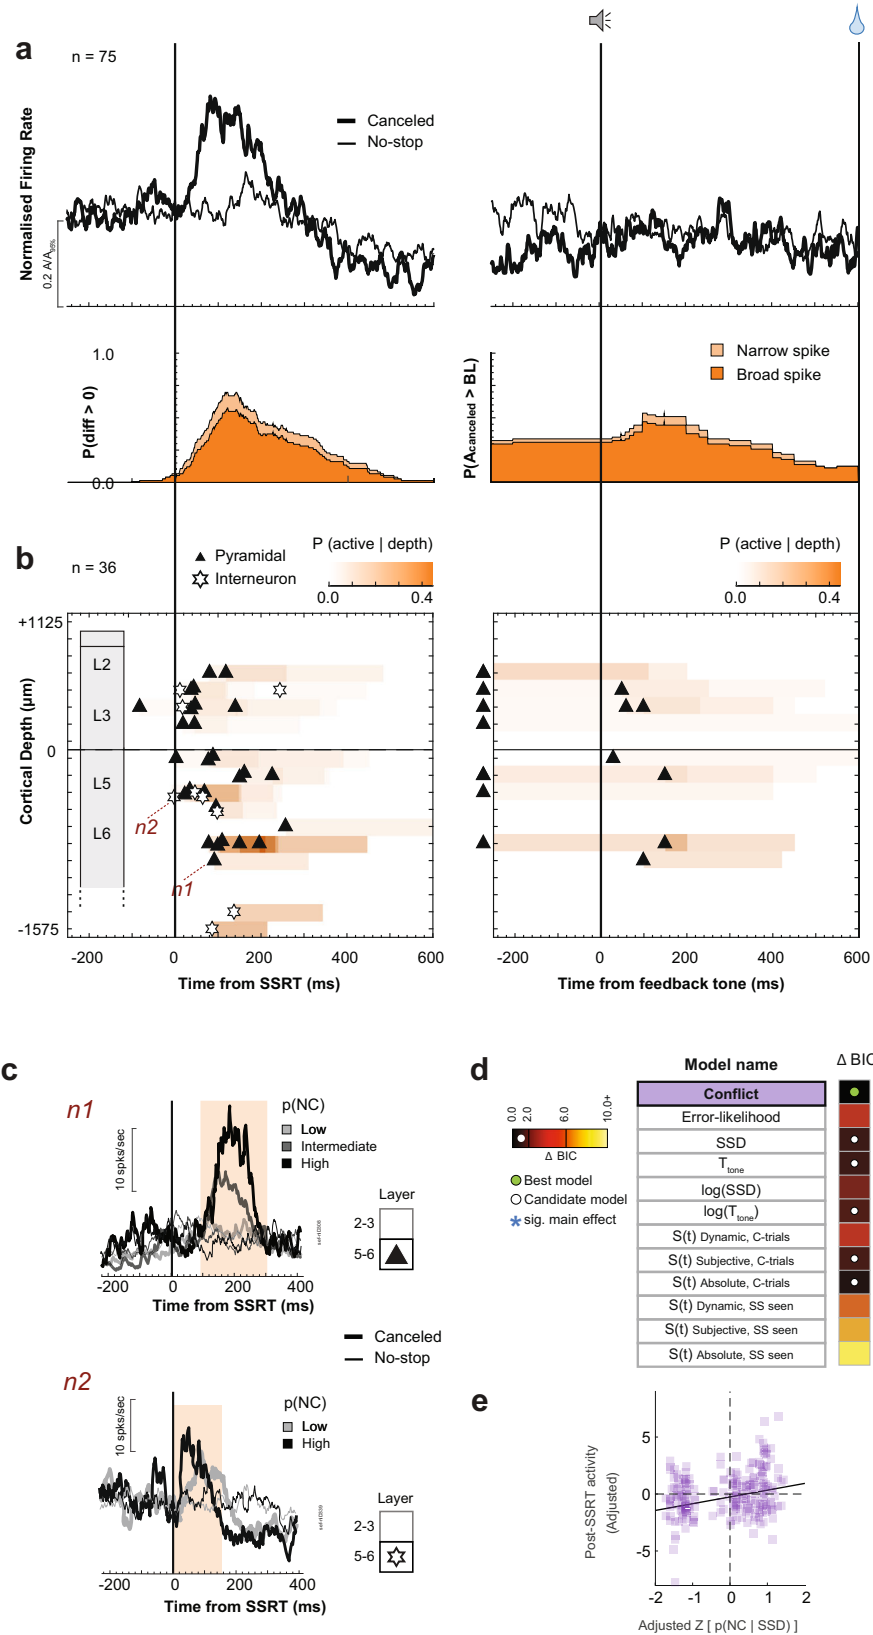
On canceled trials, a minority of these neurons produced persistent weak activity until the tone, some with a transient response thereafter (Fig. 3a, b; Supplementary Fig. 1e). The spike rate immediately before the feedback tone was unrelated to its time or anticipation. Although RT slowed after successful stopping (multiple linear regression (two-tailed) controlling for direction and session, $t(13664) = 15.7, p < 0.001$), the modulation of only a few neurons (7/75) covaried with RT (multiple linear regression controlling for SSD, $p < 0.05$). Therefore, Conflict neurons were not involved in RT adjustments.

Over half of Conflict neurons (57%) exhibited multiplexing with reinforcement- or error-related signals reported previously¹⁷ (Supplementary Table 7). Some produced higher discharge rates on unrewarded trials (previously identified as Loss neurons); some, had higher discharge rates on rewarded trials (Gain neurons). However, multiplexing incidence did not differ from that predicted by the sampling

prevalence of these signals (Chi-square test of homogeneity (one-tailed), $\chi^2(3, N = 575) = 1.02, p = 0.791$). The vast majority (65/75) were not modulated when stopping failed (Supplementary Fig. 2b, Supplementary Table 7), reinforcing previous findings that conflict and error monitoring are distinct^{16,17}.

To investigate the microcircuit contribution of Conflict neurons, we examined their spike waveform duration and distribution across the layers (Supplementary Fig. 1a). Neurons were distinguished as broad- and narrow-spiking, which may identify pyramidal neurons that can send extrinsic connections and intrinsic inhibitory interneurons, respectively⁴⁰. Most Conflict neurons (63/75) had broad spikes, but the incidence did not exceed that observed in the overall recording sample (Chi-square test of homogeneity (one-tailed), $\chi^2(1, N = 575) = 0.205, p = 0.152$).

Conflict neurons were observed in all recording sites but more prevalently at some (Chi-square test of homogeneity (one-tailed), $\chi^2(4, N = 575) = 11.6, p = 0.020$). Those recorded in sessions with perpendicular penetrations (36/75) revealed the spatiotemporal progression of the conflict signal across time and cortical depth (Fig. 3b). These neurons were sampled from all layers with an incidence corresponding to the sampling distribution (Chi-square test of homogeneity (one-tailed), $\chi^2(4, N = 293) = 4.28, p = 0.369$; Fig. 3b; Supplementary Fig. 1a; Supplementary Table 1). The onset time of the modulation did not



differ between L2/3 and L5/6 (Mann–Whitney *U* test (two-tailed), $U(13, 23) = 181$, $z = -1.94$, $p = 0.052$). A larger proportion of concurrently activated Conflict neurons was observed in L5/6 (Fig. 3b). The few neurons modulating with the tone were observed in all layers.

Thus, complementing previous observations¹⁶, particular neurons in SEF modulate in a manner consistent with signaling the co-activation

of gaze-shifting (GO) and gaze-holding (STOP) mechanisms previously interpreted as conflict^{31,41}. These neurons are distributed across all SEF layers and are predominantly broad-spiking, as expected from the sampling distribution. This conflict signal has also been observed in dopamine (DA) neurons in dorsolateral substantia nigra pars compacta (SNpc) during saccade countermanding⁴². The SNpc neurons

Fig. 3 | Conflict neuron time–depth organization. **a** Average spike rate (top) and recruitment (bottom) of broad- (dark) and narrow-spiking (light) neurons on canceled (thick) and latency-matched no-stop-signal (thin) trials, aligned on SSRT (left) and tone (right), normalized to 95th percentile within respective intervals. SSRT-aligned recruitment was the difference between trials indicated by $p(\text{diff} > 0)$. Tone-aligned recruitment was the difference in spiking on canceled trials (A_{canceled}) relative to baseline (BL) lowest spiking value ± 500 ms from the tone. Modulations after tone were not analyzed. Post-saccadic spiking on no-stop-signal trials before tone can exceed that on canceled trials. **b** Time–depth plot showing latency and recruitment across depth from perpendicular penetrations. Symbols mark the beginning of modulation for broad- (triangles) and narrow-spiking (stars) neurons. Color map indicates percentage through time at each depth relative to sampling density. Solid horizontal line marks L3–L5 boundary. The lower boundary of L6 is not discrete. **c** Modulation on canceled (thick) relative to latency-matched no-stop-signal (thin) trials for lower, intermediate, and higher $p(\text{NC|SSD})$ of two

representative neurons: n1 had broad spikes in L6; n2 had narrow spikes in L5 (n2) (identified in **b**). Shaded interval highlights significant differences in spiking across conditions. **d** Model comparison table listing each tested model. The heatmap shows the difference in BIC values (ΔBIC) for each model compared to the model with the lowest BIC value (black fill) with hotter colors corresponding to lower ΔBIC values. Asterisks (*) indicate models with a significant main effect. The green circle indicates the best-fit model; the white indicates candidate models ($\Delta\text{BIC} < 2$). Spike rate variation after SSRT was best predicted by the conflict model. Full statistics in Supplementary Table 3. **e** Significant variation of spiking (residualized and adjusted for spiking across neurons) as a function of $p(\text{NC|SSD})$ proxy for conflict (normalized z -scale). Each point plots average spike density and mean $p(\text{NC|SSD})$ across all trials in early-, mid-, or late-SSD bins for each of 75 neurons. For 11 neurons no reliable estimate of spike density for the late-SSD bin was obtained due to too few trials. These data points were not included.

modulate significantly earlier than those in SEF but accounting for the long conduction delay, a dopaminergic signal cannot cause the SEF modulation (Supplementary Fig. 9; Supplementary Table 9).

Event timing neurons. The classification pipeline identified another group of 84 neurons with ramping activity following target presentation on all trials with an abrupt reduction in discharge rate after SSRT on canceled trials (Fig. 4; Supplementary Fig. 1c–e). This modulation cannot contribute to reactive response inhibition because for nearly all neurons (76/84) it happened after SSRT (Supplementary Figs. 1d and 5a). This suppression occurred 62 ± 58 ms after SSRT, significantly earlier than Conflict neurons (Mann–Whitney U test (two-tailed), $U(84, 75) = 6001$, $z = -2.44$, $p = 0.014$). By -150 ms after SSRT, nearly all neurons had suppressed spiking (Fig. 4a). Noting the similarity of the ramping to earlier descriptions of time-keeping neurons^{4,15}, we focused on the ramping activity preceding the suppression.

Model comparison revealed that the ramping magnitude varied best with $\log(\text{SSD})$ (Mixed-effects linear regression (two-tailed) grouped by neuron, $t(250) = 12.62$, $p = 0.001$) with higher activity for longer SSD durations (Fig. 4e; Supplementary Table 4). Absolute SSD earned weak support ($\Delta\text{BIC} = 2.7$), but hazard rate models were rejected ($\Delta\text{BIC} > 6$). Once successful stopping occurred, these neurons were suppressed. On no-stop-signal and noncanceled saccade trials, the ramping activity occurred in other epochs followed by a decline following the saccade or feedback tone (Supplementary Fig. 2b). Because the discharge rate decreased sharply after SSRT on canceled but not on noncanceled stop-signal trials (Fig. 4d, Supplementary Figs. 1e and 2b), we conjectured that these neurons are sensitive to the timing of events leading to successful stopping and not the timing of the stop-signal appearance per se.

Inherent to the task, following SSD and SSRT on canceled trials monkeys had to maintain fixation on the stop-signal for a variable but predictable duration (T_{tone} , Fig. 1a). Following the post-SSRT suppression, the activity in a subset of these neurons (38/84) exhibited a second ramping period preceding the tone, whereupon the spike rate decreased (Fig. 4; Supplementary Fig. 5b). This ramping had lower slope than that before SSRT. In a few neurons, the decrease in activity after the tone followed a brief transient response (Fig. 4a).

The variation in ramping dynamics before the tone was best explained by time-based models (T_{tone} ; Mixed-effects linear regression (two-tailed) grouped by neuron, $t(112) = 3.41$, $p < 0.001$; Fig. 4h, Supplementary Table 4) with weak support for hazard rate models ($3.0 < \Delta\text{BIC} < 4.4$). The T_{tone} and $\log(T_{\text{tone}})$ models were indistinguishable ($\Delta\text{BIC} < 0.1$) (Fig. 4g). The termination of this ramping activity was synchronized with the feedback tone and not the time at which fixation was interrupted following the feedback (Supplementary Fig. 5c).

Because the ramping activity of these neurons co-varies with SSD and T_{tone} followed by suppression when the interval elapses, we conjecture that these neurons signal event timing. Event Timing neurons

were identified by ramping before SSRT, but $\sim 45\%$ also encoded the timing of the feedback tone. This suggests that the timing signal can exhibit different specificities in different neurons.

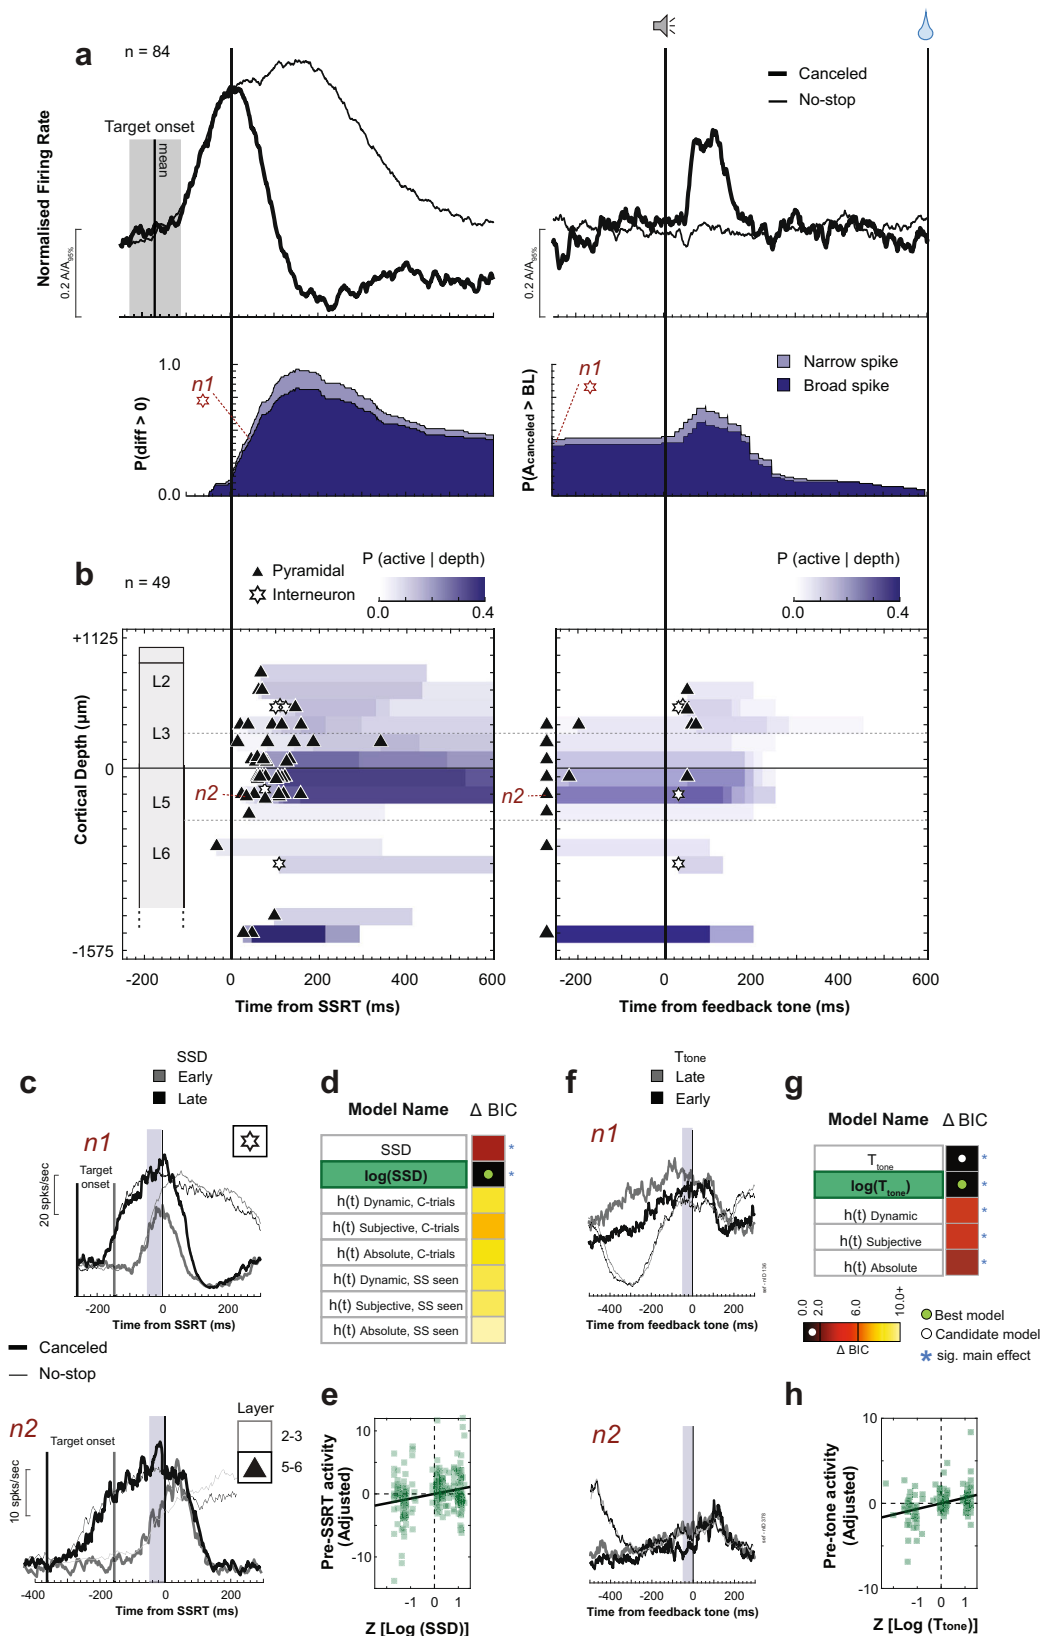
Next, we examined how the activity of these neurons relates to adjustments in behavior. Replicating previous findings, we found that RT was influenced by the SSD experienced in the previous trial, with slower RTs following longer SSD (Linear regression model (two-tailed), $t(83) = 2.64$, $p = 0.010$)⁴³. However, the peak of the peri-SSRT ramping activity of only two neurons predicted RT in the trial following successful stopping (multiple linear regression (two-tailed) for each neuron controlling for SSD, $p < 0.05$). Therefore, this putative time-keeping signal does not influence the slowing of RT after canceled trials.

Neurons identified with event timing multiplexed with performance monitoring signals reported previously¹⁷. Event timing was significantly more likely to be observed in Error and Gain neurons compared to Loss neurons (Chi-square test of homogeneity (one-tailed), $\chi^2(3, N = 575) = 44.86$, $p < 0.001$, Supplementary Table 7). The nature of the ramping before SSRT was invariant across multiplexing associations. Ramping after the feedback tone until reward delivery was observed in Gain neurons (24/84).

To elucidate the microcircuit contribution of Event Timing neurons, we examined their spike waveform duration and distribution across the layers. The majority (72/84) were broad-spiking similar in proportion to the sampling distribution (Chi-square test of homogeneity (one-tailed), $\chi^2(1, N = 575) = 3.75$, $p = 0.053$). These neurons were found in all penetrations but were more common in some sites ($\chi^2(4, N = 575) > 39.3$, $p < 0.001$; Supplementary Table 1). The time–depth organization of these neurons was revealed in the sample from perpendicular penetrations (49/84; Fig. 4b). Those ramping before SSRT were found across all layers in proportion to the sampling distribution (Chi-square test of homogeneity (one-tailed), $\chi^2(4, N = 293) = 7.33$, $p = 0.120$; Fig. 4b; Supplementary Fig. 1a; Supplementary Table 1). However, those with ramping after SSRT until the tone were significantly more concentrated in lower L3 and L5 ($\chi^2(2, N = 293) = 10.37$, $p = 0.006$; Supplementary Table 1). The suppression time after SSRT or the tone did not vary across layers.

Thus, neurons in SEF exhibited ramping activity that can signal the time preceding critical events for successful task performance. These results show that these neurons are distributed across all SEF layers and are predominantly broad-spiking as expected from the sampling distribution. Also, the time-related signals in the SEF can have different functional specificities⁴⁴ and multiplex with error and reinforcement signals in different layers.

Goal maintenance neurons. We identified another class of 54 facilitated neurons with significantly greater discharge rate on canceled compared to latency-matched no-stop-signal trials after SSRT (Fig. 5). This neuron class was distinguished from Conflict neurons based on



the duration of facilitation (*K*-means clustering, Fig. 1c, Supplementary Fig. 1d) and by other differences described below.

The facilitation was not a response to the stop-signal, because it did not occur on noncanceled trials (Supplementary Fig. 2b). It cannot contribute to reactive response inhibition because it arose after SSRT for effectively all neurons (53/54; Fig. 5a; Supplementary

Fig. 6a). On average, the facilitation began 120 ± 85 ms after SSRT not significantly different from that of Conflict neurons (Mann-Whitney *U* test (two-tailed), $U(75, 54) = 4509$, $z = -1.75$, $p = 0.081$) but significantly later than the suppression in Event Timing neurons (Mann-Whitney *U* test (two-tailed), $U(84, 54) = 4940.5$, $z = -3.91$, $p < 0.001$). The peak recruitment of these

Fig. 4 | Event timing neuron time–depth organization. Conventions as in Fig. 3. **a** Average spike rate and suppression time of neurons with ramping spiking after the target (mean: vertical line; min-max: gray-shaded rectangle) followed by suppression on successfully canceled relative to latency-matched no-stop-signal trials. SSRT-aligned recruitment was the difference in spiking between canceled and no-stop-signal trials. **b** Time–depth plots. Horizontal dashed lines highlight where Event Timing neurons with pre-tone ramping were concentrated in lower L3 and L5. **c** Modulation on successfully canceled relative to latency-matched no-stop-signal trials for early (lighter) and late (darker) SSDs of two representative neurons: n1 had narrow spikes but no layer assignment; n2 had broad spikes in L5. Spiking in shaded

50 ms before SSRT (shaded) was used for analysis. **d** Model comparison table for pre-SSRT activity. The best model is highlighted in green. Variation of spiking was best predicted by $\log(\text{SSD})$. Full statistics in Supplementary Table 4. **e** Significant variation of spiking activity before SSRT as a function of $\log(\text{SSD})$ with 84 neurons contributing 252 samples. **f** Modulation of neurons n1 and n2 with ramping before the tone. Shaded 50 ms indicates epoch analyzed. **g** Model comparison table for pre-tone activity. Variation of spiking was best-predicted $\log(T_{\text{tone}})$. Full statistics in Supplementary Table 4. **h** Significant variation of spiking activity before tone as a function of $\log(T_{\text{tone}})$ with 38 neurons contributing 144 samples across early, intermediate, or late T_{tone} .

neurons reached -95% after -300 ms, later than that of the Conflict (-110 ms) and Event Timing neurons (-150 ms) (Fig. 5a).

The variation in activity of these neurons was explained best by the surprise model ($S(t)_{\text{Subjective, SS seen}}$; Mixed-effects linear regression (two-tailed) grouped by neuron, $t(151) = -3.91$, $p < 0.001$; Fig. 5d; Supplementary Table 5) with other surprise and time-based models also candidates ($\Delta\text{BIC} < 2$) and weak support for $\log(\text{SSD})$ ($\Delta\text{BIC} = 3.0$) and conflict ($\Delta\text{BIC} = 4.1$) models. The error-likelihood model was rejected ($\Delta\text{BIC} > 6.0$). The variation of activity varied inversely with surprise and positively with T_{tone} and $\log(T_{\text{tone}})$ (Fig. 5c, e; Supplementary Table 5). Thus, the modulation of these neurons was further distinguished from the conflict signal by the different (and opposite) relationship to performance and task parameters (Supplementary Fig. 6e).

On canceled trials after SSRT a large fraction of these neurons (40/54) produced persistent spiking; a fraction exceeding Conflict neurons (Chi-square test of homogeneity (one-tailed), $\chi^2(1, N = 129) = 27.3$, $p < 0.001$). The sustained activity attenuated -300 ms after the feedback tone that cued successful performance (Fig. 5a). Attenuation after the tone was also observed on no-stop-signal trials. The spike rate immediately before the feedback tone was unrelated to any factor related to its time or anticipation. Consistent with the indirect contribution of SEF to saccade initiation, the termination of this modulation was time-locked to the tone and not when monkeys stopped fixating on the stop-signal (on canceled trials) or the target (on no-stop-signal trials). Hence, this signal is not directly involved in gaze-holding (Supplementary Fig. 6c).

Based on the evidence above and previous findings identifying SEF signals with working memory^{11,12}, we conjecture that these neurons contribute to maintaining a representation of task goals (e.g., sustain unblinking fixation) for the successful completion of the task. Consistent with this hypothesis, when the monkey broke fixation too early, the facilitation after SSRT was reduced significantly in a subset of neurons with enough data (14/54; Supplementary Fig. 6f). Therefore, we refer to these neurons as Goal Maintenance neurons.

The modulation of only a minority (7/54) of these neurons covaried with RT on the subsequent trial (multiple linear regression (two-tailed) controlling for SSD, $p < 0.05$) which precludes this signal from contributing to adjustments of RT.

Goal Maintenance neurons multiplexed with reinforcement and error signals¹⁷. The vast majority were previously classified as Loss neurons because although the activity of most of these neurons was suppressed after the feedback tone cued success when it cued failure, activity increased¹⁷ (Supplementary Fig. 6d). The prevalence of this multiplexing pattern exceeded chance (Chi-square test of homogeneity (one-tailed), $\chi^2(3, N = 575) = 19.43$, $p < 0.001$; Supplementary Table 7).

To elucidate the microcircuit contribution of Goal Maintenance neurons, we examined spike duration and distribution across the layers. Over one-third of Goal Maintenance neurons were narrow-spiking, a proportion exceeding chance sampling (21/54; Chi-square test of homogeneity (one-tailed), $\chi^2(1, N = 575) = 9.27$, $p = 0.002$). They were found in all penetrations but significantly more commonly at certain sites ($\chi^2(4, N = 575) > 39.3$, $p < 0.001$, Supplementary Table 1). Perpendicular penetrations revealed the time-depth organization of 34

Goal Maintenance neurons (Fig. 5b). The distribution of these neurons across cortical layers was significantly different from the sampling distribution ($\chi^2(4, N = 293) = 11.24$, $p = 0.024$, Supplementary Fig. 1a, Supplementary Table 1) with significantly more in L2/3 relative to L5/6 (Fig. 5b, $\chi^2(1, N = 293) = 10.37$, $p = 0.001$). Their laminar distribution was also significantly different from those of Conflict ($\chi^2(1, N = 70) = 11.54$, $p < 0.001$) and Event Timing neurons ($\chi^2(1, N = 83) = 5.49$, $p = 0.019$). Those in L2/3 modulated significantly earlier than those in L5/6 (L2/3: 85 ± 64 ms, L5/6: 193 ± 101 ; Mann–Whitney U test (two-tailed), $U(26,8) = 388$, $z = -2.7$, $p = 0.007$).

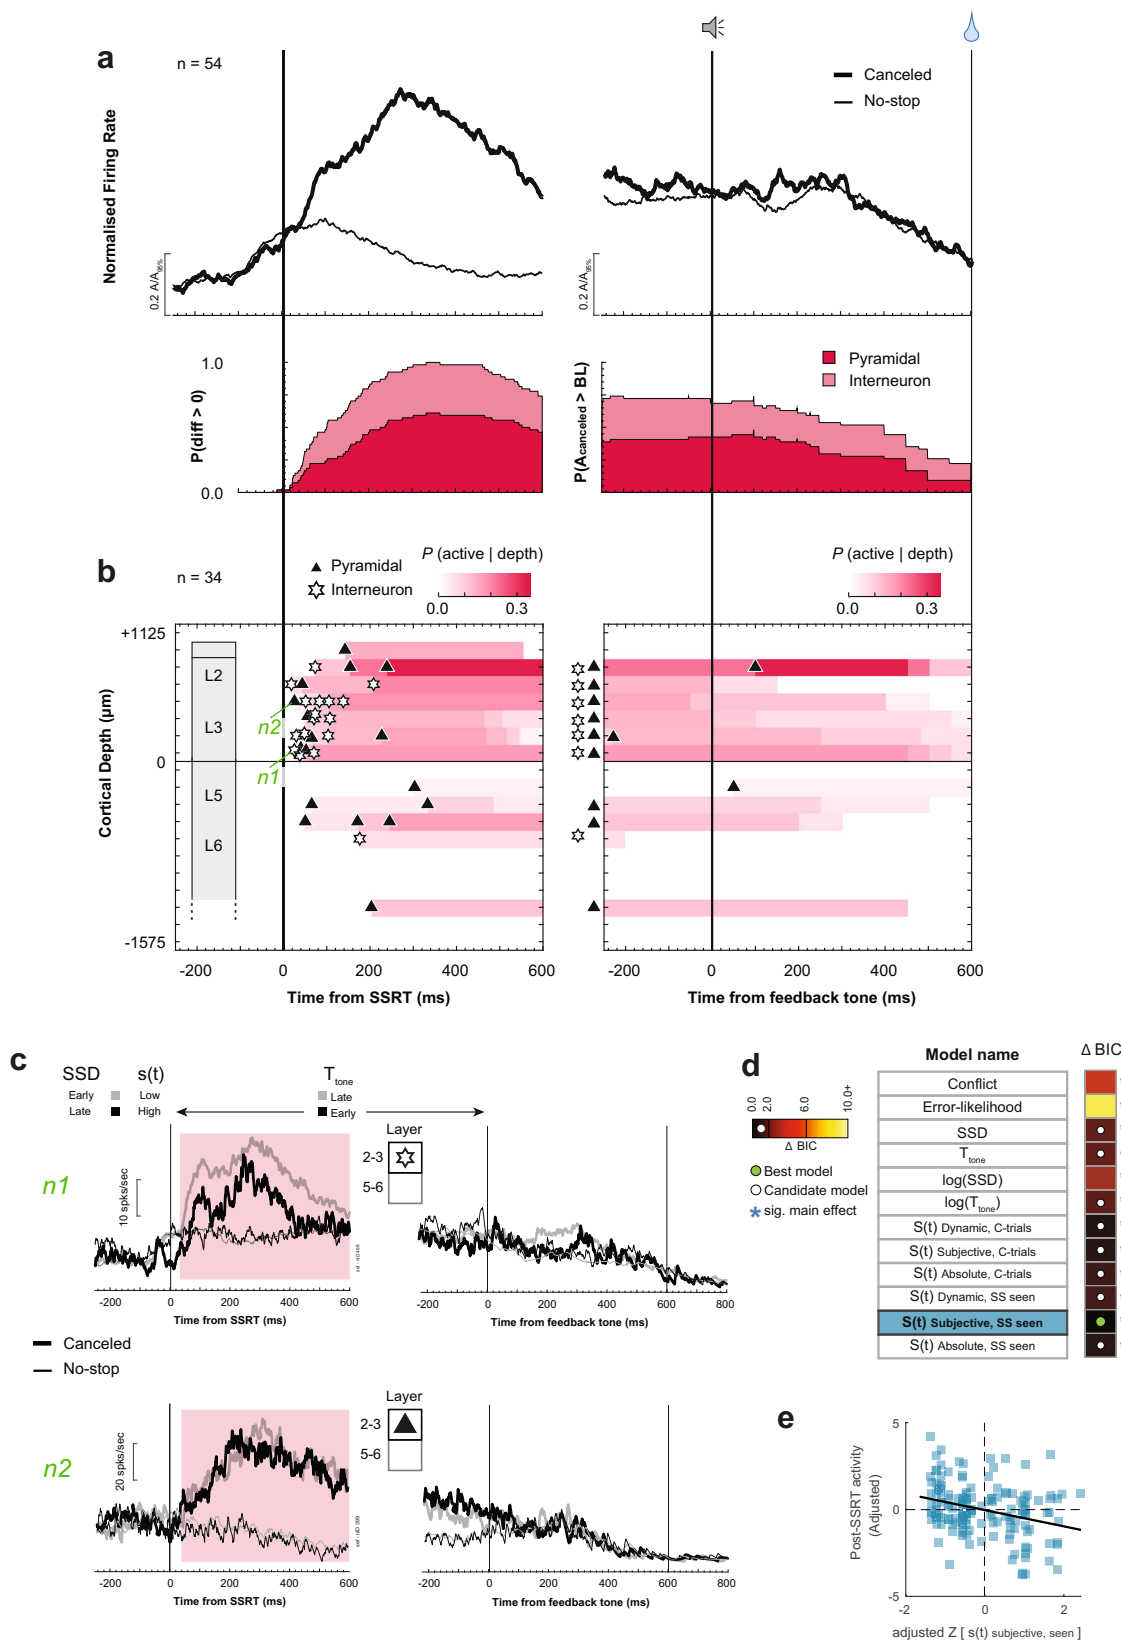
Thus, consistent with previous studies^{11,12}, neurons in SEF produced activity sufficient to enable a working memory representation of the goal of saccade inhibition through time tuned by experienced intervals. These results show that these neurons are most common in L2/3 and a relatively higher proportion have narrow spikes.

Functional classification of N2/P3 ERP related to response inhibition.

To determine whether macaque monkeys produce ERP components associated with response inhibition homologous to humans⁷, we sampled EEG from an electrode located over MFC (10–20 Fz) while recording neural spiking in SEF (Fig. 6a). To isolate signals associated with response inhibition by eliminating components associated with visual responses and motor preparation, we measured the difference in polarization on canceled trials and latency-matched no-stop-signal trials for each SSD (Fig. 6b). Homologous to humans, we observed an enhanced N2/P3 sequence associated with successful stopping. The conclusions drawn from the results presented below do not differ if the analyses are performed on the raw EEG polarization on canceled trials instead of the difference between conditions in these intervals.

The N2, characterized as a negative deflection homologous to the human N2, began -150 ms and peaked at 222 ± 17 ms after the stop-signal, well after the visual ERP polarization (Supplementary Fig. 7a). The N2 was observed after SSRT, too late to index reactive response inhibition. Furthermore, the N2 peak time across sessions was significantly better aligned on stop-signal presentation than on SSRT, further dissociating the N2 from reactive inhibition (F -test comparison of variances (two-tailed), $F(28,28) = 0.29$, $p = 0.002$; Supplementary Fig. 7c). Variation in the amplitude of the N2 was only explained by the error-likelihood model with the largest negativity associated with the lowest error-likelihood (mixed-effects linear regression (two-tailed) grouped by session, $t(85) = 2.42$, $p = 0.018$; Fig. 6d, Supplementary Table 6). Conflict, time-based, and surprise models were rejected (non-significant main effect, and $\Delta\text{BIC} > 3.0$; Fig. 6c). This result adds to the inconclusive evidence for the frontal N2 association with conflict monitoring and response inhibition⁷.

The N2 was followed by a robust P3 beginning -300 ms and peaking 358 ± 17 ms after the stop-signal, homologous to the human P3⁷ (Fig. 6a, b). The peak polarization time was better synchronized on the stop-signal than on SSRT (F -test comparison of variances (two-tailed), $F(28,28) = 0.44$, $p = 0.034$; Supplementary Fig. 7c). Variation in the amplitude of the P3 was best described by the $\log(T_{\text{tone}})$ on canceled trials, with P3 polarization increasing with T_{tone} (Mixed-effects linear regression (two-tailed) grouped by session, $t(85) = 3.72$, $p < 0.001$; Fig. 6c, e; Supplementary Table 6). All time-based models



were candidates ($\Delta BIC < 1.30$). The conflict ($\Delta BIC = 2.90$) and most surprise models ($4.5 < \Delta BIC < 5.7$) received weak support, and error-likelihood and one surprise model were rejected ($\Delta BIC > 6.0$).

Association of N2/P3 with neural spiking. We examined how neural spiking related to concomitant ERP^{17,45}. Appreciating that EEG arises

from -10^6 neurons and spikes are too brief to create scalp EEG⁴⁶, we evaluated whether the N2/P3 complex can be a biomarker of layer-specific neural spiking.

The N2 coincided generally with the peak recruitment of Conflict and of Event Timing neurons, and the P3 with the peak recruitment of Goal Maintenance neurons (Fig. 6b). The relationship between neural

Fig. 5 | Goal maintenance neuron time–depth organization. Conventions as in Fig. 3. **a** Average spike rate and recruitment through time of neurons with persistent activity on canceled relative to latency-matched no-stop-signal trials. **b** Time–depth plot. **c** Modulation on successfully canceled relative to latency-matched no-stop-signal trials for shorter and longer T_{tone} of two representative neurons: n1 had narrow spikes in L3; n2 had broad spikes L3. Spiking of both neurons decreased after the tone. Shaded interval highlights significant differences in spiking across

conditions. **d** Model comparison table for post-SSRT activity. Variation of spiking was best predicted by the surprise model $S(t)_{\text{Subjective, SS seen}}$. Full statistics are in Supplementary Table 5. **e** Significant variation of spiking activity after SSRT as a function of $S(t)_{\text{Subjective, SS seen}}$ with 54 neurons contributing 162 samples across early, intermediate, or late T_{tone} . For 9 neurons no reliable estimate of spike density for the late-SSD bin was obtained due to too few trials. These data points were not included.


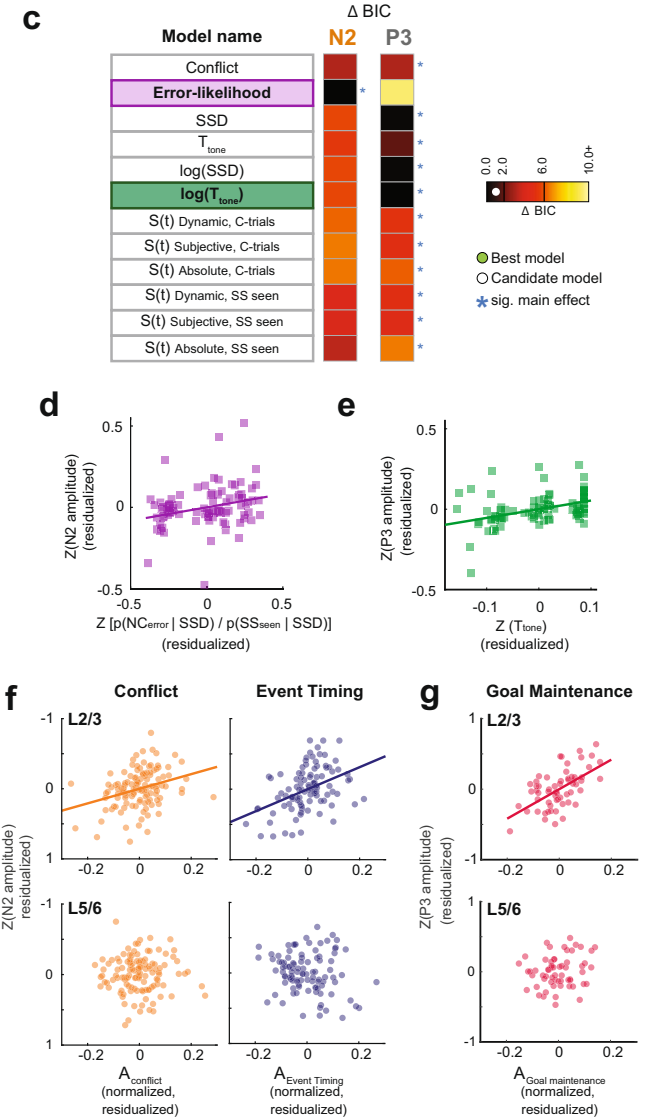


Fig. 6 | Event-related potentials for successful response inhibition. Conventions as in Fig. 3. **a** Grand average z-transformed EEG on canceled (thick) and latency-matched no stop-signal (thin) trials. **b** Difference functions (top) remove stimulus-evoked ERP to highlight N2 and P3 components in 3 SSD bins. Shaded intervals show ± 50 ms sampling interval around N2 (orange) and P3 (gray) peaks. Concomitant recruitment of the three neuron classes (bottom). **c** Model comparison shows N2 amplitude variation was best described by the error-likelihood, and P3 amplitude was best described by $\log(T_{\text{tone}})$. Full statistics in Supplementary Table 6. **d** Significant variation of N2 amplitude as a function of $p(\text{N}_{\text{error}}|\text{SSD})/p(\text{SS}_{\text{seen}}|\text{SSD})$ in 29 sessions contributing 87 points across early, intermediate, or late SSD. **e** Significant variation of P3 amplitude as a function of T_{tone} in 29 sessions contributing 87 points across T_{tone} . **f** Partial regression between N2 amplitude and



spike rate for Conflict (left) and Event Timing (right) neurons in L2/3 (top) and L5/6 (bottom) for sessions with both L2/3 and L5/6 neurons sampled. Ordinate scale plots, with EEG convention, residual from fixed-effects-adjusted ERP amplitude controlling for activity in the opposite layer. Abscissa scale plots residual fixed-effects-adjusted neuronal discharge rate in the identified layer controlling for the activity in the opposite layer and stop-signal delay. Each point plots the average EEG voltage and associated spiking rate in one of 20 bins with equal numbers of trials per session. Plotted are 120 points from 6 sessions for Conflict (left) and 100 points from 5 sessions for Event Timing (right) neurons. N2 amplitude variation was predicted by spiking rate variation of Conflict and Event Timing neurons in L2/3 but not in L5/6. **g** P3 amplitude variation was predicted by spiking rate variation of Goal Maintenance neurons in L2/3 but not in L5/6.

events in SEF and cranial voltages is both biophysical and statistical. The cranial voltage produced by synaptic currents associated with a given spike must follow Maxwell's equations applied to the brain and head, regardless of the timing of the different events. Hence, we

counted the spikes of the three classes of neurons separately in L2/3 and in L5/6 during the 100 ms spanning the peak of the ERP and tested multiple linear regression models with activity in upper layers (L2/3) and lower layers (L5/6) of each neuron class as predictors. Only

successfully canceled trials were included in this analysis. We found that variation in N2 voltage is not associated with the spiking of Goal Maintenance neurons (Multiple linear regression (two-tailed) with L2/3 and L5/6 activity as predictors; L2/3: $t(57) = -1.28$, $p = 0.206$; L5/6: $t(57) = 0.52$, $p = 0.605$; Supplementary Fig. 8a; Supplementary Table 8). However, it was predicted by the spiking in L2/3 but not in L5/6 of Conflict (L2/3: $t(117) = -3.6$, $p < 0.001$; L5/6: $t(117) = 0.046$, $p = 0.963$) and of Event Timing (L2/3: $t(97) = -4.60$, $p < 0.001$; L5/6: $t(97) = 1.67$, $p = 0.097$) neurons (Fig. 6f). When the discharge rate of the L2/3 neurons was higher, the N2 exhibited more negativity. N2 polarization was also predicted by the spiking in L2/3 but not in L5/6 of other neurons that were not modulated on canceled trials and so were not described in this manuscript (L2/3: $t(317) = -2.51$, $p = 0.012$; L5/6: $t(317) = -1.60$, $p = 0.110$; Supplementary Fig. 8a). Conversely, variation in P3 polarization was predicted by the spiking activity of Goal Maintenance neurons in L2/3 but not L5/6 (multiple linear regression (two-tailed) with L2/3 and L5/6 activity as predictors (L2/3: $t(57) = 5.46$, $p < 0.001$; L5/6: $t(57) = 1.47$, $p = 0.148$; Fig. 6g; Supplementary Fig. 8c, d; Supplementary Table 8), with higher spike rates associated with greater P3 positivity. P3 amplitude was unrelated to the spiking of Conflict (L2/3: $t(117) = 0.44$, $p = 0.660$; L5/6: $t(117) = -0.49$, $p = 0.624$), Event Timing (L2/3: $t(97) = -1.19$, $p = 0.236$; L5/6: $t(97) = -0.77$, $p = 0.440$), or unmodulated neurons (L2/3: $t(317) = -1.11$, $p = 0.269$; L5/6: $t(317) = 0.054$, $p = 0.956$; Supplementary Fig. 8c; Supplementary Table 8).

Discussion

These results offer further insights into the cortical microcircuitry supporting executive control in primates. Model-based analysis of the latency, temporal dynamics, and variation in the strength of neural spiking across the neuron sample revealed functionally distinct and theoretically informative classes of neurons with distinct biophysical and laminar properties. Moreover, a bridge between these neurophysiological findings and human electrophysiology was established through the specific associations observed between the N2 and P3 ERP observed in response inhibition tasks and classes of neurons in particular cortical layers.

The utility of these findings is amplified by their complementarity with our previous description of the laminar organization of error and reward processing in SEF¹⁷. Based on the results presented in this paper, we will discuss how SEF can contribute to conflict monitoring, time-keeping, and goal maintenance. Coupled with the current knowledge about the connectivity of SEF^{47–49}, our findings detailing the laminar distribution of neurons signaling response conflict, event timing, and maintaining goals suggest several specific hypotheses and research questions about how SEF and associated structures accomplish response inhibition and executive control (Fig. 7). Also, complementing our earlier description of the source of the ERN¹⁷, we now report a macaque homolog of the N2/P3 ERP components associated with response inhibition. These results demonstrate one potential cortical source of these ERP components.

Conflict, surprise, salience, and dopamine

In this study, we report a population of SEF neurons with pronounced, transient facilitation after successful response inhibition (SSRT). These neurons were predominantly broad-spiking, proportional to the sampling distribution, and found in all layers. Their spike rate during canceled trials was best described by the conflict model, operationalized by the probability of generating noncanceled saccades— $p(\text{NC|SSD})$. The mechanism producing responses in this task is a well-understood network of gaze-shifting and gaze-holding neurons in the frontal eye field (FEF) and superior colliculus (SC)^{27,50}. Neurocomputational models demonstrate how this network can instantiate the GO and STOP processes^{29,30,51}. Noncanceled saccades happen when the gaze-holding STOP units do not interrupt the rise to the threshold of the gaze-shifting GO units. Reactive inhibition happens only if the

STOP unit interrupts the GO unit, which must be brief and potent. Thus, in successful canceled trials, the GO and STOP units (corresponding to gaze-shifting and gaze-holding neurons) are in an unstable state of co-activation, corresponding to the original, formal definition of conflict³¹. The multiplicative conflict between GO and STOP accumulator units scales with $p(\text{NC|SSD})$ and peaks following SSRT (Fig. 2a) but it is unrelated to adjustments in RT. Therefore, we conjecture that this modulation signals the difficulty of the stopping process, which can then be incorporated with other information to drive adaptive changes in behavior.

Recent findings from the nigrostriatal dopamine system of monkeys performing saccade countermanding⁴² offer an alternative interpretation for these neurons. The modulation of DA neurons in dorsolateral SNpc scales with $p(\text{NC|SSD})$ just like the Conflict neurons. Although alternative models were not tested in that study, the spiking of DA neurons has been identified with salience or surprise^{42,52}. We tested the surprise hypothesis by quantifying the moment-by-moment expectancy of the stop-signal given the experienced distribution of SSD and probability of stop-signal occurrence^{37,38}. This neural modulation was explained almost as well by surprise as by conflict. However, it could not be explained by error-likelihood. While these SEF data align with the Conflict hypothesis, they do not exclude surprise or salience hypotheses. From the perspective of reinforcement theory, a phasic DA signal can be an eligibility trace broadcast to SEF and other regions to associate reinforcement with successful cancellation after the infrequent stop-signal. To be most useful, such an eligibility trace must be salient and may be surprising.

Event timing and goal maintenance

We found neurons encoding the timing of task events. In our version of the stop-signal task, knowledge of the timing of the stop-signal and of the feedback tone was important. To earn the reward, monkeys must hold their gaze stable for an extended period, which required preventing eye movements and blinks that would interrupt the camera-based eye tracker. This entails learning and exploiting regularities in the timing of task events^{43,53}. A contribution of SEF and nearby areas in action timing and time production tasks has been demonstrated^{14,15,54}. We extend that description to this stop-signal task in terms of time keeping and goal maintenance.

A distinct group of SEF neurons produced ramping spike rates. When the saccade was inhibited, this ramping was interrupted by a pronounced suppression. These neurons were described previously with no explanation⁸. Our results rule out the possibility that these neurons control movement initiation because the suppression occurred too late to contribute to gaze-holding. Also, this ramping activity did not encode the expectancy for the stop-signal arising from the temporal distribution of SSDs and the probability of stop-signal appearance. Instead, they were best described by SSD.

The task design exposed a second period of ramping before the feedback tone in roughly half of these neurons which reached higher levels for longer durations. Our discovery of an association between the spiking rate and the log-transformed duration of the elapsed time motivates a more integrated interpretation framed by a body of research on time keeping^{14,15,33,34,36,44,54}. We interpret the ramping activity as representing the timing of task events, like neurons in the basal forebrain that signal event timing depending on surprise, salience, and uncertainty⁴⁴. The sharp suppression in activity can reset the system to track the timing of subsequent events. Although the stop-signal occurred randomly and response inhibition was accomplished stochastically, the feedback tone was certain to happen. Therefore, we conjecture that neurons exhibiting ramping activity before both SSRT and the feedback tone encode the timing of expected salient events regardless of certainty or expected response. In contrast, neurons with no ramping activity before the tone can encode events that are less certain in occurrence or consequence. These differences were

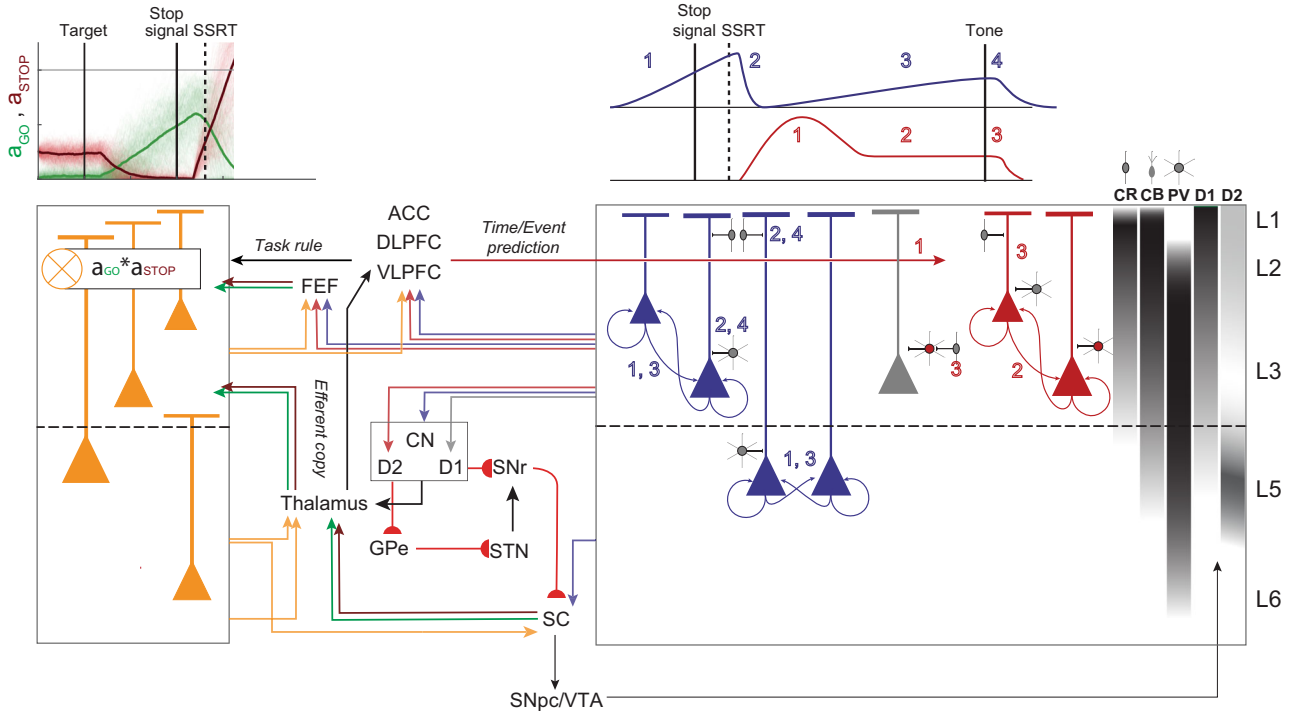


Fig. 7 | Executive control circuitry. Conflict (orange), Event Timing (dark blue), and Goal Maintenance (dark red) neurons are illustrated with selected anatomical connections and laminar densities of calretinin (CR), calbindin (CB), and parvalbumin (PV) neurons²¹ and of D1 and D2 receptors²² (right). Dopamine projections from SNpc and VTA modulate all computations in SEF. *Left*, Conflict signal can arise through coincident inputs of gaze-holding (STOP, dark red) and gaze-shifting neurons (GO, green) directly from FEF and indirectly via thalamus from SC terminating in L2 and L3. Conflict signals integrated across apical and basal dendrites can be sent to multiple cortical and subcortical structures. *Right*, Schematic profiles of Event Timing and Goal Maintenance activity with numbered phases. Event Timing neurons can receive inputs from DLPFC and ACC informing them about an upcoming event. Ramping results from recurrent connections (1, dark blue). SEF can receive information about the stop-signal appearance and successful stopping from VLPFC and DLPFC and from Conflict neurons. This signal can suppress the ramping activity via inhibitory connections onto Event Timing neurons (2, dark

blue), resetting these neurons for the next ramping phase (3, dark blue), which is terminated by the appearance of the feedback tone (4). Event Timing neurons can project to the caudate nucleus (CN) to inform the fronto-striatal reinforcement learning loop about experienced event timing. Goal Maintenance neurons can delay unwanted movement through the push-pull basal ganglia circuitry. Pyramidal neurons can project to the indirect (D2) pathway and inhibitory neurons can project to other pyramidal neurons, unobserved in this study (gray), that project to the direct (D1) pathway. Inputs from DLPFC and ACC terminating in L2/3 can inform SEF of the anticipated timing of task events for successful completion of the task based on the experienced SSD. These inputs can produce the phasic response in Goal Maintenance neurons (1, red) followed by persisting activity via recurrent connections with balanced excitation and inhibition (2, red). The feedback tone, integrated with the task rule from DLPFC, terminates operant control on behavior through CR and CB inhibition of the sustained activity (3). Further details in the text.

reinforced by the distinct laminar distribution of the two groups of neurons.

Another class of neurons produced persistent activity on canceled trials after SSRT, for most neurons lasting until the feedback tone. These neurons were not gaze-holding neurons contributing to response inhibition because this facilitation occurred too late to be involved in reactive inhibition. In neurons with enough trials, we observed weaker modulation when monkeys aborted trials after canceling the saccade. Similar signals have been observed in tasks with a cue indicating reinforcement probability^{55,56}. However, the activity of these SEF neurons was clearly not explained by error-likelihood, which is the inverse of success probability. Instead, time-based and surprise models could explain this modulation. The inverse relationship between spike rate and surprise implies the contribution of inhibitory neurons with spiking rates directly proportional to surprise. We did not sample many such neurons, but narrow-spiking Conflict neurons can serve this role. The direct relationship between spike rate and T_{tone} , on the other hand, resembles a common motif for encoding duration³⁴. Furthermore, this activity was linked to events occurring after but not before successful stopping. Therefore, we believe that this modulation is explained most parsimoniously by T_{tone} predicted by the experienced SSD and not surprise. To earn a reward on canceled trials, monkeys needed to sustain unblinking fixation on the stop-

signal until the feedback tone (Fig. 1a). Hence, we conjecture that these neurons contribute to goal maintenance. This conjecture is consistent with an original theory of response inhibition²⁶ and previous evidence linking SEF to working memory^{11,12} and working memory to time estimation^{57,58}. Future work can discriminate time-based and surprise parameters and evaluate the link between surprise and goal maintenance.

Cortical microcircuitry of executive control

By combining the laminar distribution of the neuron classes described in this study with the anatomical, histological, and neurophysiological properties of SEF neurons, we offer hypotheses about mechanisms by which such signals can be generated and influence other neurons, layers, or brain areas (Fig. 7).

To signal conflict, SEF can be informed about the dynamic state of gaze-shifting and gaze-holding through inputs from FEF and oculomotor thalamic nuclei. Based on previous conjectures⁵ and recent biophysical modeling⁵⁹ we hypothesize that the integration of information producing the modulation of these neurons is derived through synaptic integration across apical and basal dendrites. The circuitry sufficient for signaling prediction errors¹⁹ can signal the occurrence of conflict in this task. The presence of this signal in all layers enables it to interact with all intrinsic processes and possibly influence all cortical

and subcortical efferent targets. For example, a thalamic input that a saccade has been canceled can change the corollary discharge communicated through the thalamus⁶⁰. Consistent with its original conception, communicating conflict (or salience or surprise or just difficulty) to multiple areas simultaneously can coordinate adaptive changes in behavior⁶¹.

The unexpected parallels between SEF and SNpc modulation patterns invite consideration of cause and effect. SEF is innervated by DA neurons in SNpc⁶². Whilst SNpc DA neurons modulated significantly earlier than the SEF Conflict neurons (Supplementary Fig. 9), the estimated arrival times of DA spikes in SEF were not significantly different from the modulation times of the Conflict neurons after accounting for the very slow conduction of DA axons⁶³ (~100 ms conduction time from SNpc to SEF, Supplementary Fig. 9) and second-messenger delay. Therefore, we infer that this transient modulation in SEF cannot be caused by DA inputs. Conversely, because axon terminals from SEF are rare in SNpc^{47,49}, SEF neurons are unlikely to directly cause the modulation of the SNpc DA neurons. Instead, other investigators have shown that the phasic DA activation is delivered by the SC⁶⁴. Through the Conflict neurons in L5, SEF can influence SC directly⁴⁷. Curiously, though, the modulation specifically after SSRT scaling with related performance parameters has not been observed in SC⁵⁰.

While Event Timing neurons were found in all layers, those encoding timing regardless of event predictability or action were most common in L3 and L5 with broad spikes consistent with pyramidal neurons. This laminar differentiation demonstrates that the timing of different events can be conveyed by different layer-specific extrinsic connections. The timing signal can be sent via cortico-cortical connections to other cortical areas to influence motor, cognitive, and limbic processes. Also, these neurons can contribute to fronto-striatal pathways to learn and update a representation of the temporal structure of the task^{54,65,66}. Axon terminals from SEF are dense in the caudate nucleus⁴⁸, arising from pyramidal neurons in L3 and L5^{67,68}. In fact, neurons with this pattern of modulation have been described in a recent investigation of the caudate nucleus of monkeys performing saccade countermanding⁴². Our finding that the suppression in the SEF Event Timing neurons occurred after SSRT, but significantly earlier than those previously reported in the caudate nucleus suggests a primary role of the cortex in this signaling (Supplementary Fig. 9)⁴².

The rapid suppression of the ramping activity after SSRT merits consideration. One source can be intracortical inhibition from the narrow-spiking, putative parvalbumin (PV) neurons that we observed. Another source can be the very small calbindin (CB) and calretinin (CR) neurons concentrated in L2/3 that are innervated by the dorsolateral prefrontal cortex (DLPFC) and selectively inhibit pyramidal neurons⁶⁹, although our methods are unlikely to sample their spikes. We note that although SEF is an agranular structure with weak interlaminar inhibitory connections²², CR neurons in L2/3 can potentially inhibit L5 neurons through specialized projections on the apical dendrites⁷⁰. This inhibition must be informed about the presence of the stop-signal and inhibition of the saccade. We observe that such a signal is available in the Conflict neurons. However, the suppression of Event Timing neurons occurred significantly earlier than the facilitation of the Conflict neurons. Further research can resolve these cortical interactions.

Goal Maintenance neurons were mainly found in L2/3. Inputs to these layers from DLPFC, ventrolateral prefrontal cortex (VLPFC), and anterior cingulate cortex (ACC) can signal task rules and the expected time of the secondary reinforcer when executive control can be released. Input from these areas can inform SEF of the anticipated reward based on the experienced stop-signal delay, contingent on successful stopping. Dopaminergic release in SEF from SNpc and ventral tegmental area (VTA), with similar time-predicting signals⁷¹, can enhance these influences through the higher density of D1 relative to D2 receptors in L2/3⁷². The sustained discharge can be maintained

through recurrent activation within SEF and between other structures¹². Also, many Goal Maintenance neurons had narrow spikes, consistent with PV inhibitory neurons, which can balance excitation and inhibition necessary for the maintenance of persistent activity in recurrent networks^{73,74}. The auditory feedback tone, integrated with the task rule from DLPFC, cues the termination of operant control on behavior, resulting in the inhibition of pyramidal and inter-neurons by CR and CB neurons. This results in the termination of the sustained activity.

Pyramidal Goal Maintenance neurons can encourage the suppression of movements through projections to the indirect pathway D2 neurons in the striatum^{67,68}. Intrinsic inhibitory Goal Maintenance neurons can suppress the movement by inhibiting pyramidal neurons projecting the direct pathway D1 and to the frontal eye field. As PV neurons in primates largely lack extrinsic connections⁷⁵, we propose that this can be mediated by the inhibition of other excitatory neurons (unidentified neurons and possibly Gain neurons identified in ref. 17) that send projections to these motor structures (gray neurons in Fig. 7, right panel). Therefore, Goal Maintenance neurons can achieve their role by altering the balance in the push-pull mechanism mediated by the direct (D1) and indirect (D2) pathways. This function is consistent with the observation that many of these neurons also exhibited higher activity after the feedback tone on unrewarded trials, previously described, and this activity influenced post-error adjustments in RT¹⁷. This is while the facilitation of response in Goal Maintenance neurons did not influence post-canceling changes in RT. Therefore, it is possible that these signals have different influences on their efferent targets depending on the task epoch.

We note that neurons with facilitated activity after SSRT were described in an investigation of the caudate nucleus of monkeys performing saccade countermanding⁴². The facilitation in the caudate nucleus coincided with that measured in SEF (Supplementary Fig. 9). The parallel between SEF and the striatum in patterns of modulation associated with proactive but not reactive inhibition is surprisingly, but satisfyingly, clear.

Event-related potentials

We showed that macaque monkeys exhibit an N2/P3 ERP complex homologous to humans⁷. Disagreement persists about the frontal N2 and P3 index^{76,77}. We found that the amplitude of the N2 during the stop-signal task varied most with the likelihood of error associated with experienced SSDs, and not conflict or SSD as previously suggested^{7,78}. Consistent with previous reports of P3 indexing expectation and timing of behavior⁷⁶, we found that P3 amplitude co-varied most with the expected time of the feedback tone.

Variation in N2 and P3 polarization was predicted by the spiking of specific neuron classes in L2/3 and not L5/6. N2 magnitude was unrelated to the spiking of Goal Maintenance neurons but co-varied with the spiking of Conflict and Event Timing neurons, as well as the spiking of other neurons that did not modulate around the time of successful stopping. In contrast, P3 amplitude was predicted by the spiking of Goal Maintenance but not Conflict or Event Timing neurons. Also, N2 timing coincided with maximal modulation of Conflict and Event Timing neurons while P3, with Goal Maintenance neurons. A relationship between L2/3 spiking and these ERPs may appear trivial because the upper layers are closer to the surface EEG electrodes, but the result merits attention for several reasons. First, action potentials are too brief to generate the EEG, so the association with L2/3 spiking entails an association with coherent synaptic potentials⁷⁹. Second, EEG polarization is related to the strength and orientation as much as the proximity of a dipole, and biophysical models of EEG sources assume that the larger L5 pyramidal neurons are the major contributor⁴⁶. Finally, because ERPs arise from multiple sources, a dipole in one region can be canceled by a dipole oppositely oriented in another region. Therefore, it is unlikely that these L2/3 neurons are directly

causing these ERPs. It is more likely that the same synaptic potentials that result in the activation of these L2/3 neurons are also generating the EEG. Our results establish the N2 and P3 as possible biomarkers of the activity of neurons in the upper layers of SEF serving executive control functions. These results indicate that cognitive ERPs reflect diverse neuro-computational processes, rendering unitary and exclusive hypotheses incomplete.

Incidence and multiplexing of signals

Here and in our previous report, we distinguished specific categories of neural signals. Different numbers of units were sampled in each category. Knowing that neural sampling with extracellular recording is biased in various ways, we cannot infer with confidence the importance of a process or the magnitude of computational contribution based on the number of units sampled. However, we will consider the reliability of the categories and the mixture of signals produced by single neurons.

The three signals reported here were multiplexed in single neurons with previously reported error and reinforcement signals¹⁷. That is, some neurons produced one kind of functional signal around the time of successful stopping and another at the time of reinforcement. Such multiplexing has been observed previously and can appear for several reasons^{80,81}. First, modulation patterns may be too weak and variable to distinguish classes of neurons. Our selection criteria for neurons to analyze avoided this by including only neurons with distinct patterns of modulation. Second, frontal lobe neurons support diverse inputs and outputs from multiple cortical areas and subcortical nuclei. Therefore, a neuron can participate in partially overlapping but distinct networks such that in one state neurons broadcast one signal to some efferent targets, while in another state they broadcast another signal to other efferent targets. Theories about mixed-selectivity and dynamical systems have emphasized state-dependent dynamics^{6,80}, but they have not incorporated the specificity of laminar properties derived from specialized connectivity. Third, our classification of signals was based on response dynamics around the time of successful stopping, but we know of no theoretical or empirical prohibition against neurons modulating in association with multiple events. Ultimately, different neurons in different layers receive different inputs and have different outputs. Therefore, understanding the laminar distribution of signals reported in this study is a necessary step toward formulating more specific hypotheses about how neural networks function⁸².

In conclusion, the present results add to the first catalog of the laminar functional architecture of an agranular frontal lobe area. Pioneering insights into the microcircuitry and mechanisms of the primary visual cortex began by describing the properties of neurons in different layers¹⁸. Contrasts of the results reported here with primary sensory areas will reveal the degree of computational uniformity across cortical areas. Being a source contributing to ERPs indexing performance monitoring and executive control, details about laminar processing in SEF offer unprecedented insights into the microcircuitry of executive control. These results validate the tractability of formulating neural mechanism models of performance monitoring and executive control, especially when constrained by formal²⁶, algorithmic^{29,30}, and spiking network⁵¹ models of performance of a task with clear clinical relevance⁸³.

Methods

Animal care and surgical procedures

Data was collected from one male bonnet macaque (Eu, *Macaca radiata*, 8.8 kg, 6 y.o.) and one female rhesus macaque (X, *Macaca mulatta*, 6.0 kg, 8 y.o.) performing a countermanding task^{21,25}. All procedures were in accordance with the National Institutes of Health Guidelines, the American Association for Laboratory Animal Care Guide for the Care and Use of Laboratory Animals and approved by the

Vanderbilt Institutional Animal Care and Use Committee in accordance with the United States Department of Agriculture and Public Health Service policies. Surgical details have been described previously⁸⁴. Briefly, magnetic resonance images (MRIs) were acquired with a Philips Intera Achieva 3 T scanner using SENSE Flex-S surface coils placed above or below the animal's head. T1-weighted gradient-echo structural images were obtained with a 3D turbo field echo anatomical sequence (TR = 8.729 ms; 130 slices, 0.70 mm thickness). These images were used to ensure Cilux recording chambers were placed in the correct area. Chambers were implanted normal to the cortex (Monkey Eu: 17°; Monkey X: 9°; relative to stereotaxic vertical) 1 mm right of the midline, 30 mm (Monkey Eu) and 28 mm (Monkey X) anterior to the interaural line.

Acquiring EEG

EEG was recorded from the cranial surface with electrodes located over the MFC. Electrodes were referenced to linked ears using ear-clip electrodes (Electro-Cap International). The EEG from each electrode was amplified with a high-input impedance head stage (Plexon) and bandpass filtered between 0.7 and 170 Hz. Trials with blinks within 200 ms before or after the analysis interval were removed.

Cortical mapping and electrode placement

Chambers implanted over the medial frontal cortex were mapped using tungsten microelectrodes (2–4 M Ω , FHC, Bowdoin, ME) to apply 200 ms trains of biphasic micro-stimulation (333 Hz, 200 μ s pulse width). The SEF was identified as the area from which saccades could be elicited using <50 μ A of current^{85,86}. In both monkeys, the SEF chamber was placed over the left hemisphere. The dorsomedial location of the SEF makes it readily accessible for linear electrode array recordings across all cortical layers. A total of five penetrations were made into the cortex—two in monkey Eu, and three in monkey X. Three of these penetration locations were perpendicular to the cortex. In monkey Eu, the perpendicular penetrations sampled activity at site P1, located 4 mm lateral to the midline and 31 mm anterior to the interaural line. In monkey X, the perpendicular penetrations sampled activity at sites P2 and P3, located 4 mm lateral to the midline and 29 and 30 mm anterior to the interaural line, respectively.

Acquiring neural spiking

Spiking activity and local field potentials were recorded using a 24-channel Plexon U-probe with 150 μ m between contacts, allowing sampling from all layers. The U-probes were 100 mm in length with 30 mm reinforced tubing, 210 μ m probe diameter, 30° tip angle, with 500 μ m between the tip and first contact. Contacts were referenced to the probe shaft and grounded to the headpost. We used custom-built guide tubes consisting of 26-gauge polyether ether ketone (PEEK) tubing (Plastics One, Roanoke, VA) cut to length and glued into 19-gauge stainless steel hypodermic tubing (Small Parts Inc., Logansport, IN). This tubing had been cut to length, deburred, and polished so that they effectively support the U-probes as they penetrated the dura and entered the cortex. The stainless-steel guide tube provided mechanical support, while the PEEK tubing electrically insulated the shaft of the U-probe, and provided an inert, low-friction interface that aided in loading and penetration.

Microdrive adapters were fit to recording chambers with <400 μ m of tolerance and locked in place at a single radial orientation (Crist Instruments, Hagerstown, MD). After setting up hydraulic microdrives (FHC, Bowdoin, ME) on these adapters, pivot points were locked in place by means of a custom mechanical clamp. Neither guide tubes nor U-probes were removed from the microdrives once recording commenced within a single monkey. These methods ensured that we were able to sample neural activity from precisely the same location relative to the chamber on repeated sessions.

Electrophysiology data were processed with unity-gain high-input impedance head stages (HST/32o25-36P-TR, Plexon). Spiking data were bandpass filtered between 100 Hz and 8 kHz and amplified 1000 times with a Plexon preamplifier, filtered in software with a 250 Hz high-pass filter, and amplified an additional 32,000 times. Waveforms were digitized at 40 kHz from -200 to 1200 μ s relative to voltage threshold crossings. Thresholds were typically set at 3.5 standard deviations from the mean. All data were streamed to a single data acquisition system (MAP, Plexon, Dallas, TX). Time stamps of trial events were recorded at 500 Hz. Single units were sorted online using a software window discriminator and refined offline using principal components analysis implemented in Plexon offline sorter.

Cortical depth and layer assignment

The retrospective depth of the electrode array relative to gray matter was assessed through the alignment of several physiological measures. Firstly, the pulse artifact was observed on a superficial channel which indicated where the electrode was in contact with either the dura mater or epidural saline in the recording chamber; these pulsed visibly in synchronization with the heartbeat. Secondly, a marked increase of power in the gamma frequency range (40–80 Hz) was observed at several electrode contacts, across all sessions. Previous literature has demonstrated elevated gamma power in superficial and middle layers relative to deeper layers⁸⁷. Thirdly, an automated depth alignment procedure was employed which maximized the similarity of CSD profiles evoked by passive visual stimulation between sessions²¹. Further details can be found in ref. 21.

Further support for the laminar assignments was provided by an analysis of the depths of SEF layers measured in histological sections visualized with Nissl, neuronal nuclear antigen (NeuN), Gallyas myelin, acetylcholinesterase (AChE), non-phosphorylated neurofilament H (SMI-32), and the calcium-binding proteins parvalbumin (PV), calbindin (CB), and calretinin (CR)^{17,21}. Additional information about the laminar structure was assessed through the pattern of cross-frequency phase-amplitude coupling across SEF layers²³. Owing to variability in the depth estimates and the indistinct nature of the L6 border with white matter, some units appeared beyond the average gray-matter estimate; these were assigned to the nearest cellular layer.

Acquiring eye position

Eye position data were collected at 1 kHz using an EyeLink 1000 infrared eye-tracking system (SR Research, Kanata, Ontario, Canada). This was streamed to a single data acquisition system (MAP, Plexon, Dallas, TX) and combined with other behavioral and neurophysiological data streams.

Data collection protocol

Visual stimulus generation, event timing, and task control were done with TEMPO (version 16.12, set 49.20; Reflective Computing). The same protocol was used across monkeys and sessions. In each session, the monkey sat in an enclosed primate chair with its head restrained 45 cm from a CRT monitor (Dell P1130, background luminance of 0.10 cd/m²). The monitor had a refresh rate of 70 Hz, and the screen subtended $46^\circ \times 36^\circ$ of the visual angle. After advancing the electrode array to the desired depth, they were left for 3–4 h until recordings stabilized across contacts. This led to consistently stable recordings with single units typically held indefinitely. Once these recordings stabilized, an hour of resting-state activity in near-total darkness was recorded. This was followed by the passive presentation of visual flashes followed by periods of total darkness in alternating blocks. Finally, the monkey performed approximately 2000 trials of the saccade countermanding (stop-signal) task per session.

Countermanding task

The countermanding (stop-signal) task utilized in this study has been widely used previously¹. Briefly, trials were initiated when monkeys fixated on a central point. Following a variable time period, drawn from an aging function to avoid anticipation of the visual stimulus⁵³, the center of the fixation point was removed leaving an outline. Simultaneously, a peripheral target was presented to the left or right of the screen.

On no-stop-signal trials, the monkey was required to shift its gaze to the target. Fixation on the target was required for 600 ms until an auditory tone sounded, whereupon monkeys could shift their gaze anywhere. The fluid reward was delivered 600 ms later. On stop-signal trials, comprising less than half of all trials, the center of the fixation point was re-illuminated after a variable stop-signal delay (SSD). An initial set of SSDs, typically separated by ~ 50 ms for Monkey Eu (45 ± 15 ms) and by ~ 100 ms for monkey X (115 ± 17 ms), were selected for each recording session. The selection of SSDs was adjusted to the idiosyncrasies of each subject to ensure performance satisfying key criteria for the stop-signal tasks. Different SSD values were used for the two subjects to account for between-subject differences in stopping performance¹. To ensure that monkeys failed to countermand on $\sim 50\%$ of stop-signal trials, SSD was adjusted through an adaptive staircasing procedure. When a monkey failed to inhibit a response, the SSD was decreased by 1, 2, or 3 steps (randomly drawn) to increase the likelihood of success on the next stop trial. When a monkey canceled the saccade, SSD was increased by 1, 2, or 3 steps (randomly drawn) to decrease the likelihood of success on the next stop trial. On stop-signal trials, the monkey was required to maintain fixation on the central point until the tone sounded, whereupon monkeys could shift gaze anywhere. The fluid reward was delivered 600 ms later. By design, the duration from the target presentation until the tone was a fixed interval of 1500 ms. Thus, as SSD increased, the duration of fixation decreased (Fig. 1a).

Performance on this task is characterized by the probability of not canceling a saccade as a function of the SSD (the inhibition function) and the distribution of latencies of correct saccades in no-stop-signal trials and of noncanceled error saccades in stop-trials. Performance of the stop-signal task is explained as the outcome of a race between a GO and a STOP process²⁶. The race model provides an estimate of the duration of the covert STOP process, and the time taken to accomplish response inhibition, known as stop-signal reaction time (SSRT)^{29,30,51}. SSRT was calculated using two approaches—the conventional weighted-integration method and the more recent Bayesian Ex-Gaussian Estimation of Stop-Signal RT distributions (BEEST)⁸⁸ (Supplementary Figs. 4a, 5a, 6a). Compared to the weighted integration method, the Bayesian approach provides estimates of the variability in SSRT and the fraction of trigger failures for a given session⁸⁸. Individual parameters were estimated for each session. Based on previous implementations of this approach, the priors were bounded uniform distributions ($\mu_{Go}, \mu_{Stop}: U(0.001, 1000)$; $\sigma_{Go}, \sigma_{Stop}: U(1, 500)$; $\tau_{Go}, \tau_{Stop}: U(1, 500)$; pTF: $U(0,1)$). The posterior distributions were estimated using Metropolis-within-Gibbs sampling ran multiple through three chains. We ran the model for 5000 samples with a thinning of 5. None of our conclusions depend on the choice of SSRT calculation method. Full summary statistics of this analysis can be found in the data repository.

Analysis of EEG

Methods paralleling those used in human studies were used. The N2 and P3 were obtained from average EEG synchronized on stop-signal presentation. Peak N2 was the time when the mean ERP reached maximal negativity in a 150–250 ms window after the stop-signal. Peak P3 was the time when the mean ERP was in a 250–400 ms window after the stop-signal. The amplitude of the N2 and P3 was quantified as the mean Z-transformed voltage for each SSD in a ± 50 ms window around

the maximal ERP deflection determined for each session. Indistinguishable results were obtained with wider (± 75 ms), and narrower (± 25 ms) windows or just the instantaneous maximal polarization. To characterize the polarizations associated with response inhibition, a difference ERP (Δ ERP) was obtained by subtracting from the ERP recorded on canceled trials the ERP recorded on RT-matched no-stop-signal trials.

Analysis of neural spiking

Spike widths exhibited a bimodal distribution (Supplementary Fig. 1a), and neurons were distinguished as narrow- (peak-to-trough duration ≤ 250 μ s) or broad-spikes (>250 μ s). Measurements of neural spiking were based on spike density functions (SDF) produced by convolving the spike train with a kernel resembling a postsynaptic potential defined by

$$\text{SDF}(t) = \left(1 - \exp\left(\frac{-t}{\tau_g}\right)\right) \times \exp\left(\frac{-t}{\tau_d}\right) \quad (3)$$

with growth time constant (τ_g) of 1 ms, and decay time constant (τ_d) of 20 ms, corresponding to the values measured for excitatory postsynaptic potentials. The area of the kernel was set to equal 1. To analyze spiking activity associated with successful stopping, we compared the activity on canceled trials and on no-stop-signal trials with RT greater than SSD + SSRT. This latency-matching compares trials in which countermanding was successful with trials in which countermanding would have been successful had the stop signal been presented. Neurons were distinguished by patterns of modulation consisting of periods of facilitation or suppression using a consensus clustering algorithm²⁸ (Fig. 1c, Supplementary Fig. 1c–e). The input to this analysis pipeline was the SDF on canceled trials and on latency-matched no-stop-signal trials during the 100 ms preceding SSRT and 200 ms following SSRT. Results did not change much if interval durations were changed.

To prevent outlying values from exerting excessive influence, population spike density plots were obtained by scaling the SDF of each neuron by the 95% confidence interval between the 2.5% lowest rate and the 97.5% highest rate in one of two intervals. The first interval was a 600 ms window centered on SSRT on canceled and on no-stop-signal trials. The second interval was -100 to $+300$ ms relative to the feedback tone.

To identify spiking modulation, we applied methods previously employed. First, we calculated a difference function (Δ SDF), the difference between the SDF on canceled and latency-matched no-stop-signal trials. Periods of statistically significant modulation were identified based on multiple criteria—(a) the difference function must exceed by at least 2 standard deviations a baseline difference measured in the 100 ms interval before the target appeared, (b) the difference must occur from 50 ms before to 900 ms after the stop-signal, and (c) the difference must persist for at least 100 ms (or for 50 ms if the difference exceeded baseline by 3 standard deviations). As commonly found in the MFC, some neurons exhibited low spiking rates. To obtain reliable estimates of modulation times, we also convolved the SDF with a square 8 ms window. The modulation intervals were validated by manual inspection.

To determine modulation associated with the systematically variable timing of the feedback tone on canceled trials, the SDF was compared against the minimum value found between 500 ms before and 900 ms after the tone. Focusing on modulation occurring only during the period of operant control on behavior, modulations beginning less than 300 ms after the tone were not included. For comparisons across neurons and sessions, Z-transformed SDF or Δ SDF was used.

Mixed-effects models

We fit variation in spike counts or EEG voltage to models of task events and performance outcomes (Fig. 1, Supplementary Table 2, Supplementary Fig. 3). To determine which performance measure accounted best for the variation of neural measures, the performance and neural quantities were averaged within groups of early-, mid-, and late-SSD trials. SSD values greater than -350 ms were not included because too few canceled trials were obtained. The analysis of the facilitation after SSRT was based on Δ SDF (Figs. 3 and 5), but the major conclusions held if the analysis used SDF. The analysis of the modulation before SSRT or the feedback tone (Fig. 4) was based on the SDF of canceled trials. Before SSRT the SDF of canceled and no-stop-signal trials was not different. Before the feedback tone, the interval was variable on canceled trials but not on no-stop-signal trials, and longer on canceled relative to no-stop-signal trials.

Each model was defined by one parameter. A limited number of SSDs (typically between 3–7, divided into 3 bins) offered sufficient trials for analysis. Because of the dependence of the tested parameters on SSD, the values for each parameter varied with SSD. Although performance and behavioral parameters could be correlated, the non-linear relations between them and session-wise variations empowered effective model comparison (Supplementary Fig. 3b).

Mixed-effects models of Δ SDF, SDF, or Δ ERP values in relation to the various performance measures were compared using Bayesian Information Criteria (BIC), which provides a criterion for model selection. We report the results of the most basic version of each model with the main effect term corresponding to the performance parameter and random intercepts grouped by a neuron (for spiking) or session (for EEG). The values for each performance parameter were z-transform normalized for a fair comparison between models related to different quantities. All models had the same degrees of freedom, allowing direct comparison of BIC values between models. The smallest BIC identified the best model. Δ BIC ($\text{BIC}_{\text{best}} - \text{BIC}_{\text{competing}}$) quantified the fit of the other models relative to the best with Δ BIC < 2 were considered as candidate models, $2 < \Delta$ BIC < 6 earning weak support, and Δ BIC > 6 supporting rejection of the model^{89,90}. Models with non-significant main effects were also rejected. More complex versions of these models resulted in similar conclusions. Mixed-effects models were performed using MATLAB's Statistical Toolbox.

Relating RT and neural spiking

All neurons were tested. For facilitated neurons, we counted spikes immediately following SSRT, within the neuron-specific modulation window. For suppressed neurons, we counted spikes in a 20 ms interval centered on the peak of the ramp prior to the suppression. Because both spiking activity and RT varied with SSD, a multiple linear regression model controlled for SSD to test whether spike rates varied with RT.

Relating N2/P3 and neural spiking

We used the method described previously to establish the relationship between spiking activity and the ERN¹⁷. Single-trial spiking was the mean convolved spike data for that trial recorded from neurons in L2/3 and in L5/6 of perpendicular penetrations within ± 50 ms of the N2 and P3 peaks. To account for variations in ERP voltage and spike counts across sessions, a fixed-effects adjustment was performed by centering each distribution on its mean and dividing by its most extreme value. To measure the N2/P3 amplitudes robustly, we grouped rank-ordered single-trial ERP values into 20 successive bins. From trials in each bin, we calculated the mean N2 and mean P3 magnitude (dependent variable), the mean spike count in the upper and lower layers (independent variables), and the average SSD, on Canceled trials. Data from all sessions were combined for a pooled partial correlation. Each point in Fig. 6 plots the paired values of the mean normalized ERP voltage and normalized activity for each of the 20 bins from every session. The

statistical relationship between ERP magnitude and spiking activity was quantified through multiple linear regression on normalized data pooled across sessions. Two factors were considered: (1) spiking activity in L2/3, and (2) spiking activity in L5/6. Both spiking activity and ERPs also varied with SSD but the inclusion of SSD as a predictor did not change the results.

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Other aspects of the behavioral and neural dataset have been published previously^{17,21,23}. The processed data used in this study are openly available online through the Open Science Framework (OSF). <https://doi.org/10.17605/OSF.IO/3RZJX>. Available online via: <https://osf.io/3rzjx/>. The raw data analyzed in the current study are available from the corresponding author on reasonable request.

Code availability

The analysis codes used in this study are openly available online through the Open Science Framework (OSF). <https://doi.org/10.17605/OSF.IO/3RZJX>. Available online via: <https://osf.io/3rzjx/>.

References

- Verbruggen, F. et al. A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *Elife* **8**, e46323 (2019).
- Emeric, E. E. et al. Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vis. Res.* **47**, 35–49 (2007).
- Kolling, N. et al. Value, search, persistence and model updating in anterior cingulate cortex. *Nat. Neurosci.* **19**, 1280–1285 (2016).
- Shenhav, A., Cohen, J. D. & Botvinick, M. M. Dorsal anterior cingulate cortex and the value of control. *Nat. Neurosci.* **19**, 1286–1291 (2016).
- Cohen, M. X. A neural microcircuit for cognitive conflict detection and signaling. *Trends Neurosci.* **37**, 480–490 (2014).
- Fu, Z. et al. The geometry of domain-general performance monitoring in the human medial frontal cortex. *Science* **376**, eabm9922 (2022).
- Kok, A., Ramautar, J. R., De Ruyter, M. B., Band, G. P. & Ridderinkhof, K. R. ERP components associated with successful and unsuccessful stopping in a stop-signal task. *Psychophysiology* **41**, 9–20 (2004).
- Stuphorn, V., Brown, J. W. & Schall, J. D. Role of supplementary eye field in saccade initiation: executive, not direct, control. *J. Neurophysiol.* **103**, 801–816 (2010).
- Stuphorn, V. & Schall, J. D. Executive control of countermanding saccades by the supplementary eye field. *Nat. Neurosci.* **9**, 925–931 (2006).
- Pouget, P. et al. Neural basis of adaptive response time adjustment during saccade countermanding. *J. Neurosci.* **31**, 12604–12612 (2011).
- Reinhart, R. M. et al. Homologous mechanisms of visuospatial working memory maintenance in macaque and human: properties and sources. *J. Neurosci.* **32**, 7711–7722 (2012).
- Bastos, A. M., Loonis, R., Kornblith, S., Lundqvist, M. & Miller, E. K. Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. *Proc. Natl Acad. Sci. USA* **115**, 1117–1122 (2018).
- Kawaguchi, N. et al. Surprise signals in the supplementary eye field: Rectified prediction errors drive exploration-exploitation transitions. *J. Neurophysiol.* **113**, 1001–1014 (2015).
- Wang, J., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
- Ohmae, S., Lu, X., Takahashi, T., Uchida, Y. & Kitazawa, S. Neuronal activity related to anticipated and elapsed time in macaque supplementary eye field. *Exp. Brain Res.* **184**, 593–598 (2008).
- Stuphorn, V., Taylor, T. L. & Schall, J. D. Performance monitoring by the supplementary eye field. *Nature* **408**, 857–860 (2000).
- Sajad, A., Godlove, D. C. & Schall, J. D. Cortical microcircuitry of performance monitoring. *Nat. Neurosci.* **22**, 265–274 (2019).
- Hubel, D. H. & Wiesel, T. N. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* **195**, 215–243 (1968).
- Bastos, A. M. et al. Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
- Shipp, S. The importance of being agranular: a comparative account of visual and motor cortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **360**, 797–814 (2005).
- Godlove, D. C., Maier, A., Woodman, G. F. & Schall, J. D. Microcircuitry of agranular frontal cortex: testing the generality of the canonical cortical microcircuit. *J. Neurosci.* **34**, 5355–5369 (2014).
- Beul, S. F. & Hilgetag, C. C. Towards a “canonical” agranular cortical microcircuit. *Front. Neuroanat.* **8**, 165 (2014).
- Ninomiya, T., Dougherty, K., Godlove, D. C., Schall, J. D. & Maier, A. Microcircuitry of agranular frontal cortex: contrasting laminar connectivity between occipital and frontal areas. *J. Neurophysiol.* **113**, 3242–3255 (2015).
- Rapan, L. et al. Multimodal 3D atlas of the macaque monkey motor and premotor cortex. *Neuroimage* **226**, 117574 (2021).
- Hanes, D. P. & Schall, J. D. Countermanding saccades in macaque. *Vis. Neurosci.* **12**, 929–937 (1995).
- Logan, G. D. & Cowan, W. B. On the ability to inhibit thought and action - a theory of an act of control. *Psychol. Rev.* **91**, 295–327 (1984).
- Hanes, D. P., Patterson, W. F. 2nd & Schall, J. D. Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J. Neurophysiol.* **79**, 817–834 (1998).
- Lowe, K.A. & Schall, J.D. Functional categories of visuomotor neurons in macaque frontal eye field. *eNeuro* **5**, eNeuro.0131-0118.2018 (2018).
- Boucher, L., Palmeri, T. J., Logan, G. D. & Schall, J. D. Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychol. Rev.* **114**, 376–397 (2007).
- Logan, G. D., Yamaguchi, M., Schall, J. D. & Palmeri, T. J. Inhibitory control in mind and brain 2.0: blocked-input models of saccadic countermanding. *Psychol. Rev.* **122**, 115–147 (2015).
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S. & Cohen, J. D. Conflict monitoring and cognitive control. *Psychol. Rev.* **108**, 624–652 (2001).
- Brown, J. W. & Braver, T. S. Learned predictions of error likelihood in the anterior cingulate cortex. *Science* **307**, 1118–1121 (2005).
- Kim, J., Ghim, J. W., Lee, J. H. & Jung, M. W. Neural correlates of interval timing in rodent prefrontal cortex. *J. Neurosci.* **33**, 13834–13847 (2013).
- Tallot, L. & Doyere, V. Neural encoding of time in the animal brain. *Neurosci. Biobehav. Rev.* **115**, 146–163 (2020).
- Janssen, P. & Shadlen, M. N. A representation of the hazard rate of elapsed time in macaque area LIP. *Nat. Neurosci.* **8**, 234–241 (2005).
- Coull, J. & Nobre, A. Dissociating explicit timing from temporal expectation with fMRI. *Curr. Opin. Neurobiol.* **18**, 137–144 (2008).
- Starkweather, C. K., Babayan, B. M., Uchida, N. & Gershman, S. J. Dopamine reward prediction errors reflect hidden-state inference across time. *Nat. Neurosci.* **20**, 581–589 (2017).
- O’Reilly, J. X. et al. Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proc. Natl Acad. Sci. USA* **110**, E3660–E3669 (2013).

39. Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* **84**, 279 (1977).
40. Lemon, R. N., Baker, S. N. & Kraskov, A. Classification of cortical neurons by spike shape and the identification of pyramidal neurons. *Cereb. Cortex* **31**, 5131–5138 (2021).
41. Schall, J. D. & Boucher, L. Executive control of gaze by the frontal lobes. *Cogn. Affect. Behav. Neurosci.* **7**, 396–412 (2007).
42. Ogasawara, T., Nejime, M., Takada, M. & Matsumoto, M. Primate nigrostriatal dopamine system regulates saccadic response inhibition. *Neuron* **100**, 1513–1526 e1514 (2018).
43. Nelson, M. J., Boucher, L., Logan, G. D., Palmeri, T. J. & Schall, J. D. Nonindependent and nonstationary response times in stopping and stepping saccade tasks. *Atten. Percept. Psychophys.* **72**, 1913–1929 (2010).
44. Zhang, K., Chen, C. D. & Monosov, I. E. Novelty, salience, and surprise timing are signaled by neurons in the basal forebrain. *Curr. Biol.* **29**, 134–142.e133 (2019).
45. Westerberg, J. A., Schall, M. S., Maier, A., Woodman, G. F. & Schall, J. D. Laminar microcircuitry of visual cortex producing attention-associated electric fields. *Elife* **11**, e72139 (2022).
46. Nunez, P. L. & Srinivasan, R. *Electric Fields of the Brain: The Neurophysics of EEG* (Oxford University Press, USA, 2006).
47. Huerta, M. F. & Kaas, J. H. Supplementary eye field as defined by intracortical microstimulation: connections in macaques. *J. Comp. Neurol.* **293**, 299–330 (1990).
48. Parthasarathy, H. B., Schall, J. D. & Graybiel, A. M. Distributed but convergent ordering of corticostriatal projections: analysis of the frontal eye field and the supplementary eye field in the macaque monkey. *J. Neurosci.* **12**, 4468–4488 (1992).
49. Shook, B. L., Schlag-Rey, M. & Schlag, J. Primate supplementary eye field: I. Comparative aspects of mesencephalic and pontine connections. *J. Comp. Neurol.* **301**, 618–642 (1990).
50. Paré, M. & Hanes, D. P. Controlled movement processing: Superior colliculus activity associated with countermanded saccades. *J. Neurosci.* **23**, 6480–6489 (2003).
51. Lo, C. C., Boucher, L., Paré, M., Schall, J. D. & Wang, X. J. Proactive inhibitory control and attractor dynamics in countermanding action: a spiking neural circuit model. *J. Neurosci.* **29**, 9059–9071 (2009).
52. Bromberg-Martin, E. S., Matsumoto, M. & Hikosaka, O. Dopamine in motivational control: rewarding, aversive, and alerting. *Neuron* **68**, 815–834 (2010).
53. Errington, S. P. & Schall, J. D. Express saccades during a countermanding task. *J. Neurophysiol.* **124**, 484–496 (2020).
54. Egger, S. W., Le, N. M. & Jazayeri, M. A neural circuit model for human sensorimotor timing. *Nat. Commun.* **11**, 3933 (2020).
55. Schultz, W., Carelli, R. M. & Wightman, R. M. Phasic dopamine signals: from subjective reward value to formal economic utility. *Curr. Opin. Behav. Sci.* **5**, 147–154 (2015).
56. So, N. Y. & Stuphorn, V. Supplementary eye field encodes option and action value for saccades with variable reward. *J. Neurophysiol.* **104**, 2634–2653 (2010).
57. Gibbon, J., Church, R. M. & Meck, W. H. Scalar timing in memory. *Ann. NY Acad. Sci.* **423**, 52–77 (1984).
58. Polti, I., Martin, B. & van Wassenhove, V. The effect of attention and working memory on the estimation of elapsed time. *Sci. Rep.* **8**, 6690 (2018).
59. Herrera, B., Sajad, A., Woodman, G. F., Schall, J. D. & Riera, J. J. A minimal biophysical model of neocortical pyramidal cells: implications for frontal cortex microcircuitry and field potential generation. *J. Neurosci.* **40**, 8513–8529 (2020).
60. Atsma, J., Majij, F., Corneil, B. D. & Medendorp, W. P. No perisaccadic mislocalization with abruptly cancelled saccades. *J. Neurosci.* **34**, 5497–5504 (2014).
61. Botvinick, M. M. Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cogn. Affect. Behav. Neurosci.* **7**, 356–366 (2007).
62. Williams, S. M. & Goldman-Rakic, P. S. Widespread origin of the primate mesofrontal dopamine system. *Cereb. Cortex* **8**, 321–345 (1998).
63. Grace, A. A. & Bunney, B. S. Nigral dopamine neurons: Intracellular recording and identification with L-dopa injection and histo-fluorescence. *Science* **210**, 654–656 (1980).
64. Redgrave, P. & Gurney, K. The short-latency dopamine signal: a role in discovering novel actions? *Nat. Rev. Neurosci.* **7**, 967–975 (2006).
65. Buhusi, C. V. & Meck, W. H. What makes us tick? Functional and neural mechanisms of interval timing. *Nat. Rev. Neurosci.* **6**, 755–765 (2005).
66. Emmons, E. et al. Temporal learning among prefrontal and striatal ensembles. *Cereb. Cortex Commun.* **1**, tgaa058 (2020).
67. Griggs, W. S. et al. Flexible and stable value coding areas in caudate head and tail receive anatomically distinct cortical and subcortical inputs. *Front. Neuroanat.* **11**, 106 (2017).
68. Saint-Cyr, J. A., Ungerleider, L. G. & Desimone, R. Organization of visual cortical inputs to the striatum and subsequent outputs to the pallido-nigral complex in the monkey. *J. Comp. Neurol.* **298**, 129–156 (1990).
69. Medalla, M. & Barbas, H. Synapses with inhibitory neurons differentiate anterior cingulate from dorsolateral prefrontal pathways associated with cognitive control. *Neuron* **61**, 609–620 (2009).
70. Gidon, A. & Segev, I. Principles governing the operation of synaptic inhibition in dendrites. *Neuron* **75**, 330–341 (2012).
71. Fiorillo, C. D., Newsome, W. T. & Schultz, W. The temporal precision of reward prediction in dopamine neurons. *Nat. Neurosci.* **11**, 966–973 (2008).
72. Lidow, M. S., Goldman-Rakic, P. S., Gallager, D. W. & Rakic, P. Distribution of dopaminergic receptors in the primate cerebral cortex: quantitative autoradiographic analysis using [³H]raclopride, [³H]spiperone and [³H]SCH23390. *Neuroscience* **40**, 657–671 (1991).
73. Lim, S. & Goldman, M. S. Balanced cortical microcircuitry for maintaining information in working memory. *Nat. Neurosci.* **16**, 1306–1314 (2013).
74. Wang, X. J., Tegner, J., Constantinidis, C. & Goldman-Rakic, P. S. Division of labor among distinct subtypes of inhibitory neurons in a cortical microcircuit of working memory. *Proc. Natl Acad. Sci. USA* **101**, 1368–1373 (2004).
75. Tremblay, R., Lee, S. & Rudy, B. GABAergic interneurons in the neocortex: from cellular properties to circuits. *Neuron* **91**, 260–292 (2016).
76. Huster, R. J., Messel, M. S., Thunberg, C. & Raud, L. The P300 as marker of inhibitory control—fact or fiction? *Cortex* **132**, 334–348 (2020).
77. Huster, R. J., Enriquez-Geppert, S., Lavallee, C. F., Falkenstein, M. & Herrmann, C. S. Electroencephalography of response inhibition tasks: functional networks and cognitive contributions. *Int. J. Psychophysiol.* **87**, 217–233 (2013).
78. Smith, J. L., Smith, E. A., Provost, A. L. & Heathcote, A. Sequence effects support the conflict theory of N2 and P3 in the Go/NoGo task. *Int. J. Psychophysiol.* **75**, 217–226 (2010).
79. Buzsáki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.* **13**, 407–420 (2012).
80. Rigotti, M. et al. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
81. Ramakrishnan, A. et al. Cortical neurons multiplex reward-related signals along with sensory and motor information. *Proc. Natl Acad. Sci. USA* **114**, E4841–E4850 (2017).

82. Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F. & Ostojic, S. The role of population structure in computations through neural dynamics. *Nat. Neurosci.* **25**, 783–794 (2022).
83. Thakkar, K. N., Schall, J. D., Boucher, L., Logan, G. D. & Park, S. Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biol. Psychiatry* **69**, 55–62 (2011).
84. Godlove, D. C., Garr, A. K., Woodman, G. F. & Schall, J. D. Measurement of the extraocular spike potential during saccade countermanding. *J. Neurophysiol.* **106**, 104–114 (2011).
85. Schlag, J. & Schlag-Rey, M. Evidence for a supplementary eye field. *J. Neurophysiol.* **57**, 179–200 (1987).
86. Schall, J. D. Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *J. Neurophysiol.* **66**, 559–579 (1991).
87. Maier, A., Adams, G. K., Aura, C. & Leopold, D. A. Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation. *Front. Syst. Neurosci.* **4**, <https://doi.org/10.3389/fnsys.2010.00031> (2010).
88. Matzke, D., Love, J. & Heathcote, A. A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behav. Res. Methods* **49**, 267–281 (2017).
89. Raftery, A. E. Bayesian model selection in social research. *Sociol. Methodol.* **25**, 111–163 (1995).
90. Kass, R. E. & Raftery, A. E. Bayes factors. *J. Am. Stat. Assoc.* **90**, 773–795 (1995).

Acknowledgements

The authors thank D. Godlove for data collection, G. Luppino, M. Matsumoto, N. Palomero-Gallagher, and L. Rapan for sharing data; J. Elsey, M. Feurtado, M. Maddox, S. Motorny, J. Parker, D. Richardson, M. Schall, C.R. Subraveti, L. Toy, B. Williams, and R. Williams for animal care and other technical assistance; and Z. Fu, M. Matsumoto, P. Redgrave, U. Rutishauser, E. Sigworth, A. Tomarken, and G. Woodman for helpful discussions. Imaging data were collected at the Vanderbilt Institute of Imaging Science. This work was supported by R01-MH55806 (J.D.S.), R01-EY019882 (G. Woodman), P30-EY08126 (J.D.S. and D. Calkins), NSERC RGPIN-2022-04592 (J.D.S.), CIHR Post-Doctoral Fellowship (A.S.), by Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience, and by the Vision Science to Application (VISTA) program at York University.

Author contributions

Experimental design, J.D.S. Data collection, J.D.S. Data analysis, A.S. and S.E. Interpretation and preparation of the manuscript, A.S., S.E., and J.D.S.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-022-33942-1>.

Correspondence and requests for materials should be addressed to Jeffrey D. Schall.

Peer review information *Nature Communications* thanks Hansem Sohn and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022

Functional architecture of executive control and associated event-related potentials in macaques

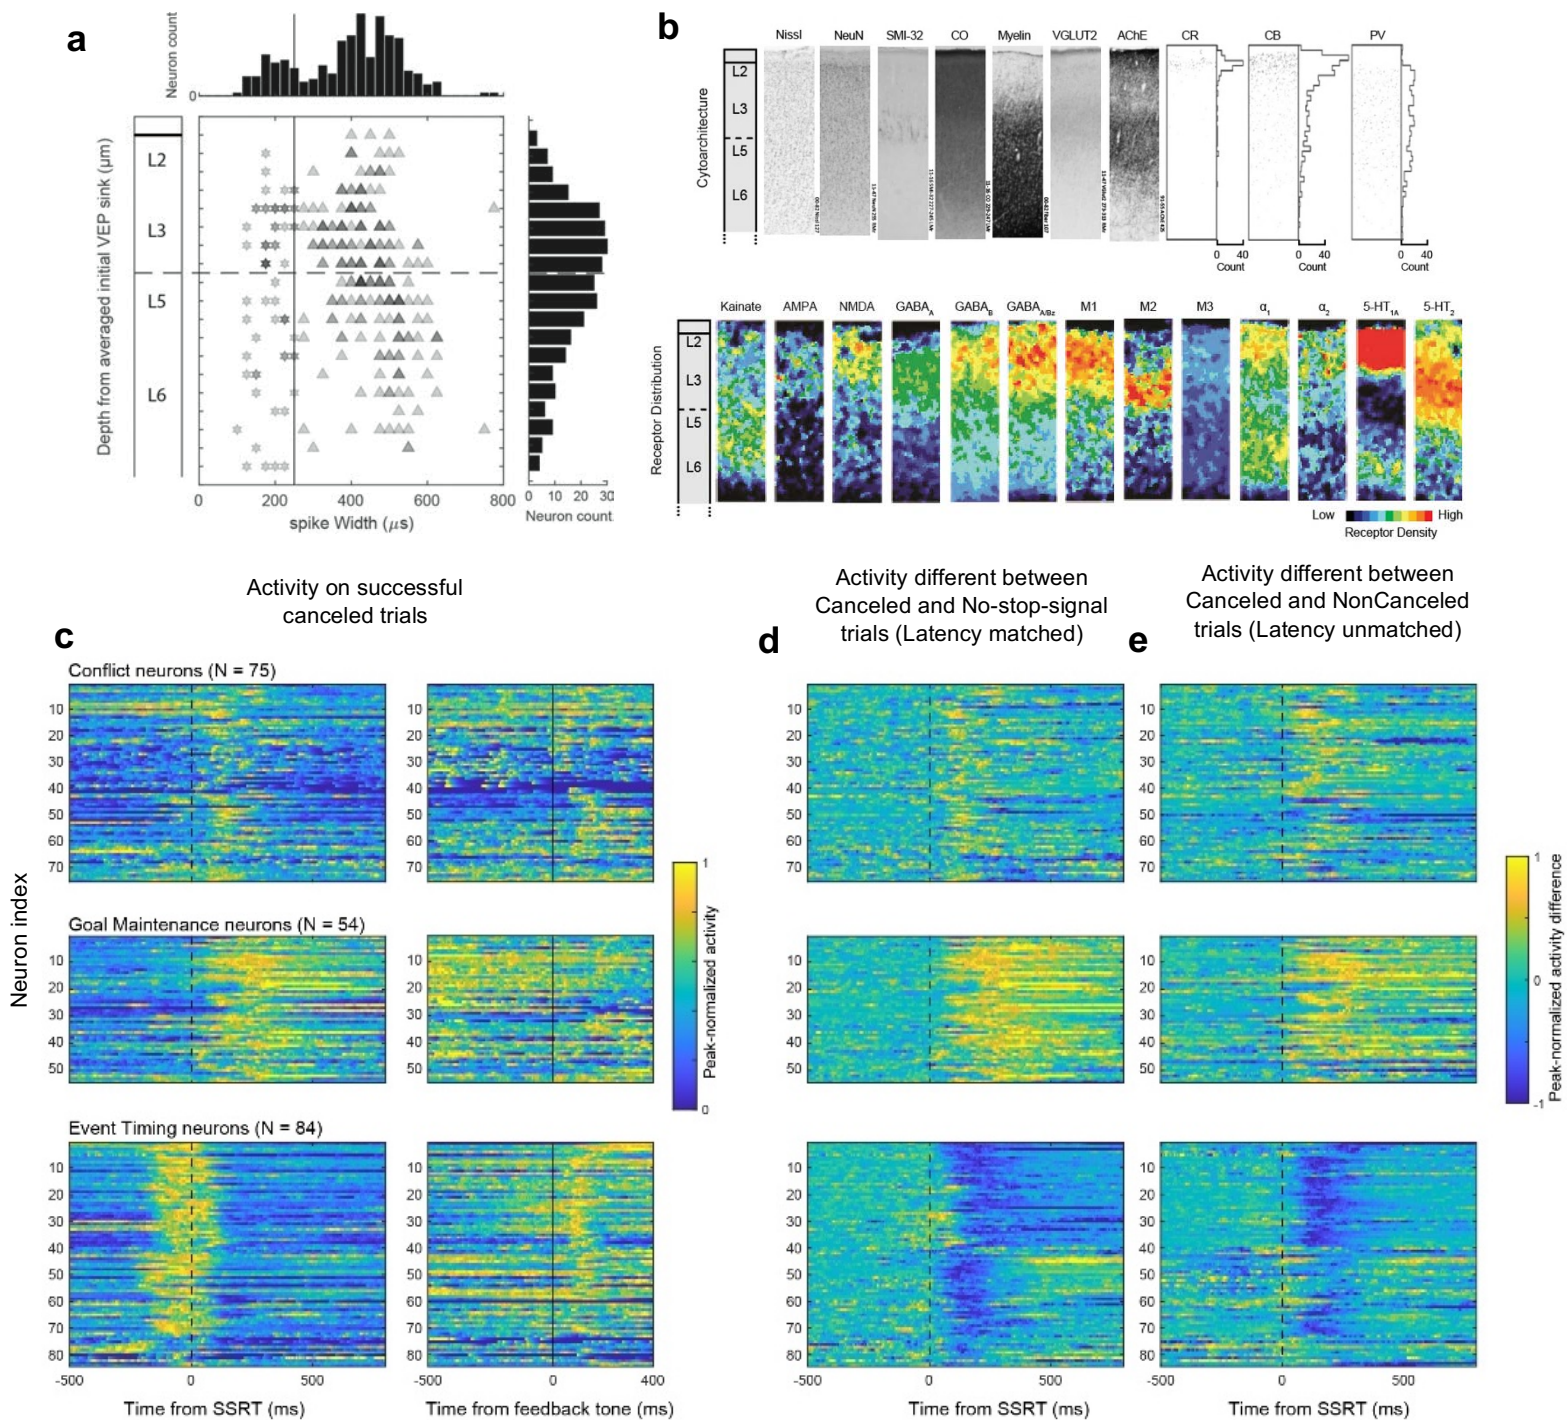
Amirsaman Sajad^{1*}, Steven P. Errington^{1*}, & Jeffrey D. Schall^{1,2}

¹ Department of Psychology, Vanderbilt Vision Research Center, Center for Integrative & Cognitive Neuroscience, Vanderbilt University, Nashville, TN

² Centre for Vision Research, Vision Science to Application, Department of Biology, York University, Toronto, ON

* Authors contributed equally to this work.

Supplementary Figure 1 | SEF laminar structure and neuron classification.



a, Distributions of depths of units and spike widths sampled. Horizontal histogram shows the spike widths of all neurons sampled ($n = 575$), which exhibited bimodal values. The vertical line marks the $250 \mu\text{s}$ separation criterion used. Scatterplot shows variation of spike width across depth for neurons sampled in perpendicular penetrations ($n = 293$). Narrow spiking neurons are indicated by stars and broad spiking neurons by triangles. The number of neurons at each time-depth indicated by gray scale (darker indicating higher count). The width of the spikes narrower than $250 \mu\text{s}$ does not vary with depth, and the incidence of encountering narrow spikes parallels the density of parvalbumin (PV) neurons¹. The width of spikes wider than $250 \mu\text{s}$ increases from L3 to L6, which parallels the size of pyramidal neurons. Also, the

incidence of isolated neurons decreases with depth, which parallels the density of pyramidal neurons. Vertical histogram shows the depths of all neurons sampled during the saccade countermanding task from the sessions with penetrations oriented perpendicular to the cortical layers. The color scale used for the time-depth plots in Fig. 2, 3, 5, and 5 indicates the number of neurons of each class sampled relative to the entire sampling distribution. Adapted from ².

b, Laminar structure of SEF. Top row, from left to right are shown sections through SEF stained for Nissl, NeuN, nonphosphorylated neurofilament H (SMI-32), cytochrome oxidase (CO), Gallyas myelin, vesicular glutamate transporter 2 (VGLUT2), acetylcholinesterase (AChE), calretinin (CR), calbindin (CB), and PV. Counts of CR, CB, and PV-stained neurons are adapted from ¹. In SEF, L1 is ~200 μm thick with some CR and CB but no PV neurons, and weak staining for CO, myelin, and VGLUT2. L2 is ~300 μm thick with dense stellate and few pyramidal neurons, the highest density of CR and CB neurons, but no SMI-32, stronger CO staining, fascicular myelin fibers, slightly stronger VGLUT2, and modest AChE. L3 is ~700 μm thick with a superficial sublayer with smaller pyramidal neurons and weak SMI-32 staining, very few CR, and modest densities of CB and PV neurons, stronger CO staining and denser myelin, VGLUT2, and AChE. A deeper sublayer is characterized by larger pyramidal neurons, pronounced SMI-32 staining, vanishingly few CR, less dense CB and modestly dense PV neurons, weaker CO, denser myelin and VGLUT2, and denser AChE. No granular L4 is evident in SEF. L5 is ~300 μm thick with large pyramidal neurons but inconsistent SMI-32 staining, no CR and fewer CB but modest density of PV neurons, lighter CO, denser myelin, lighter VGLUT2, and diminishing AChE staining. L6 is ~700 μm thick with smaller pyramidal neurons, light SMI-32, no CR, vanishingly few CB and low density of PV neurons, with lighter CO, still denser myelin, lighter VGLUT2 and sparse AChE staining.

Lower row, sections processed for receptor autoradiography and color coded to visualize the laminar densities of various receptors adapted from ³. The color scale maps to densities in fmol/mg protein with (minimum, maximum) levels as indicated for the following receptors: kainate (20, 1200), AMPA (20, 750), NMDA (200, 2500), GABA_A (100, 3200), GABA_B (150, 3500), GABA_{A/Bz} (200, 3500), muscarinic M₁ (100, 1200), muscarinic M₂ (10, 350), muscarinic M₃ (100, 1200), adrenergic α_1 (50, 700), adrenergic α_2 (20, 800), serotonergic 5-HT_{1A} (20, 700), and serotonergic 5-HT₂ (100, 500). Receptor densities reveal pronounced differences between L2/3 and L5/6 and other differences distinguishing L2 from L3 and L5 from L6. The variation in laminar structure can guide the investigation and interpretation of functional architecture of SEF.

c-e, Heat-maps representing normalized spike-density function for all neurons reported in this paper. These neuron classes emerged from our neuron classification pipeline (**Fig. 1c**). In a sample of 575 neurons, we took the following steps to identify distinct types of neurons. First, we identified 271 with significant changes in spiking rate following presentation of the stop-signal until 200 ms after stop-signal reaction time (SSRT) on successfully canceled trials. Second, the spike density function on canceled and latency-matched no-stop-signal trials in intervals 0–100 ms before and 0–200 ms after SSRT were submitted to an unsupervised consensus clustering pipeline ⁴. 5 clusters were identified. Neurons were distinguished by a relative increase or decrease in discharge rate on canceled compared to latency-matched no-stop-signal trials. Cluster 1 included 146 neurons; cluster 2, 65 neurons; cluster 3, 28 neurons; and cluster 4, 16 neurons, and cluster 5, 16 neurons. Clusters 1 and 2 formed the majority of neurons and were analyzed further. The clustering results were verified by manual curation, resulting in some cluster 3, 4, and 5 neurons being included in clusters 1 and 2 and removal of neurons with poor signal quality. Therefore, this clustering procedure justified the consideration of two general types of modulation: facilitation (i.e., relatively higher activity on canceled trials) and suppression (i.e., relatively lower activity on canceled trials). Within each cluster we observed heterogeneity in the latency, duration, and magnitude of modulation after SSRT and around the time of feedback tone. Tone modulation was indexed by the contrast (difference divided by sum) between spike counts 200 ms before and 200 ms after the feedback tone. Each neuron was located in a 3D space based on these values. K-means clustering of the normalized values distinguished two clusters among facilitated neurons primarily based on modulation duration. Arriving at $k = 2$ clusters was justified by the Elbow method (analysis of within sum of distance to center) and the Silhouette score (inset panel, **Fig 1c**). K-means clustering did not distinguish among the suppressed neurons (not shown). Thus, this rigorous classification approach resulted in three neuron classes based only on the pattern of modulation around the time of SSRT with no presumptions about their functional role. The distinct patterns of modulation were reinforced by qualitative examination of clear differences among population spike density functions (**Fig. 3, 4, 5**). The functional roles and further distinctions based on different patterns of modulation around the time of tone are described in the main text.

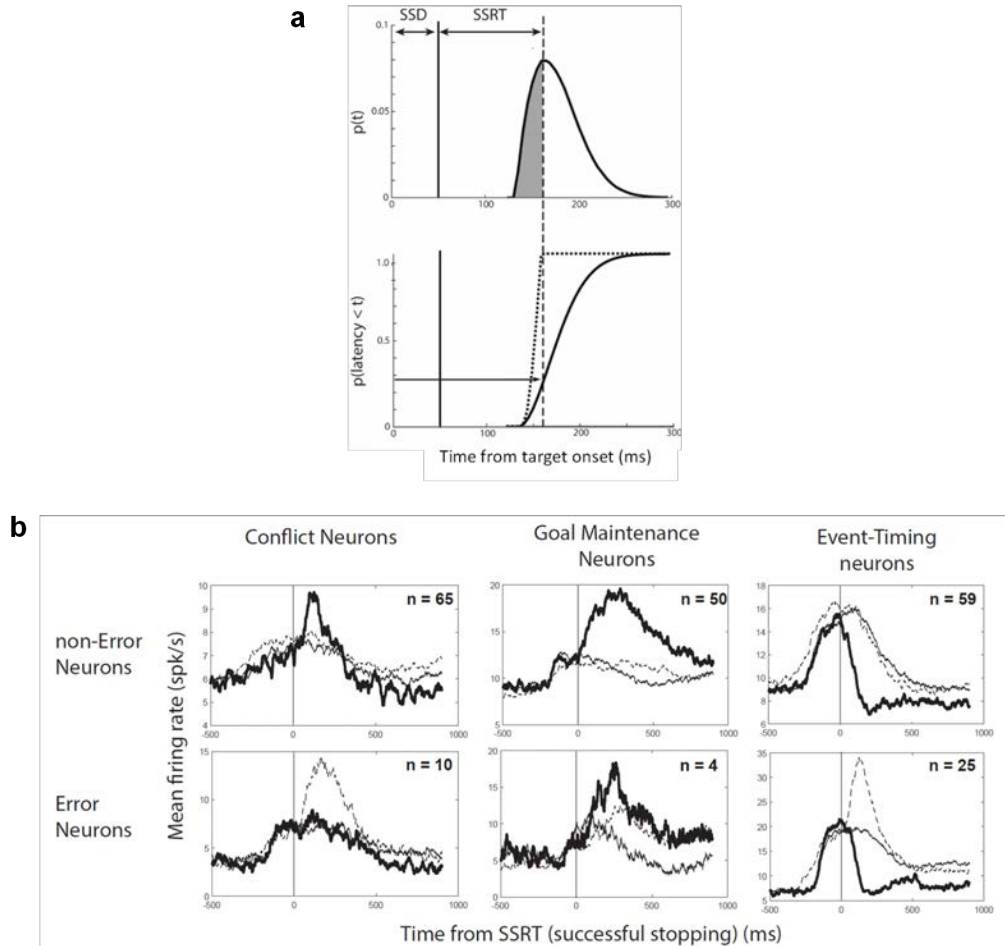
c, Heat-maps representing normalized spike-density function on successful canceled trials for all neurons reported in this paper. The spike-density function for each neuron is normalized to its peak activity ± 300 ms around SSRT. The left panel plots activity aligned on stop-signal reaction time (SSRT) (dashed black line). The right panel plots activity aligned on the feedback tone (solid line). Each row corresponds to activity of one neuron with higher discharge rates in hotter, and lower discharge rates in cooler colors. The time course of activation that distinguished conflict, event-timing, and goal maintenance neurons is evident in the three pairs of panels. Similarities and differences between these SEF neurons and

DA neurons in SNpc and neurons in striatum sampled during saccade countermanding is afforded by comparison with Fig. 2 and 5 of ⁵.

d, Heat-maps representing the difference in spike-density function on successful canceled trials relative to latency-matched no-stop-signal trials, aligned on SSRT, selected for the early stop-signal delay (SSD). The difference function for each neuron is normalized to the maximum deviation. Hot colors represent facilitation and cold colors represent suppression on canceled relative to no-stop-signal trials. The patterns of activation that distinguish the three major categories of neurons are more evident across the three panels. Similarities and differences between these SEF neurons and DA neurons in SNpc and neurons in striatum sampled during saccade countermanding is afforded by comparison with Fig. 2 and 5 of ⁵.

e, Heat-maps representing the difference in spike-density function on successful canceled trials relative to noncanceled trials (despite the inherent difference in underlying RT), aligned on SSRT, selected for an intermediate SSD with enough trials of both types. The patterns of activation that distinguish the three major categories of neurons are more evident across the three panels. Similarities and differences between these SEF neurons and DA neurons in SNpc and neurons in striatum sampled during saccade countermanding is afforded by comparison with Fig. 2 and 5 of ⁵.

Supplementary Figure 2 | Rationale for analysis of neural activity.

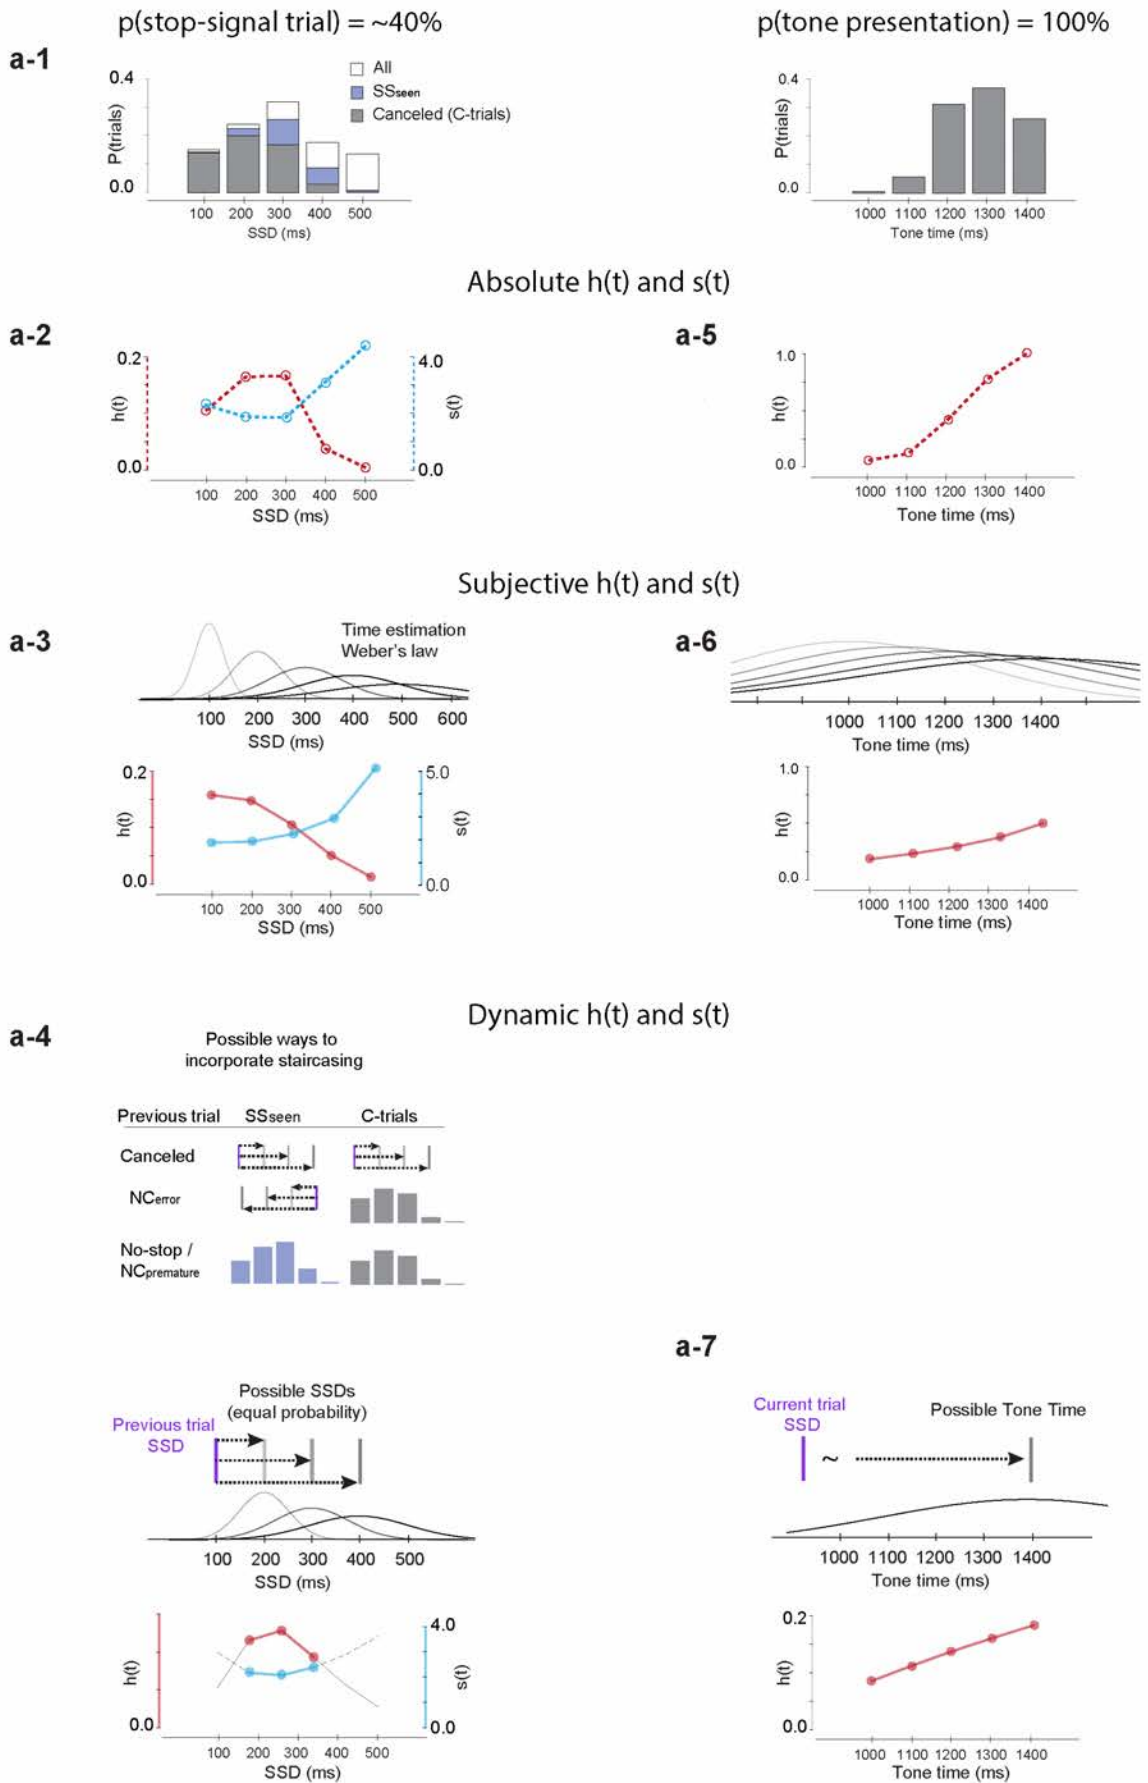


a, The primary analysis was based on the principles of the Logan race model, which have been used extensively⁶⁻⁸. Diagram of density (top) and cumulative (bottom) distributions of RT on trials with no stop-signal portrays the application of the race model for a given SSD. Because performance is the outcome of a race, at a given SSD, the finish time of the STOP process (SSRT) is the RT of trials with no stop-signal at which the fraction of the all trials equals $p(\text{NC} | \text{SSD})$ (shaded in density function and dotted in cumulative distribution). The race model entails that the no-stop-signal trials with $\text{RT} > \text{SSD} + \text{SSRT}$ are so slow that they would have been canceled if the stop-signal had appeared. Therefore, we compared neural activity on canceled trials when response inhibition happened with activity on no-stop-signal trials when no response inhibition happened but with RT long enough that if the stop-signal had been presented, the saccade would have been canceled. By latency-matching canceled trials to no-stop-signal trials with $\text{RT} > \text{SSD} + \text{SSRT}$, the dynamics of the neural processes governing saccade initiation are equated. Likewise, noncanceled trials can be latency matched to no-stop-signal trials with $\text{RT} < \text{SSD} + \text{SSRT}$.

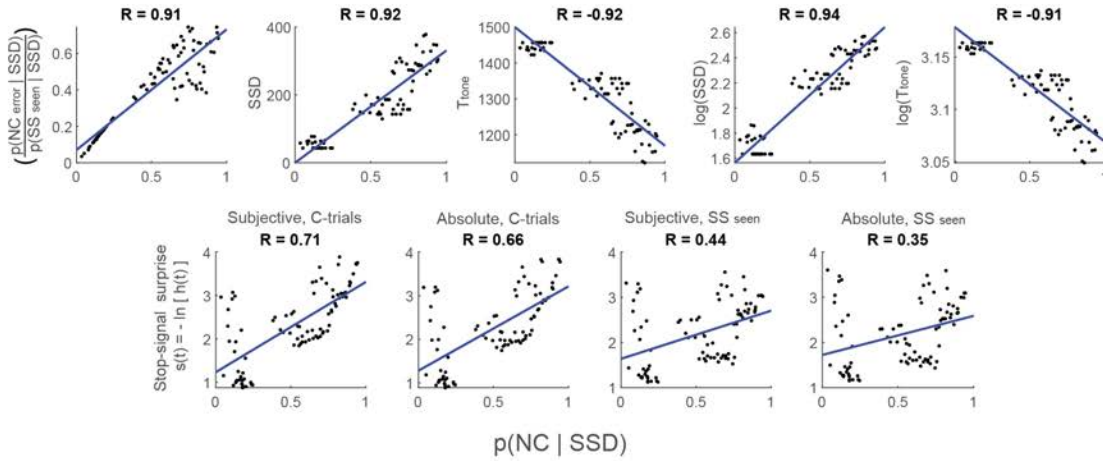
b, Activity on no-stop-signal (thin solid line), canceled (thick solid line), and error noncanceled (thin dashed) trials for neurons classes with post-SSRT modulation. Modulation related to successful stopping was determined based on the difference in activity between canceled and latency-matched no-stop-signal trials. Noncanceled trials were not used for analyzing modulation related to response inhibition for two reasons. First, failures of response inhibition in this task happen when RT is too short, which is the opposite of how successful inhibition occurs. Second, noncancelled trials elicit additional processes related to error detection. To illustrate, we have divided each neuron class into those without (top) and with (bottom) error-related modulation. The minority of Conflict (10/75), Goal Maintenance (4/54) and Event timing neurons (25/84) produced an error signal. Also, post-saccadic modulation in noncanceled and no-stop-signal trials is similar for neurons without but not with error-related modulation.

Supplementary Figure 3 | Time parameters and their relationships used to test neural modulation.

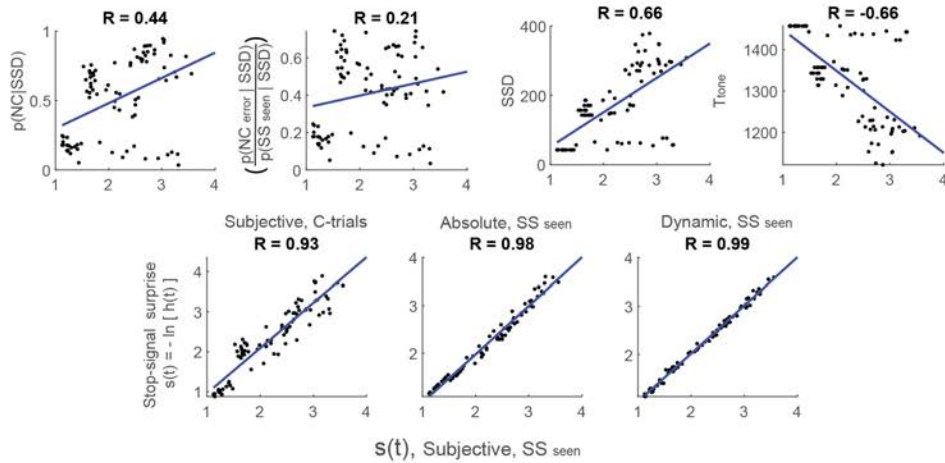
a




b Comparison between Conflict model and other models tested for post-SSRT facilitation, N2, and P3




Comparison between S(t) Subjective, SSseen model and other models tested for post-SSRT facilitation, N2, and P3



Comparison between log(SSD) model and other models tested for pre-SSRT activity



Comparison between log(Ttone) model and other models tested for pre-Tone activity



a, Patterns of neural spiking were analyzed in relation to multiple hazard rate (expectancy), and surprise measures. Stop-signals were scheduled to be presented at a range of times (**a-1**, left, white bars). At most times the stop-signal was seen (SS_{seen} , blue bars), which included canceled trials (gray bars) and explicit error noncanceled trials with $RT > SSD$ (NC_{error}). By design, on canceled trials, the interval between presentation of the visual target and of the feedback tone was fixed at 1,500 ms. Thus, the interval from stop-signal to success tone (T_{tone}) was just $T_{\text{tone}} = 1500 - SSD$ (**a-1**, right). Monkeys can learn temporal regularities to form moment-by-moment expectations for task events. Neural responses can encode this expectation prospectively, before the event, or surprise about violation of these expectations, retrospectively, after the event. Expectation was operationalized by hazard rate. Surprise was calculated as the Shannon Information, derived from the hazard rate. Guided by previous research⁹⁻¹³, three kinds of hazard rate were calculated based on different assumptions: (i) the absolute hazard rate, (ii) subjective hazard rate, and (iii) dynamic hazard rate.

Absolute hazard rate was derived from the distribution of time (t) values as $h(t) = f(t) / [1-F(t)]$, where $f(t)$ is the density distribution and $F(t)$ is the cumulative distribution. Monkeys can gain knowledge only from SSD that are experienced, so we used only trials in which stop-signal was seen before the response (SS_{seen} trials). We also considered the possibility that the knowledge from task structure is gained in a trial-specific manner, only restricted to canceled trials (C-trials). Now, because stop-signal presentation was probabilistic with ~40% likelihood, the hazard function was calculated with the conditional probability of stop-signal appearance for each session. This conditional probability accounted for the fact that on a given trial, because of the uncertainty associated with the current trial being a stop-signal trial, if time passes and no stop-signal is observed, the belief about the current trial being a stop-signal trial drops, reducing the expectation for stop-signal appearance. Consequently, conditional hazard rate does not approach 1.0 and is not monotonic. Because our approach relied on mean responses across many trials having different patterns of preceding trial history within each SSD bin, we assumed variations in the inferred state only change due to passage of time within a trial. A conditional probability was not necessary for the hazard rate of T_{tone} because the feedback tone occurred with 100% certainty. Representative plots of the absolute hazard rates of SSD and T_{tone} are shown in **a-2** and **a-5** (red) and the associated surprise for SSD is shown in **a-2** (blue). The quantities shown here are based on C-trials, corresponding to $h(t)_{\text{Absolute, C-trials}}$ and $s(t)_{\text{Absolute, C-trials}}$. We did not analyze the period after the tone, therefore surprise for tone was not considered.

Subjective hazard rate was derived from time (t) values that incorporated the imprecision in time perception. The estimation of elapsed time has increasing uncertainty over longer intervals¹⁴. To account for this imprecision, $f(t)$ was convolved with a Gaussian kernel with standard distribution proportional to the elapsed time:

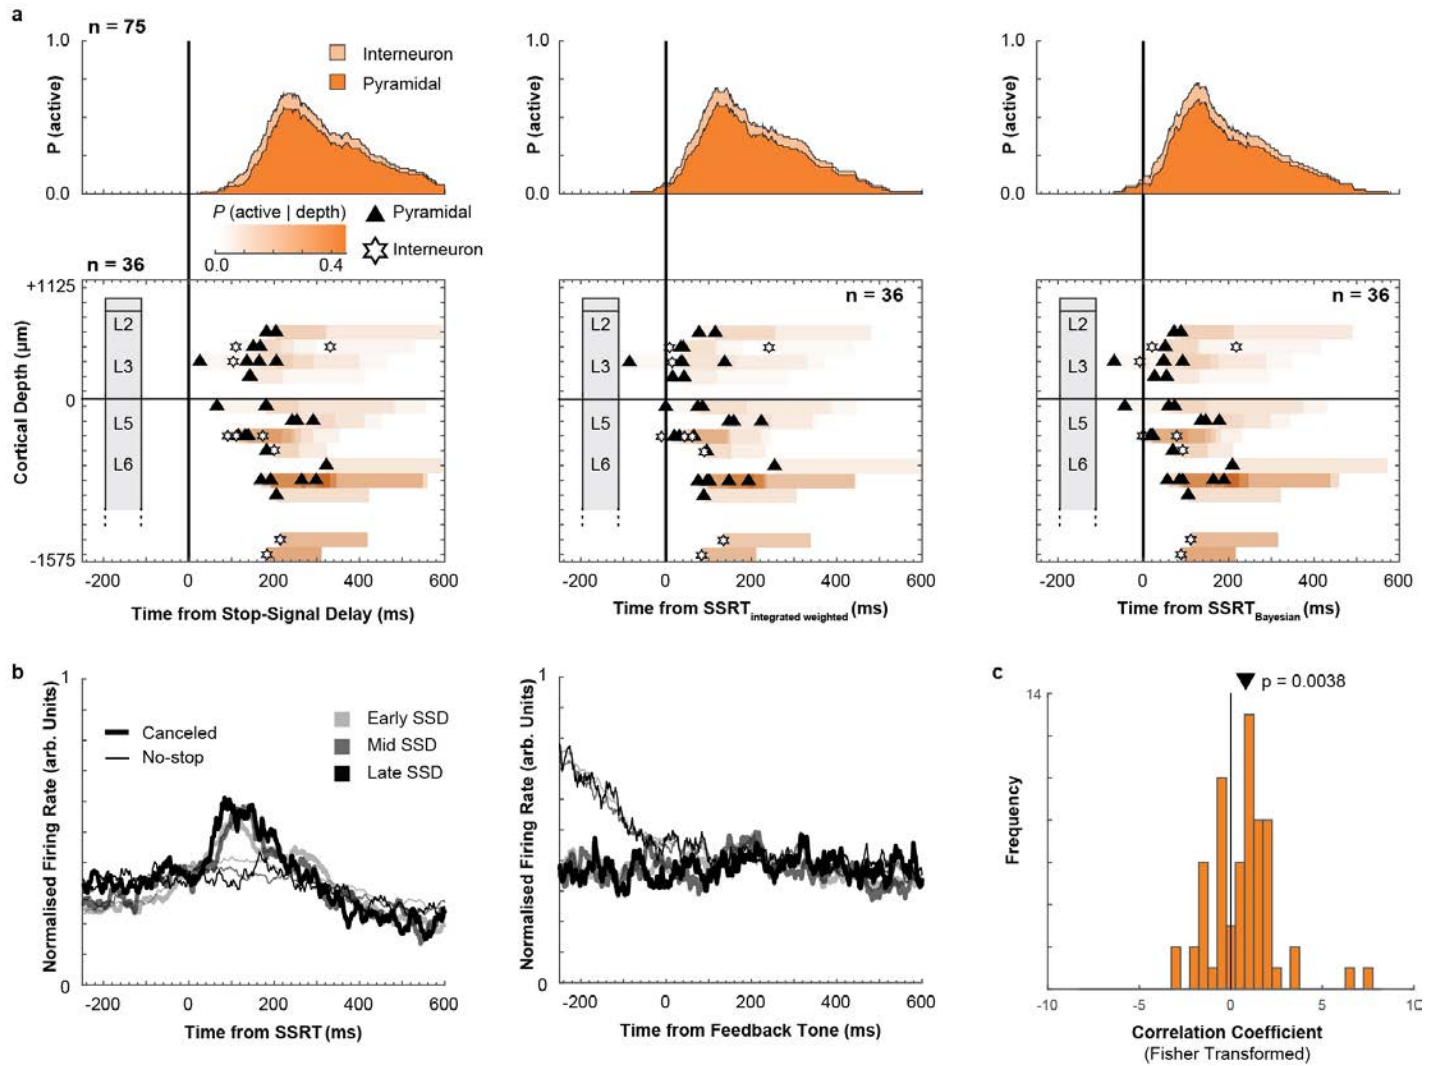
$$g(t) = \frac{1}{\sigma t \sqrt{2\pi}} \int_{-\infty}^{\infty} f(x) e^{-\frac{(t-x)^2}{2\sigma^2 t}} dx$$

The coefficient of variation, σ , was set to 0.26 based on previous research^{11, 14} (**a-3**, **a-6**, top). Representative plots of the subjective hazard rates of SSD and T_{tone} are shown in **a-3** and **a-6** (red) with associated surprise for SSD in **a-3** (blue). The quantities shown here are based on C-trials, thus corresponding to $h(t)_{\text{Subjective, C-trials}}$ and $s(t)_{\text{Subjective, C-trials}}$.

Dynamic hazard rate was derived from time values incorporating both imprecision in time estimation and the restricted range of possible SSD or T_{tone} values based on knowledge of regularities in the preceding intervals experienced¹². This was important because to ensure that monkeys failed to cancel on ~50% of all stop-signal trials, SSD was adapted dynamically in a staircase algorithm. After noncanceled trials, SSD was decreased. After canceled trials, SSD was increased. To discourage monkeys from exploiting the stair-casing algorithm, SSD was adjusted in steps of 1, 2, or 3 intervals randomly selected with uniform 1/3 probability. Nevertheless, we explored the possibility that monkeys can incorporate this knowledge to predict the timing of the upcoming SSD. A prediction of SSD would influence estimation of T_{tone} because $T_{\text{tone}} = 1500 - SSD$. The left panel of **a-4** portrays how the preceding trial restricted possible SSD values in the current trial. Separate calculations were done with C-trials and SS_{seen} trials. For $h(t)_{\text{Dynamic, C-trials}}$ stair-casing knowledge was only incorporated following canceled trials. For $h(t)_{\text{Dynamic, } SS_{\text{seen}}}$ stair-casing knowledge was incorporated for both canceled trials and NC_{error} trials. For all other trials, absolute $f(t)$ was used. Representative plots of the subjective hazard rates of SSD and T_{tone} are shown in **a-4** and **a-7** (red, $h(t)_{\text{Dynamic, C-trials}}$) with associated surprise for SSD (blue, $s(t)_{\text{Dynamic, C-trials}}$). The gray traces in lower panel of **a-4** corresponds to the calculated $h(t)$ and $s(t)$, and the colored portion of the trace corresponds to possible values at the SSDs restricted by the staircasing algorithm.

b, Some behavioral and task measures were correlated, but random and systematic variations within and across sessions enabled differentiation. The relationship of neural activity before SSRT and before the tone to each of the behavioral and task measures was compared in a set of mixed-effects models using Bayesian Information Criteria (BIC)¹⁵. Top two rows of panels plot numerical relationships between indicated measures (y-axis) and $p(NC | SSD)$. Four of the six surprise models are shown. Third and fourth rows plot relationships between indicated measures and surprise $s(t)_{\text{subjective, } SS_{\text{seen}}}$. Fifth row plots relationships of SSD and four of six stop-signal hazard rate measures to $\log(SSD)$. The bottom row plots the relationships of T_{tone} and two of three hazard rate measures to $\log(T_{\text{tone}})$.

Supplementary Figure 4 | Conflict Neurons.

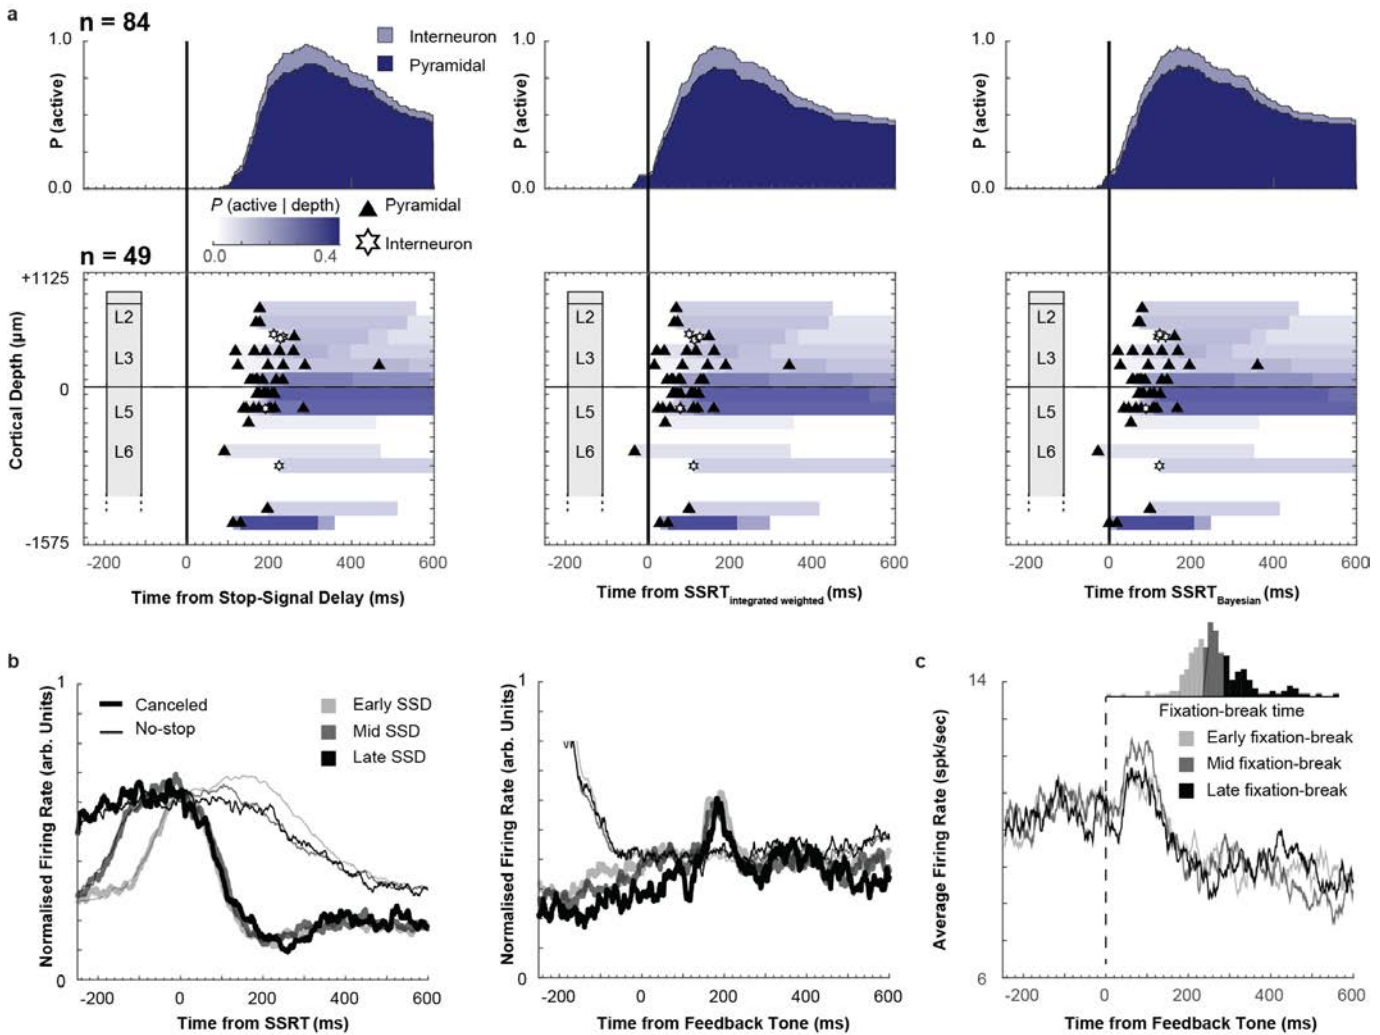


a, Recruitment and time-depth plots aligned on presentation of the stop-signal (left), SSRT calculated by the method of integration (middle, same as **Fig 3**), and SSRT calculated by the Bayesian Estimation of Ex-Gaussian Stop-Signal reaction time (BEESTS) (right). The patterns supporting the conclusions do not vary with the method of alignment. This neural signal nearly exclusively follows SSRT.

b, Normalized population spike-density functions aligned on SSRT (left) and feedback time (right) for canceled (thick) and latency-matched no-stop-signal trials (thin) subsampled from trials with low (lightest), middle (intermediate), and high (darkest) $p(\text{NC} | \text{SSD})$. The scaling of the modulation with SSD and $p(\text{NC} | \text{SSD})$ is evident across the sample. Note that the activity of neurons aligned on tone, on no-stop-signal trials during the pre-tone interval can be larger than that on canceled trials because of post-saccadic activity bleeding into this time period.

c, Distribution of correlation coefficients of post-SSRT modulation as a function of $p(\text{NC} | \text{SSD})$ across neurons. Correlation coefficients were transformed into a normal distribution using Fisher-transformation. The distribution was shifted significantly positively (Two-tailed Wilcoxon test, $p = 0.004$). The black triangle indicates the mean of the distribution. This plotting format clarifies the correspondence between the modulation patterns of SEF and of brainstem dopamine neurons of monkeys performing saccade countermanding tasks ⁵.

Supplementary Figure 5 | Event-Timing Neurons.

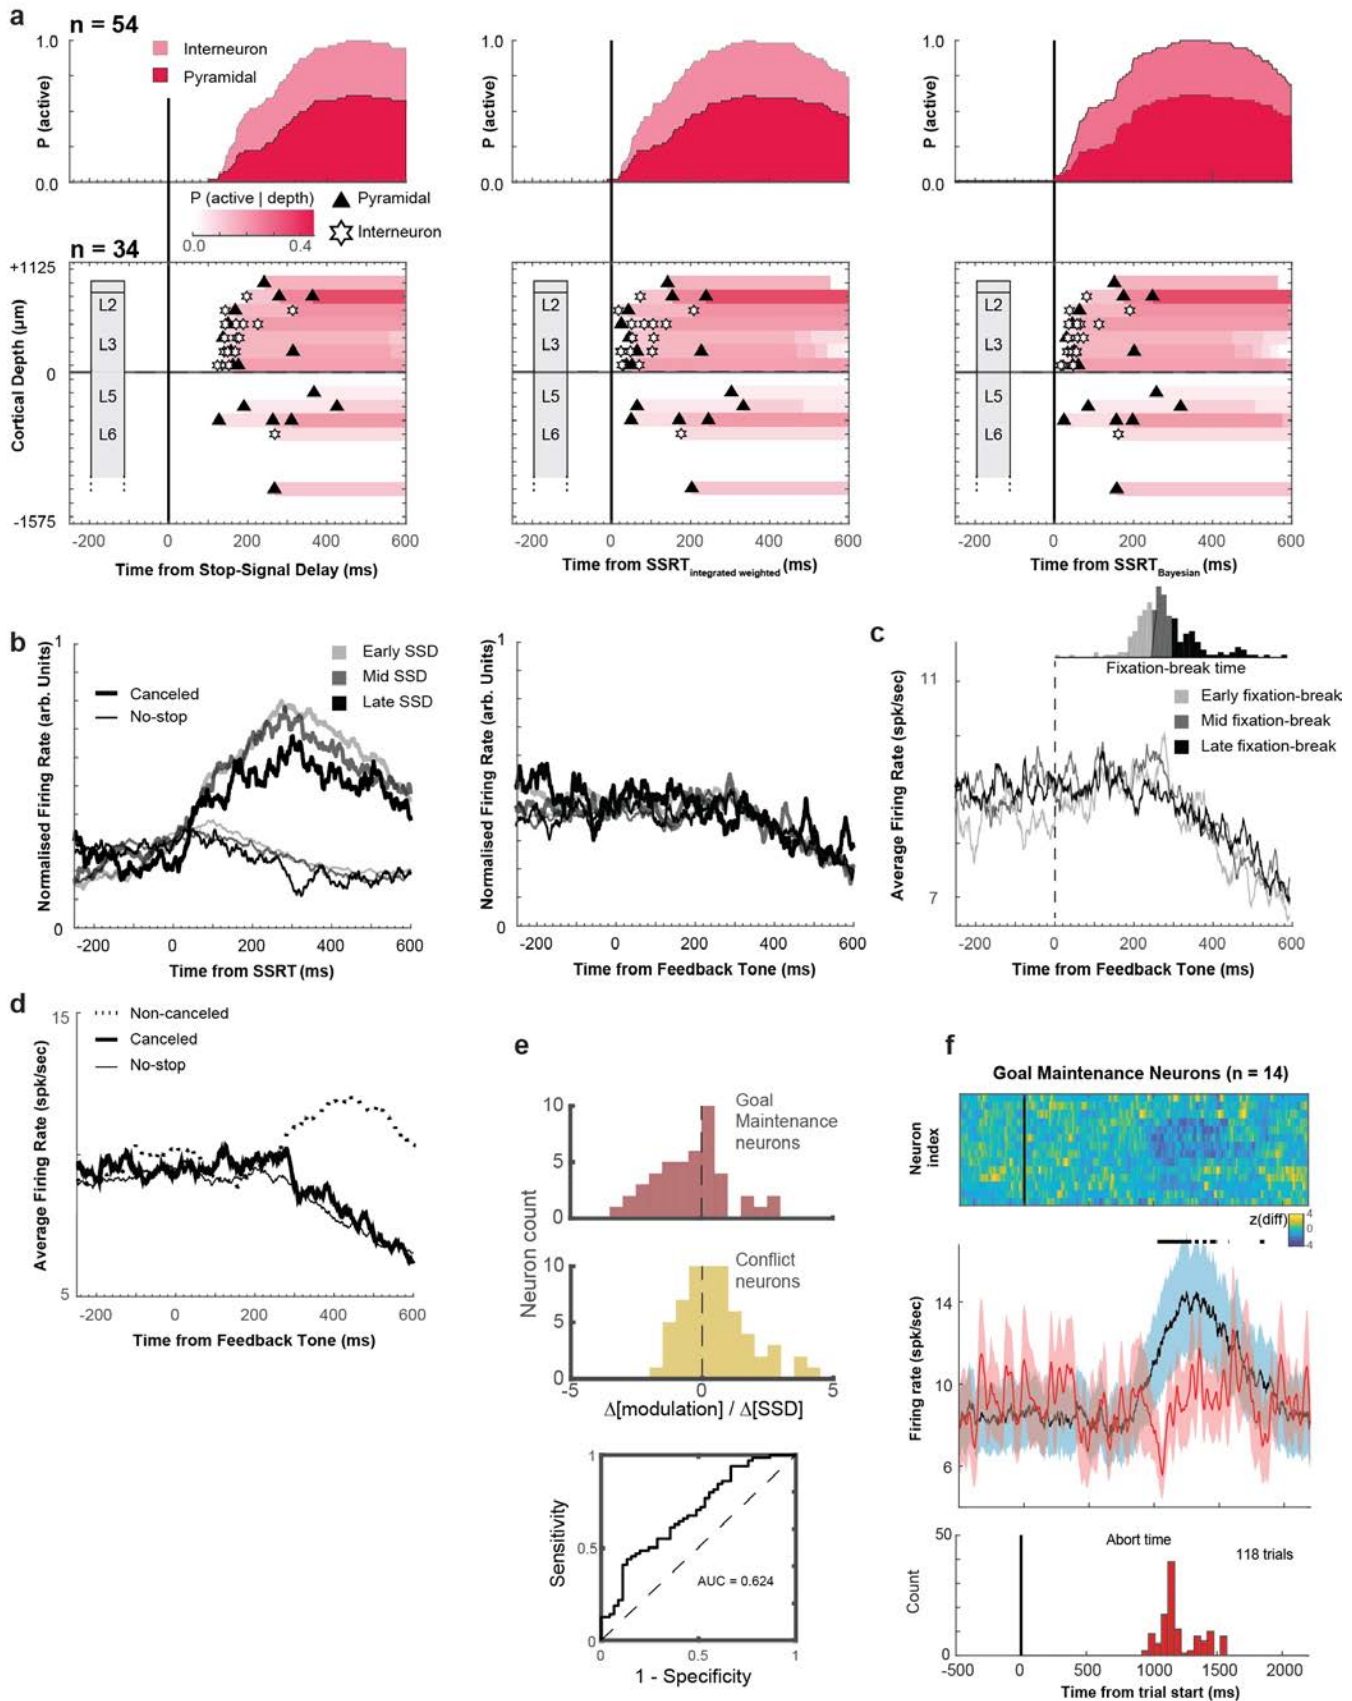


a, Recruitment and time-depth plot aligned on stop-signal (left), SSRT calculated based on the method of integration (middle, same as **Fig 4**), and SSRT by BEESTS (right). The patterns supporting the conclusions reported in the main text do not vary with the method of alignment. This neural signal nearly exclusively follows SSRT and persists until the tone.

b, Normalized population spike-density functions aligned on SSRT (left) and feedback time (right) for canceled (thick) and latency-matched no-stop-signal trials (thin) subsampled from trials with low (lightest), middle (intermediate), and high (darkest) SSD. The timing of the modulation across SSD is evident across the sample. Note that the activity of neurons aligned on tone, on no-stop-signal trials during the pre-tone interval can be larger than that on canceled trials because of post-saccadic activity bleeding into this time period.

c, Average spike-density function for subset of neurons ($n = 84$) aligned on feedback tone for short (light gray), medium (dark gray), and long (black) periods until monkeys shifted gaze from the central point or blinked (inset plots distribution of the time of the first saccade or blink following the feedback tone in one session). The temporal dynamics of the tone-aligned activity did not depend on the time at which fixation was broken following the feedback tone. This pattern was observed on almost all individual neurons (not shown).

Supplementary Figure 6 | Goal Maintenance Neurons



a, Recruitment and time-depth plots aligned on stop-signal (left), SSRT calculated based on the method of integration (middle, same as **Fig 5**), and SSRT by BEESTS (right). The patterns supporting the conclusions reported in the main text do not vary with the method of alignment. This neural signal grows after SSRT and persists until after the tone.

b, Normalized population spike-density functions aligned on SSRT (left) and feedback time (right) for canceled (thick) and latency-matched no-stop-signal trials (thin) subsampled from trials with low (lightest), middle (intermediate), and high (darkest) SSD. The scaling of the modulation across SSD is evident across the sample.


c, Average spike-density function for subset of neurons ($n = 54$) aligned on feedback tone for short (light gray), medium (dark gray), and long (black) periods until monkeys shifted gaze from the central point or blinked (inset plots distribution of the time of the first saccade or blink following the feedback tone in one session). The temporal dynamics of the tone-aligned activity did not depend on the time at which fixation was interrupted following the feedback tone. This pattern was observed on almost all individual neurons (not shown).

d, Population spiking-density function on canceled (thick solid), no-stop-signal (thin), and noncanceled (thick dashed) trials. A large proportion of Goal Maintenance neurons were also classified as Loss neurons, with higher activity on unrewarded compared to rewarded trials.

e, Rate of change of spiking as a function of SSD for Goal Maintenance (red) and Conflict (yellow) neurons. With longer SSD, Goal Maintenance neurons tended to be less active, but Conflict neurons were more active. The distinction between Goal Maintenance and Conflict neurons on this measure is confirmed in the ROC plot.

f, Goal Maintenance neurons with reduced activity when fixation was interrupted. The contrast of neural spiking between trials in which reward was delivered after saccades were canceled and trials when fixation was interrupted by a saccade or a blink before the feedback tone. Matching trials for SSD, 14/54 Goal Maintenance neurons had at least 5 aborted trials. The difference across conditions is plotted as a heat-map (top) with one neuron per row and average \pm SEM spike density (middle) for successful canceled trials (black) and aborted trials (red), smoothed for visualization purposes, with periods of significant difference (two-tailed Wilcoxon test, $p < 0.05$, no multiple comparison correction) indicated by black horizontal lines above the SDF plots. The histogram plots the abort times.

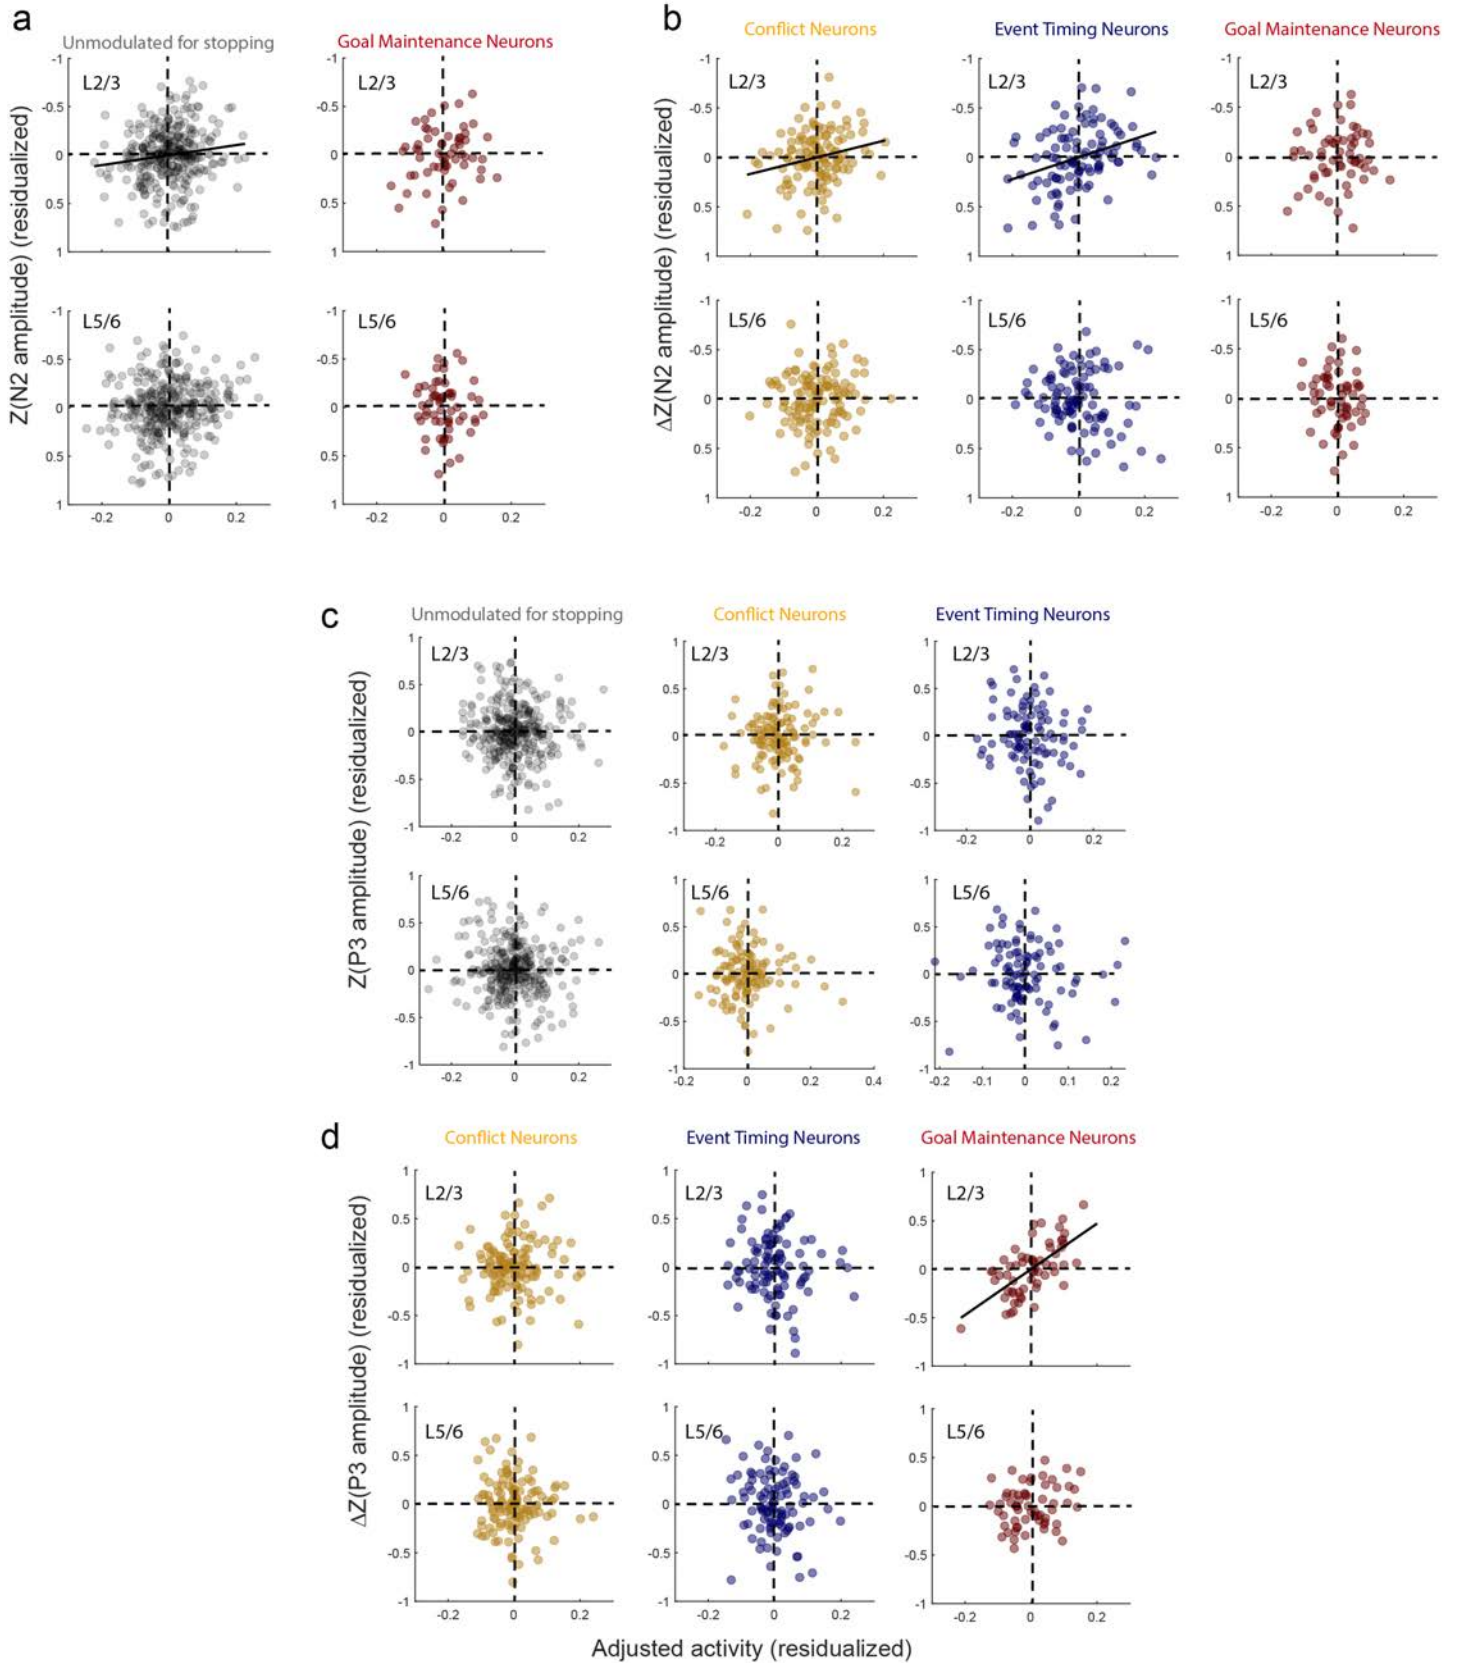
Supplementary Figure 7 | N2/P3 Characteristics.



a, ERP aligned on target presentation (top), and stop-signal (bottom) for early (left), intermediate (middle), and late (right) SSD. At all SSD, the visually-evoked negativity (green dashed line) occurs earlier than the peak of N2 by at least 40 ms. Thus, the N2 is not just a visual response.

b, N2 and P3 difference plots for all sessions (top) and sessions from monkey Eu (middle) and X (bottom) with ± 50 ms period around the peak shaded. The N2 and P3 exhibited similar features across monkeys. **c**, Cumulative distributions of N2 (left) and P3 (right) peak times aligned on stop-signal (top), SSRT by method of integration (middle) and SSRT by BEESTS (bottom). The distributions are narrower when aligned on stop-signal.

Supplementary Figure 8 | N2/P3 relationship with spiking.




a, N2 polarization was predicted by the spiking of SEF neurons that were not modulated around the time of SSRT in L2/3 (Multiple Linear regression (two-tailed) with L2/3 and L5/6 activity as predictors, $t(317) = -2.51$, $p = 0.012$) but not L5/6 (left) but was unrelated to the spiking of Goal Maintenance neurons in L2/3 or L5/6 (right).

b, Difference of N2 polarization between canceled and latency-matched no-stop-signal trials was predicted by spiking of Conflict ($t(117) = -2.63$, $p = 0.010$) and Event timing neurons in L2/3 ($t(97) = -3.57$, $p < 0.001$) but not L5/6 and not by spiking of Goal maintenance neurons.

c, P3 polarization was unrelated to the spiking of SEF unmodulated neurons, Conflict neurons, or Event timing neurons in L2/3 or L5/6.

d, Difference of P3 polarization between canceled and latency-matched no-stop-signal trials was not predicted by spiking of Conflict or Event timing neurons but was predicted by the spiking of Goal maintenance neurons in L2/3 ($t(57) = 6.26$, $p < 0.001$) but not in L5/6. The statistics for the Linear regression models for Conflict, Goal Maintenance, and Event Timing neurons are presented in **Supplementary Tables 3, 4, and 5**.

Supplementary Figure 9 | Relationship of SEF modulation to substantia nigra pars compacta and caudate nucleus.



a, The mesocortical pathway of macaque monkeys. To infer the temporal relationship between VTA/SNpc and SEF signals, the mesocortical pathway was traced on a sagittal slice from the RH12 dataset (Dataset 6, RH12, Slide 27/56, Sagittal, Nissl; brainmaps.org). The estimated was 55.8 mm. Based on measurements of the conduction velocity of DA axons in rodents and assuming similar values for primates (0.54 m/s¹⁶; 0.55 m/s^{17, 18}; 0.58 m/s¹⁹), we estimate that the conduction time of a spike from SNpc to SEF is 100.7 ms (average of 103.4 ms, 101.5 ms, 101.5 ms, and 96.3 ms, respectively).

b, Modulation times of the SEF neurons and of dopaminergic and striatal neurons. Cumulative distributions for Conflict (gold), Event Timing (blue), and Goal Maintenance (red) neurons in SEF are plotted with the latencies of dorsolateral SNpc neurons (dashed green) and of facilitated (dashed purple) and suppressed (dashed blue) neurons in the head of the caudate nucleus obtained from⁵. Also plotted is the estimated distribution of DA spike arrival times in SEF based on the conduction time from SNpc to SEF (green).

Supplementary Table 1 | Sampling distributions across recording sites and cortical depth.

Site	Number penetrations	All neurons	Conflict	Event Timing	Goal Maintenance
np1	6	140	10	29	12
np2	7	142	29	6	8
P1	6	104	12	41	1
P2	6	133	18	6	30
P3	4	56	6	2	3
Total count		575	75	84	54
$X^2(4, 575)$	test statistics	-	11.62	84.13	39.3
	p-value	-	0.020	< 0.001	< 0.001

Layer	All neurons	Conflict	Event Timing	Goal Maintenance	
L2	34	2	3	7	
L3	114	11	22	19	
L5	88	14	19	6	
L6	39	7	2	1	
L6+	18	2	3	1	
Total count		293	36	49	34
$X^2(4, 293)$	test statistics	-	4.28	7.33	11.24
	p-value	-	0.369	0.120	0.024

Layer	All neurons	Event Timing with preTone ramping
L2/upperL3	90	2
Lower L3/ L5	146	20
L6/L6+	57	3
Total count		293
$X^2(2, 293)$	test statistics	10.37
	p-value	0.006

Chi-square test for homogeneity was computed for each signal in a 5 (sites or layers) by 2 (with or without the signal) contingency matrix based on the counts of units.

Top, Count of Conflict, Event Timing, and Goal Maintenance signals across the 575 neurons sampled from five sites in monkey Eu and monkey X. Three of the five sites were sampled with perpendicular penetrations (Eu: P1, X: P2 and P3) and two were not (Eu: np1, X: np2). The probability of sampling each kind of signal varied significantly across sites.

Middle, Count of the neural signals from sites P1, P2 and P3 across cortical depth. Contacts ranged from depth 1 (shallowest) to depth 19 (deepest)^{1, 2}. Neurons were assigned to the four cellular cortical layers plus contacts located beyond the L6 boundary. Conflict and Event timing neurons were distributed homogeneously across depth. Goal Maintenance neurons were significantly more concentrated in L2/3 relative to L5/6.

Bottom, Chi-square test of laminar homogeneity for Event Timing neurons with pre-tone ramping showing that these neurons were significantly more clustered in lower L3 and L5 (Fig 3b). The boundary between upper and lower L3 and between L5 and L6 are shown in Fig 3b (horizontal dashed lines).

Supplementary Table 2 | List of considered models.

Model name	Quantity	Mathematical expression
Conflict	Probability of not canceling	$p(\text{NC} \mid \text{SSD})$
Error-Likelihood	Probability of not canceling despite seeing the stop-signal	$p(\text{NC}_{\text{error}} \mid \text{SSD}) / p(\text{SS}_{\text{seen}} \mid \text{SSD})$
Time-based	Stop-signal delay	SSD $\log(\text{SSD})$
	Feedback tone time	$T_{\text{tone}} = 1500 - \text{SSD}$ $\log(T_{\text{tone}})$

Model variants

Model name	Quantity	Mathematical expression	Trials used for learning time distribution	Additional task knowledge	Time perception inaccuracy	Descriptor
Hazard Rate	Stop-signal hazard rate	$h(t) = f(t) / [1 - F(t)]$	canceled trials	--	--	$h(t)$ Absolute, C-trials
			canceled trials	--	Weber's law	$h(t)$ Subjective, C-trials
			canceled trials	SSD staircasing	Weber's law	$h(t)$ Dynamic, C-trials
			SS seen trials	--	--	$h(t)$ Absolute, SS seen
			SS seen trials	--	Weber's law	$h(t)$ Subjective, SS seen
			SS seen trials	SSD staircasing	Weber's law	$h(t)$ Dynamic, SS seen
Hazard Rate	Feedback Tone hazard rate		canceled trials	--	--	$h(t)$ Absolute, C-trials
			canceled trials	--	Weber's law	$h(t)$ Subjective, C-trials
			canceled trials	$T_{\text{tone}} = 1500 - \text{SSD}_{\text{current trial}}$	Weber's law	$h(t)$ Dynamic, C-trials
Surprise	Stop-signal surprise	$s(t) = -\log_2[h(t)]$	canceled trials	--	--	$s(t)$ Absolute, C-trials
			canceled trials	--	Weber's law	$s(t)$ Subjective, C-trials
			canceled trials	SSD staircasing	Weber's law	$s(t)$ Dynamic, C-trials
			SS seen trials	--	--	$s(t)$ Absolute, SS seen
			SS seen trials	--	Weber's law	$s(t)$ Subjective, SS seen
			SS seen trials	SSD staircasing	Weber's law	$s(t)$ Dynamic, SS seen

The nomenclature, mathematical expression, and descriptor for different models and their variants are shown. Hazard rate and surprise models had different variants based on different assumptions indicated in the table (see **Supplementary Figure 3a**).

Supplementary Table 3 | Comparison of model fits for Conflict neurons exhibiting post-SSRT facilitation

Signal	Response Period	BIC	Δ BIC	Predictor (Z-scaled)	df	β	SE	t	p
Conflict	Transient post-SSRT	980.10	0	intercept	212	2.61	0.20	13.20	< 0.001
				p(NC SSD)		0.61	0.14	4.24	< 0.001
		983.59	3.52	intercept	212	2.59	0.20	13.30	< 0.001
				p(NC _{error} SSD) / p(SS _{seen} SSD)		0.54	0.15	3.72	< 0.001
		981.40	1.32	intercept	212	2.61	0.20	13.22	< 0.001
				SSD		0.59	0.14	4.06	< 0.001
		982.23	2.15	intercept	212	2.60	0.20	13.27	< 0.001
				log(SSD)		0.56	0.14	3.93	< 0.001
		981.40	1.32	intercept	212	2.61	0.20	13.22	< 0.001
				T _{tone}		-0.59	0.14	-4.06	< 0.001
		981.68	1.61	intercept	212	2.61	0.20	13.20	< 0.001
				log(T _{tone})		-0.59	0.15	-4.03	< 0.001
		980.86	0.78	intercept	212	2.63	0.20	13.45	< 0.001
		s(t) _{absolute, C-trials}		0.66	0.16	4.11	< 0.001		
981.56	1.49	intercept	212	2.63	0.20	13.41	< 0.001		
		s(t) _{subjective, C-trials}		0.63	0.16	4.02	< 0.001		
982.12	2.04	intercept	212	2.62	0.20	13.38	< 0.001		
		s(t) _{dynamic, C-trials}		0.63	0.16	3.94	< 0.001		
988.81	8.74	intercept	212	2.60	0.20	13.28	< 0.001		
		s(t) _{absolute, SS seen}		0.48	0.17	2.88	0.0043		
986.85	6.77	intercept	212	2.60	0.20	13.29	< 0.001		
		s(t) _{subjective, SS seen}		0.52	0.16	3.23	0.0014		
985.76	5.69	intercept	212	2.61	0.20	13.32	< 0.001		
		s(t) _{dynamic, SS seen}		0.56	0.16	3.40	< 0.001		

Statistical outcomes are given for mixed-effects modelling identifying which model best fits the neural modulation of Conflict neurons and behavioral/task parameters. The statistics for the best-fit model (top row, bold text) and competing models (lower rows) are shown. The mixed-effects model allowed for modeling random intercepts grouped by neuron for spiking measures. No consideration for multiple comparison is included in these statistics. This data contributes to the statistics table in Fig. 3d.

Supplementary Table 4 | Comparison of model fits for Event Timing neurons

Signal	Response Period	BIC	Δ BIC	Predictor (Z-scaled)	df	β	SE	t	p
Event Timing	pre-SSRT activity	1682.60	0	intercept	250	16.19	1.28	12.62	< 0.001
				log(SSD)		0.76	0.24	3.24	0.001
		1685.35	2.71	intercept	250	16.19	1.28	12.61	< 0.001
				SSD		0.66	0.24	2.77	0.006
		1690.86	8.21	intercept	250	16.19	1.28	12.67	< 0.001
				$h(t)_{\text{absolute, C-trials}}$		-0.41	0.29	-1.41	0.158
		1689.73	7.08	intercept	250	16.19	1.28	12.68	< 0.001
				$h(t)_{\text{subjective, C-trials}}$		-0.51	0.29	-1.77	0.078
		1690.69	8.04	intercept	250	16.19	1.28	12.68	< 0.001
				$h(t)_{\text{dynamic, C-trials}}$		-0.45	0.30	-1.47	0.142
		1692.80	10.15	intercept	250	16.19	1.28	12.63	< 0.001
				$h(t)_{\text{absolute, SS seen}}$		-0.08	0.35	-0.22	0.823
		1691.79	9.14	intercept	250	16.19	1.28	12.68	< 0.001
				$h(t)_{\text{subjective, SS seen}}$		-0.37	0.35	-1.03	0.303
		1691.60	8.95	intercept	250	16.19	1.28	12.68	< 0.001
				$h(t)_{\text{dynamic, SS seen}}$		-0.38	0.34	-1.12	0.264
Event Timing	pre-Tone activity	678.20	0	intercept	112	12.77	1.79	7.15	< 0.001
				log(T_{tone})		0.67	0.20	3.41	< 0.001
		678.29	0.01	intercept	112	12.77	1.79	7.15	< 0.001
				T_{tone}		0.67	0.20	3.41	0.001
		681.28	3.00	intercept	112	12.77	1.78	7.16	< 0.001
				$h(t)_{\text{absolute, C-trials}}$		0.58	0.20	2.87	0.005
		682.34	4.06	intercept	112	12.77	1.81	7.05	< 0.001
				$h(t)_{\text{subjective, C-trials}}$		1.27	0.47	2.68	0.008
		682.69	4.41	intercept	112	12.77	1.81	7.04	< 0.001
				$h(t)_{\text{dynamic, C-trials}}$		1.24	0.47	2.61	0.010

Statistical outcomes are given for mixed-effects modelling identifying which model best fits the ramping neural modulation of Event-Timing neurons and behavioral/task parameters. The statistics for the best-fit model for pre-SSRT and pre-Tone activity (upper row, bold text) and respective competing models (lower rows) are shown. The mixed-effects model allowed for modeling random intercepts grouped by neuron for spiking measures. No consideration for multiple comparison is included in these statistics. This data contributes to the statistics tables in Fig. 4d and 4g.

Supplementary Table 5 | Comparison of model fits for Goal Maintenance neurons exhibiting post-SSRT facilitation

Signal	Response Period	BIC	Δ BIC	Predictor (Z-scaled)	df	β	SE	t	p
Goal Maintenance	Sustained post-SSRT		0	intercept	151	6.52	0.73	8.87	< 0.001
				s(t)_{subjective, SS seen}		-0.65	0.17	-3.91	< 0.001
		801.67	4.10	intercept	151	6.54	0.74	8.89	< 0.001
				p(NC SSD)		-0.48	0.15	-3.26	0.0013
		805.73	8.16	intercept	151	6.56	0.74	8.90	< 0.001
				p(NC _{error} SSD) / p(SS _{seen} SSD)		-0.38	0.15	-2.50	0.013366
		799.45	1.88	intercept	151	6.54	0.73	8.91	< 0.001
				SSD		-0.53	0.15	-3.62	< 0.001
		800.55	2.98	intercept	151	6.55	0.73	8.91	< 0.001
				log(SSD)		-0.50	0.14	-3.45	< 0.001
		799.45	1.88	intercept	151	6.54	0.73	8.91	< 0.001
				T _{tone}		0.53	0.15	3.62	< 0.001
		799.37	1.80	intercept	151	6.54	0.73	8.91	< 0.001
		log(T _{tone})		0.53	0.15	3.63	< 0.001		
798.52	0.95	intercept	151	6.52	0.73	8.87	< 0.001		
		s(t) _{absolute, C-trials}		-0.59	0.16	-3.71	< 0.001		
799.02	1.45	intercept	151	6.52	0.73	8.88	< 0.001		
		s(t) _{subjective, C-trials}		-0.58	0.16	-3.77	< 0.001		
798.89	1.32	intercept	151	6.52	0.73	8.88	< 0.001		
		s(t) _{dynamic, C-trials}		-0.61	0.16	-3.84	< 0.001		
798.47	0.90	intercept	151	6.51	0.74	8.86	< 0.001		
		s(t) _{absolute, SS seen}		-0.65	0.17	-3.76	< 0.001		
798.00	0.43	intercept	151	6.52	0.74	8.86	< 0.001		
		s(t) _{dynamic, SS seen}		-0.62	0.17	-3.69	< 0.001		

Statistical outcomes are given for mixed-effects modelling identifying which model best fits the neural modulation of Goal Maintenance neurons and behavioral/task parameters. The statistics for the best-fit model (top row, bold text) and competing models (lower rows) are shown. The mixed-effects model allowed for modeling random intercepts grouped by neuron for spiking measures. No consideration for multiple comparison is included in these statistics. This data contributes to the statistics table in Fig. 5d.

Supplementary Table 6 | Comparison of model fits for ERP N2 and P3 components.

Signal	Response Period	BIC	Predictor (Z-scaled)	df	β	SE	t	p	
N2	post-SSRT	-52.10	0 intercept $p(\text{NC}_{\text{error}} \text{SSD}) / p(\text{SS}_{\text{seen}} \text{SSD})$	85	-0.120	0.023	-5.470	< 0.001	
		-49.75	2.88	intercept $p(\text{NC} \text{SSD})$	85	-0.124 0.029	0.023 0.016	-5.397 1.844	< 0.001 0.069
		-47.26	5.38	intercept SSD	85	-0.124 0.014	0.024 0.016	-5.243 0.912	< 0.001 0.364
		-47.56	5.07	intercept $\log(\text{SSD})$	85	-0.124 0.017	0.024 0.016	-5.215 1.072	< 0.001 0.28697
		-47.26	5.38	intercept T_{tone}	85	-0.124 -0.014	0.024 0.016	-5.243 -0.912	< 0.001 0.364
		-47.22	5.41	intercept $\log(T_{\text{tone}})$	85	-0.124 -0.014	0.024 0.016	-5.247 -0.894	< 0.001 0.374
		-49.57	3.06	intercept $s(t)_{\text{absolute, C-trials}}$	85	-0.124 -0.005	0.023 0.017	-5.286 -0.294	< 0.001 0.770
		-48.89	3.74	intercept $s(t)_{\text{subjective, C-trials}}$	85	-0.124 -0.002	0.024 0.017	-5.255 -0.128	< 0.001 0.898
		-48.78	3.85	intercept $s(t)_{\text{dynamic, C-trials}}$	85	-0.124 -0.011	0.023 0.017	-5.372 -0.626	< 0.001 0.533
		-46.51	6.12	intercept $s(t)_{\text{absolute, SS seen}}$	85	-0.124 -0.035	0.021 0.018	-5.828 -1.987	< 0.001 0.050
		-46.44	6.19	intercept $s(t)_{\text{subjective, SS seen}}$	85	-0.124 -0.031	0.022 0.018	-5.763 -1.767	< 0.001 0.081
		-46.77	5.86	intercept $s(t)_{\text{dynamic, SS seen}}$	85	-0.124 -0.031	0.022 0.018	-5.750 -1.727	< 0.001 0.088
		P3	post-SSRT	-94.60	0 intercept $\log(T_{\text{tone}})$	85	0.230	0.02	10.60
-91.73	2.88			intercept $p(\text{NC} \text{SSD})$	85	0.231 -0.037	0.022 0.011	10.697 -3.210	< 0.001 0.002
-85.21	9.40			intercept $p(\text{NC}_{\text{error}} \text{SSD}) / p(\text{SS}_{\text{seen}} \text{SSD})$	85	0.231 -0.022	0.022 0.012	10.693 -1.746	< 0.001 0.084
-94.56	0.05			intercept SSD	85	0.231 -0.041	0.022 0.011	10.621 -3.717	< 0.001 < 0.001
-93.02	1.59			intercept $\log(\text{SSD})$	85	0.231 -0.038	0.022 0.011	10.698 -3.441	< 0.001 < 0.001
-94.56	0.05			intercept T_{tone}	85	0.231 0.041	0.022 0.011	10.621 3.717	< 0.001 < 0.001
-88.45	6.16			intercept $s(t)_{\text{absolute, C-trials}}$	85	0.231 -0.033	0.022 0.012	10.639 -2.650	< 0.001 0.010
-90.11	4.49			intercept $s(t)_{\text{subjective, C-trials}}$	85	0.231 -0.035	0.022 0.012	10.628 -2.869	< 0.001 0.005
-89.92	4.68			intercept $s(t)_{\text{dynamic, C-trials}}$	85	0.231 -0.035	0.022 0.012	10.612 -2.836	< 0.001 0.006
-88.84	5.77			intercept $s(t)_{\text{absolute, SS seen}}$	85	0.231 -0.035	0.022 0.013	10.585 -2.574	< 0.001 0.012
-89.89	4.72			intercept $s(t)_{\text{subjective, SS seen}}$	85	0.231 -0.039	0.022 0.013	10.565 -2.924	< 0.001 0.004
-89.71	4.89			intercept $s(t)_{\text{dynamic, SS seen}}$	85	0.231 -0.038	0.022 0.013	10.571 -2.885	< 0.001 0.005

Statistical outcomes are given for mixed-effects modelling identifying which model best fits the variation in the N2 and P3 components and behavioral/task parameters. The statistics for the best-fit model (top row, bold text) and competing models (lower rows) are shown. The mixed-effects model allowed for modeling random intercepts grouped by session for ERP measures. No consideration for multiple comparison is included in these statistics. This data contributes to the statistics table in Fig. 6c

Supplementary Table 7 | Multiplexed signals between Error, Gain, and Loss with Conflict, Event Timing, and Goal Maintenance.

Signal type	N	Percent (%)	Conflict		Event Timing		Goal Maintenance	
			N	Percent (%)	N	Percent (%)	N	Percent (%)
Error signal	61	10.6	10 (5)	13.3	25 (12)	29.8	4 (1)	7.4
Gain signal	91	15.8	11 (1)	14.7	24 (9)	28.6	4 (0)	7.4
Loss signal	189	32.9	30 (4)	40.0	14 (3)	16.7	33 (1)	61.1
Other	234	40.7	32	42.7	33	39.3	15	27.8
Total Count	575		75		84		54	
X ² (3, N = 575)	test statistics		1.02		44.86		19.43	
	p-value		0.791		< 0.001		< 0.001	

Count of Conflict, Event Timing, and Goal Maintenance signals across the 575 neurons relative to previously described Error, Gain and Loss signals². Each of the 3 signals reported in this paper was placed in a contingency matrix based on the counts of previously described neurons conveying the other signals. This contingency matrix was 4 (Error, Gain, Loss, or Neither) by 2 (with or without the signal). The bottom row shows the test statistics for homogeneity based on a chi-square test.

Conflict neurons multiplexed error, gain, and loss signals just in proportion to their incidence of sampling. In contrast, Event timing neurons were significantly less likely than chance to signal Loss and more likely to signal Error and Gain. Goal Maintenance neurons were significantly more likely than chance to multiplex with the Loss signal.

Supplementary Table 8 | Multiple linear regression results for relationship between ERPs and spiking.

a

ERP	Neuronal Type	Predictor	df	β	SE	t	p
N2	Conflict	Intercept	117	~0	0.023	-0.003	0.997
		spkUPPER		-1.04	0.290	-3.6	< 0.001
		spkLOWER		0.0135	0.296	0.046	0.963
N2	Event Timing	Intercept	97	~0	0.028	-0.003	0.998
		spkUPPER		-1.56	0.339	-4.60	< 0.001
		spkLOWER		0.547	0.327	1.67	0.097
N2	Goal Maintenance	Intercept	57	~0	0.036	0.0022	0.998
		spkUPPER		-0.652	0.510	-1.27	0.206
		spkLOWER		0.377	0.725	0.521	0.605
P3	Event Timing	Intercept	117	~0	0.024	0.0038	0.997
		spkUPPER		-0.157	0.356	0.441	0.660
		spkLOWER		-0.170	0.346	-0.492	0.623
P3	Event Timing	Intercept	97	~0	0.031	0.0005	0.999
		spkUPPER		-0.521	0.437	-1.193	0.236
		spkLOWER		-0.342	0.441	-0.775	0.440
P3	Goal Maintenance	Intercept	57	~0	0.029	0.011	0.991
		spkUPPER		2.09	0.383	5.46	< 0.001
		spkLOWER		0.61	0.422	1.47	0.148

b

ERP	Neuronal Type	Predictor	df	β	SE	t	p
N2	Conflict	Intercept	117	~0	0.024	-0.011	0.991
		spkUPPER		-0.84	0.319	-2.63	0.0098
		spkLOWER		-0.299	0.328	-0.911	0.364
N2	Event Timing	Intercept	97	~0	0.029	-0.015	0.988
		spkUPPER		-1.11	0.310	-3.57	< 0.001
		spkLOWER		0.502	0.349	1.44	0.153
N2	Goal Maintenance	Intercept	57	~0	0.036	0.0029	0.998
		spkUPPER		-0.397	0.527	-0.754	0.454
		spkLOWER		0.631	0.781	0.807	0.423
P3	Event Timing	Intercept	117	~0	0.024	0.0026	0.998
		spkUPPER		-0.129	0.337	0.382	0.703
		spkLOWER		-0.216	0.376	-0.575	0.566
P3	Event Timing	Intercept	97	~0	0.031	-0.0018	0.999
		spkUPPER		-0.460	0.422	-1.09	0.279
		spkLOWER		-0.621	0.484	-1.28	0.202
P3	Goal Maintenance	Intercept	57	~0	0.027	0.020	0.984
		spkUPPER		2.35	0.377	6.26	< 0.001
		spkLOWER		0.628	0.404	1.55	0.126

a, Results based on canceled trials with scatterplots in Figs 6f-g and Supplementary Figure 8a & 8c. **b**, Results based on activity difference between canceled and latency-matched no-stop-signal trials with scatterplots in Supplementary Fig 8b,d. In each model ERP voltage was the response variable and spike rate in the upper layers L2/3 (spkUPPER) and lower layers L5/6 (spkLOWER) was the predictor. No consideration for multiple comparison is needed in these statistics.

Supplementary Table 9 | Comparisons of modulation times between neurons in SEF, caudate, and VTA.

	Conflict (90.4 ± 8.5 ms, n = 75)	Event Timing (69.8 ± 6.3 ms, n = 84)	Goal Maintenance (117.7 ± 11.3 ms, n = 54)
SNpc (32.5 ± 8.1 ms, n = 25)	p < 0.001 *	p = 0.063	p < 0.001*
SNpc + mesocortical latency (133.1 ± 8.1 ms, n = 25)	p = 0.012*	p < 0.001 *	p = 0.578
Caudate (Facilitation) (114.4 ± 15.8 ms, n = 40)	p > 0.999	p = 0.123	p > 0.999
Caudate (Suppression) (116.2 ± 11.3 ms, n = 61)	p = 0.807	p = 0.005*	p > 0.999

One-way ANOVA, Onset Latency | Neuron Class, $F(6, 249.897) = 8.291, p < 0.001$

Tukey post-hoc comparisons (* = significant at adjusted $p < 0.05$ level).

Mean ± SEM latencies for type of neurons are shown beneath the cell type labels. A one-way independent measures ANOVA showed significant differences in modulation latencies between the different neuron types in SEF and the SNpc and caudate neurons ($F(6,249.897) = 8.291, p < 0.001$). The table reports outcomes of Dunn-adjusted Bonferroni post-hoc comparisons with statistically significant differences marked by an asterisk. SNpc DA neurons modulated significantly earlier than conflict ($p < 0.001$) and goal maintenance ($p < 0.001$) neurons but were not different from event timing neurons ($p = 0.063$). However, the estimated arrival times of DA spikes in SEF were not different from the modulation times of goal maintenance ($p = 0.578$) neurons but were significantly later than the modulation of event timing ($p < 0.001$) or conflict ($p = 0.012$) neurons. Meanwhile, although suppression after SSRT in the caudate nucleus arose significantly later than the modulation of event timing neurons in SEF (suppression $p = 0.005$), other caudate activity occurred simultaneously with event timing (facilitation $p = 0.123$), conflict (facilitation $p > 0.999$; suppression $p = 0.807$) and goal maintenance (facilitation $p > 0.999$; suppression $p > 0.999$) neurons.

SUPPLEMENTARY REFERENCES

1. Godlove, D.C., Maier, A., Woodman, G.F. & Schall, J.D. Microcircuitry of agranular frontal cortex: Testing the generality of the canonical cortical microcircuit. *J Neurosci* **34**, 5355-5369 (2014).
2. Sajad, A., Godlove, D.C. & Schall, J.D. Cortical microcircuitry of performance monitoring. *Nat Neurosci* **22**, 265-274 (2019).
3. Rapan, L., *et al.* Multimodal 3D atlas of the macaque monkey motor and premotor cortex. *Neuroimage* **226**, 117574 (2021).
4. Lowe, K.A. & Schall, J.D. Functional categories of visuomotor neurons in macaque frontal eye field. *eNeuro* **5**, eNeuro.0131-0118.2018 (2018).
5. Ogasawara, T., Nejime, M., Takada, M. & Matsumoto, M. Primate nigrostriatal dopamine system regulates saccadic response inhibition. *Neuron* **100**, 1513-1526 e1514 (2018).
6. Hanes, D.P. & Schall, J.D. Countermanding saccades in macaque. *Vis Neurosci* **12**, 929-937 (1995).
7. Murthy, A., Ray, S., Shorter, S.M., Schall, J.D. & Thompson, K.G. Neural control of visual search by frontal eye field: effects of unexpected target displacement on visual selection and saccade preparation. *J Neurophysiol* **101**, 2485-2506 (2009).
8. Schmidt, R., Leventhal, D.K., Mallet, N., Chen, F. & Berke, J.D. Canceling actions involves a race between basal ganglia pathways. *Nat Neurosci* **16**, 1118-1124 (2013).
9. Daw, N.D., Courville, A.C. & Touretzky, D.S. Representation and timing in theories of the dopamine system. *Neural Comput* **18**, 1637-1677 (2006).
10. Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. *Psychol Rev* **84**, 279 (1977).
11. Janssen, P. & Shadlen, M.N. A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* **8**, 234-241 (2005).
12. Nelson, M.J., Boucher, L., Logan, G.D., Palmeri, T.J. & Schall, J.D. Nonindependent and nonstationary response times in stopping and stepping saccade tasks. *Atten Percept Psychophys* **72**, 1913-1929 (2010).
13. Starkweather, C.K., Babayan, B.M., Uchida, N. & Gershman, S.J. Dopamine reward prediction errors reflect hidden-state inference across time. *Nat Neurosci* **20**, 581-589 (2017).
14. Emmons, E., *et al.* Temporal learning among prefrontal and striatal ensembles. *Cerebral Cortex Communications* **1**, tgaa058 (2020).
15. Raftery, A.E. Bayesian model selection in social research. *Sociological methodology*, 111-163 (1995).
16. Grace, A.A. & Bunney, B.S. Nigral dopamine neurons: Intracellular recording and identification with L-dopa injection and histofluorescence. *Science* **210**, 654-656 (1980).
17. Thierry, A.M., Deniau, J.M., Herve, D. & Chevalier, G. Electrophysiological evidence for non-dopaminergic mesocortical and mesolimbic neurons in the rat. *Brain Res* **201**, 210-214 (1980).
18. Deniau, J.M., Hammond, C., Riszka, A. & Feger, J. Electrophysiological properties of identified output neurons of the rat substantia nigra (pars compacta and pars reticulata): evidences for the existence of branched neurons. *Exp Brain Res* **32**, 409-422 (1978).
19. Guyenet, P.G. & Aghajanian, G.K. Antidromic identification of dopaminergic and other output neurons of the rat substantia nigra. *Brain Res* **150**, 69-84 (1978).