NEURAL MECHANISMS OF COGNITIVE CONTROL IN THE PRIMATE MEDIAL FRONTAL CORTEX

By

Steven P. Errington

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Psychology

December 17th, 2022

Nashville, Tennessee

Approved:

Professor Geoffrey F. Woodman, Ph.D.

Professor Jeffrey D. Schall, Ph.D.

Professor Sohee Park, Ph.D.

Professor Ueli Rutishauser, Ph.D.

Professor Thilo Womelsdorf, Ph.D.

A mo ghràdh bhoidheach, Lisa

ACKNOWLEDGMENTS

It really does takes a village, and I am incredibly grateful to all of those that have supported me through this PhD.

Firstly, I want to express a wealth of gratitude and thanks to my mentor, Jeff Schall. In addition to the knowledge and skills you've given me through the past five years, you've instilled me with the principles, ethic, and rigor needed to be a good person doing good science.

This has only been amplified by getting to work with a group of incredible colleagues in the Schall Lab. First, I sincerely thank Micala Maddox and Michelle Schall, who not only trained me on the *many* nuances of working with macaques, but also provided endless support throughout my time in the Schall Lab. Second, I thank Thomas Reppert for project discussions, and Amir Sajad for his contributions to some of the work presented in this dissertation. Finally, I thank Kaleb Lowe and Jacob Westerberg for their friendship and for the countless hours of advice and guidance, in and out of the lab.

I have also been fortunate enough to have a supportive committee that has encouraged me to think more broadly and push my limits as a scientist. For this, I extend a profound thank you to Sohee Park, Ueli Rutishauser, Geoff Woodman, and Thilo Womelsdorf. I have always appreciated your feedback and guidance, and have learnt so much from each of our meetings.

I extend further thanks to members of the Department of Psychology that I have had the pleasure to interact with during my time at Vanderbilt. In the office, I thank Beth Clark, Savannah Crutchfield, Erin Duran, and Cris Zerface. I extend a special thanks to Angel Gaither whose kindness and patience was invaluable in helping me navigate the US college system. For contributing to the technical aspects of my research, I would like to thank Isaac Haniff, Pat Henry, Sergey Motorny, Darin Richardson, Chenchal Subraveti, Bruce Williams, and Roger Williams. Finally, I have also been fortunate to learn from some incredible scientists in our department. Specifically, I thank Adriane Seiffert for teaching me how to teach, and I thank Alex Maier and Gordon Logan for sharing their time and knowledge to support my learning.

I thank the contributions of monkeys Darwin, Euler, Fechner, Joule, Kelvin, Hogi, and Xena who have all contributed data to the work presented in this dissertation. I am humbly grateful for their sacrifice in the pursuit of advancing scientific knowledge. For their care and well-being, I also extend thanks to LuAnn Toy, Keith Campbell, Jeremy Parker, Mary Feurtado, and Kate Schuster.

On a personal front, I thank my parents (Helen, Mark, and Paul) and sister (Amy) for all their love and support. I especially thank my Mam for emphasizing the importance of education, teaching me empathy, and fostering a sense of curiosity that has led to me to where I am today.

Finally, I am eternally thankful to my wife, Lisa. Her sacrifice, strength, and love drive me to be the best person I can be. Words cannot begin to explain your impact, but I would never be where I am now without you by my side. We got here together, and we finished this together. I hope this does you proud.

Funding. This work was supported by National Institutes of Health Grants R01-MH-55806, R01-EY-019882, P30-EY-008126, P30-HD-01505, F31-EY-031293, T32-EY-007135, NSERC RGPIN-2022-04592, and Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience.

Publications. Several chapters within this dissertation have been previously published. The study presented in Chapter 2 was previously published in *Journal of Neurophysiology* (Errington and Schall, 2020). The study presented in Chapter 3 was previously published in *Journal of Neuroscience* (Errington et al., 2020). The study presented in Chapter 4 is currently under review at *eLife*. The study presented in Chapter 5 was previously published in *Nature Communications* (Sajad et al., 2022). Aspects of the introduction and discussion are adapted from a review done in conjunction with the Rutishauser lab at Cedar-Sinai; this review was recently accepted at Nature Reviews Neuroscience (Fu et al., 2022b).

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Background

"...to move things is all that mankind can do, and for this task, the sole executant is a muscle, whether it be whispering a syllable or felling a forest"

- Charles Sherrington, 1942

1.1 Actions and decisions

Actions are essential to our daily life and key to survival. Through our actions with the world we can seek sustenance, flee from danger, and reproduce. Indeed, actions are so integral to our existence that some argue that brains have evolved to serve the demands of action rather than cognition (Llinas, 2001). Actions can range from reactive and reflexive to voluntary and intentional (Peak, 1933). Whilst reflexive actions are observed in most nervous systems, voluntary actions form the basis of a principle considered unique to higher-order species – the ability to control and intentionally guide our own behaviors flexibly. In contrast to reflexes, these voluntary and intentional actions are purposive and endogenous: the action comes from, or is made by, an agent rather than an external cause. As such, voluntary actions are believed to result from cognitive processes that enable the agent to exert some specific, guided influence on the outside world to help us achieve a goal.

A decision is a process of deliberation that results in the selection of choice – an overt and purposeful action performed in the context of alternatives. With this, we can consider voluntary actions to result from a series of decisions about how we want to interact with the world. Cognitive science has classically considered decision-making from two conceptually distinct perspectives: perceptual and value-based. Whilst perceptual decision-making requires a judgment based on sensory information, value-based decision-making instead considers the value (or utility) of available alternatives to determine the best choice. Both forms of decision-making can be regarded from two logically and mechanistically distinct aspects: "decide that," in which a choice is made between alternative stimuli or categories, and "decide to," which is a selection between alternative actions (Schall, 2001). This distinction can be contextualized in both perceptual and value-based paradigms in which agents must first decide that an available option is best and then decide to initiate the relevant action to implement the choice. Whilst this distinction is presented in such a way as to capture a range of decisions, similar and overlapping distinctions have been made for decisions to act specifically (Brass and Haggard, 2008).

Insights into decision-making mechanisms can be revealed through reaction times – the time between the cue presentation and the overt response. Whilst reaction times have lower limits imposed by conduction delays from the light hitting the retina, traveling through neurons in the brain, and out to the muscles that generate the response (von Helmholtz, 1850, cited by Schmidgen, 2002), reaction times in decision-making paradigms are typically much longer. Unlike reflexive actions, this signifies that agents must take time to deliberate and make decisions over their voluntary actions (Donders, 1969). Correspondingly, work from cognitive science literature proposes that reaction time can be broken down into two parts: a non-decision time, reflecting the time needed for sensory and motor processing, and a decision time, reflecting the period between these processes in which an agent evaluates evidence and deliberates over available actions (Ratcliff, 1978; Ratcliff et al., 2016).

1.2 Cognitive control

When performing actions, we hope they lead to the best and most desired outcome. As such, some of our decisions about the correct course of action will require more deliberation as all available options must be evaluated and the most appropriate selected. Such deliberation is apparent when we are in unfamiliar environments or when beliefs about the potential outcomes of a decision are uncertain. Furthermore, decision-making must be adaptive: our decisions in one context may not be valid or appropriate in another. For example, deciding to turn left into a roundabout in the United Kingdom may get you where you want to go safely, but this may not have such positive consequences in the United States, where the flow of traffic is reversed. With this, successful decision-making relies on the appropriate detection and interpretation of alternative options, the suppression of irrelevant or inappropriate choices, and the inference and monitoring of the consequences of these choices. This is afforded through cognitive control. Until recently, our understanding of the mechanisms underlying decision-making and cognitive control has developed through largely separate pieces of literature. Whilst decision-making research has mainly focused on integrating information to make a response, the cognitive control literature has instead focused on factors that shape response selection.

1.2.1 Defining cognitive control

Cognitive control (also referred to as executive control or executive functioning) is classically considered to be an umbrella term for the higher-order processes that allow for the flexible allocation of cognitive resources to promote the effective decision-making required for goal-directed actions. Such a process is often considered as an integral, yet undefined, next step in our understanding of cognition; many theories of cognitive processes (such as perception, memory, and attention) define their systems in a modular fashion. These systems process information in a defined way but are engaged and controlled by an undefined and external intelligent agent –

sometimes referred to as a homunculus (Baddeley, 1996).

Although the concept of a control process was introduced in the cognitive science literature earlier (Miller et al., 1960), the first attempts at operationalizing cognitive control occurred in the mid-1970s (Shiffrin and Schneider, 1977; Posner et al., 2004). These early papers proposed two modes of cognitive processing: automatic and controlled. Automatic processing is rapid, autonomous, not limited by attentional capacity, and does not require much effort. Conversely, controlled processes are slower, more effortful, subject to inference, and were initially thought to rely on a central, higher-order, limited-capacity mechanism. However, despite these drawbacks, controlled processes were considered more flexible and could be applied in novel environments. These classical theories propose that these processing modes are dichotomous: a process is either automatic or controlled. Under such a constraint, an automatic process should occur unaffected by controlled higher-order processes. However, this proposition is not supported by literature demonstrating that seemingly automatic behaviors can be influenced by context and task goals (Hikosaka and Isoda, 2010). As such, more recent theories of information processing have since been refined, allowing for the greater flexibility and adaptability of automatic responses depending on the configuration of a cognitive system (Pessoa et al., 2003; Logan, 2003; Schneider and Logan, 2009; Arrington and Logan, 2004) .

This interaction has been used to define "control." Dennett (2015, pg. 57) proposed that control is the relation between two systems in which: "A controls B if and only if the relation between A and B is such that A can drive B into whichever of B's normal range of states A wants B to be in." In line with this concept, cognitive control is now posited more broadly as a top-down system that influences bottom-up processes. Historically, this top-down control function has been conceptualized as a set of functions that may include inhibition, working memory, and cognitive flexibility (Miyake et al., 2000). However, although cognitive control is fundamentally considered a selection of cognitive mechanisms that facilitate goal-directed behavior, many definitions of cognitive control in the literature are now mostly portrayed from a neuroscientific perspective. Specifically, cognitive control is now commonly viewed as synonymous with frontal cortex functions, with activity in these frontal regions exerting top-down modulation over bottom-up sensory processes in posterior and subcortical structures (Fuster, 2001; Miller and Cohen, 2001). In line with this, lesion studies demonstrate that damage to the frontal cortex results in impaired higher-order functions, including deficits in attention, action planning, and goal setting, amongst other cognitive processes thought to comprise cognitive control (Duncan et al., 1997).

1.3 Response Inhibition

Cognitive control is notably important when automatic or learned actions are detrimental to achieving a goal. In such situations, these actions must be stopped before being executed; if not inhibited, these processes may

be incorrectly executed outright, or interfere with task-relevant behaviors. This feature of cognitive control is referred to as response inhibition.

1.3.1 Defining response inhibition

In its simplest definition, inhibition can be considered as the ability to stop a process. However, inhibition as a term to describe specific processes has been applied inconsistently and heterogeneously throughout the literature. Part of this issue may arise due to the independent usage of the word across different conceptual levels and in various fields. Notably, inhibition has long been of interest to both physiology and cognitive science, and resultantly the application of the same term is considered from entirely different perspectives (for a historical overview of the use of the term inhibition, see Bari and Robbins, 2013). However, the emergence of neuroscience as a separate discipline over the past century has resulted in adopting the same terminology to define drastically different levels of analysis. This adoption has resulted in definitions of inhibition in neuroscience ranging from broad, all-encompassing descriptions such as: "any mechanism that reduces or dampens neuronal, mental, or behavioral activity" (Clark, 1996, p. 128); to specific, mechanistic-level descriptions like: "$GABA_B$ receptor mediated reduction in calcium current at the nerve terminal and a subsequent reduction in transmitter release" (Owens and Kriegstein, 1996, p. 717). Unfortunately, trying to force these two levels of analysis under the same umbrella term is unlikely to help advance understanding of either of them (Breese, 1899; MacLeod, 2007).

However, just as animals and plants can be classified into taxa, attempts have been made to understand how response inhibition can be broken down into conceptually similar but differentiable units. Initial proposals suggested a framework for approaching this issue, stipulating that dividing inhibition into taxa should be considered across three separate dimensions (Harnishfeger, 1995). The first dimension considers whether the inhibition is intentional; like the intentional generation of a movement, does the agent consciously intend to stop? The second dimension considers whether the process of inhibition occurs at the action or the cognitive level; does the inhibition concern a motor response or impulse, or does it concern the suppression of attentional or memory processes? The final dimension considers whether inhibition is an active process or a gating process; does inhibition occur through the constant suppression of a process, or does it simply act to stop other processes from happening? These suggestions were used as the foundation to propose a separate taxonomy of inhibitory processes in which inhibition was considered as four distinct factors: interference control, cognitive inhibition, behavioral inhibition, and oculomotor inhibition (Nigg, 2000). In this framework, inference control is proposed as a form of inhibition preventing the inference of task goals from resource demands or stimuli competing for attention within the environment. This feature may also be referred to as attentional inhibition. Conversely, cognitive inhibition is concerned with preventing self-generated and irrelevant task

goals or ideations into working memory. In this taxonomy, behavioral inhibition is a separate process for suppressing a prepared, prepotent, or cued response. An alternative label for this feature may be motoric or action inhibition. Interestingly, this feature is considered a seperately from oculomotor inhibition, responsible for the suppression of a reflexive saccade. However, there is no clear rationale for this separation beyond differences in measurements and neural systems, and the authors acknowledge that this "differentiation is tentative" and that "other frameworks might integrate oculomotor and other motor control systems" (Nigg, 2000, pg. 228).

Whilst these works provided solid foundations for dividing different aspects of inhibition at the cognitive level, Aron (2007) argued that further insight into how inhibition may be parsed could be facilitated by considering the underlying functional neuroanatomy - much like research into other aspects of cognitive control. On this basis, a unified framework of inhibitory control was proposed, highlighting the prefrontal cortex's role (Munakata et al., 2011). This framework builds on previous findings from the neuroscience literature that goal-relevant neural representations are maintained in regions of the prefrontal cortex implicated in working memory, such as the dlPFC (Fujii and Graybiel, 2003; Genovesio et al., 2006; Saito et al., 2005). These goal-relevant representations can include information about the task rules, the stimulus-response association, semantic categories, and value context. The activation of these representations in the prefrontal cortex during a task is thought to lower the activation thresholds in the more posterior, sensory-processing cortices, leading to a bias in processing task-relevant stimuli and promoting context-appropriate responding (Egner and Hirsch, 2005). In contrast to the cognitive taxa described above, this framework considered top-down inhibitory control under two forms: global and competitive inhibition. Global inhibition occurs when neurons in PFC excite GABAergic interneurons in downstream cortical areas, inhibiting the activity within that target area. This mechanism allows for the direct shutdown of processing areas that may be detrimental to achieving the task goal. Conversely, competitive inhibition occurs when neurons in PFC activate excitatory neurons in downstream cortical regions that are conducive to processing goal-relevant stimuli; this boost of processing for one pathway leads to the inhibition of competing pathways. Although formulated indirectly from observations in other neuroscientific literature, this theoretical hierarchical framework of inhibitory control has received since received some empirical support (Tiego et al., 2018). Another study supported the idea that response inhibition (or behavioral inhibition, as proposed by Nigg, 2000, or global inhibition, as proposed by Munakata et al., 2011) and attentional inhibition (or inference control, as proposed by Nigg, 2000, or competitive inhibition, as proposed by Munakata et al., 2011) are empirically independent constructs (Friedman and Miyake, 2004). However, both these constructs exhibited partial statistical dependence on a working memory construct, a hallmark of prefrontal cortex function, as predicted by the unified framework.

Broadly, previous studies aiming to define the taxa of response inhibition converge on three aspects of

inhibition: attentional inhibition, cognitive inhibition, and action inhibition (Diamond, 2013; Friedman and Miyake, 2004) . Note that these conceptual distinctions among different kinds of inhibition roughly correspond to different stages of information processing (Friedman and Miyake, 2004). Inhibition of attention will occur at an early stage of sensory processing and enables an agent to attend to, and focus on, stimuli relevant to the task goal, suppressing attention to other stimuli that may distract from this. Other literature has also referred to this feature of inhibitory control as attentional control or executive attention (Theeuwes, 2010). Inhibition of cognition, on the other hand, suppresses prepotent mental representations that are detrimental to the task goal. This can include resisting unwanted thoughts or memories, or memory traces which may result in retroactive or proactive interference (Postle et al., 2004; Anderson and Levy, 2009). This facet of inhibition is typically discussed in the context of working memory and has greater coherence with measures of working memory and is seemingly a dissociable process from the other forms of inhibition (Friedman and Miyake, 2004). Finally, inhibition of action is considered the suppression of a prepared (or in-preparation) motor command. This is the last stage before a response is actioned, at which point relevant responses must be selected and incorrect ones resisted.

1.3.2 Measuring response inhibition

This dissertation will focus on action inhibition. This form of response inhibition (also referred to in the literature as action-stopping) has typically been studied experimentally through two paradigms: the stop-signal task (Logan and Cowan, 1984) and the go/no-go task (Donders, 1969). These tasks share a common underlying design: trials in which a movement should be generated and those in which movements should be stopped. Concordantly, these two tasks are shown to map onto one common factor and share some similar neural correlates. These correlates include the SMA/pre-SMA, inferior frontal cortex (which may include the right inferior frontal gyrus and anterior insula), and parts of the anterior cingulate cortex (Bender et al., 2016; Swick et al., 2011).

However, whilst conceptually similar, these tasks differ in one key design feature: the time at which the stimuli cueing the inhibition of a response is presented. In no-go tasks, a stop cue is presented simultaneously with (or in place of) the go cue, whilst in stop-signal tasks, the timing of this stop cue is varied and occurs after the presentation of a go cue. As such, although go/no-go and stop-signal tasks are proposed to be interchangeable measures of response inhibition, they may measure distinctively different processes. This difference in task design results in a difference in the associated state of preparedness or readiness when the relevant cues appear. Several conclusions from previous literature lend evidence to this claim. First, behavioral deficits in one task are not associated with deficits in the other, and response times are significantly slowed in stop-signal tasks compared to go/no-go tasks (Raud et al., 2020; Krämer et al., 2013; Littman and

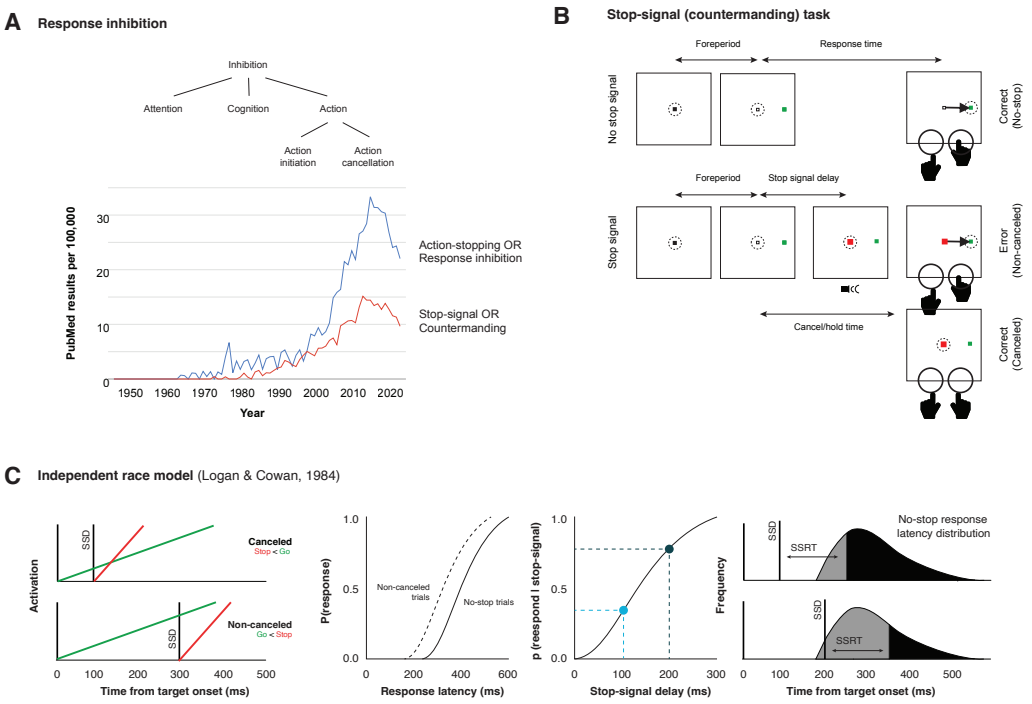

Figure 1.1: **Response inhibition and the stop-signal task**. **A**. *Top.* Differentiating inhibition. Inhibition can be considered from three perspectives (Nigg, 2000): attentional inhibition, cognitive inhibition, and action inhibition. Action inhibition can be divided further into inhibition of action initiation, action cancellation. *Bottom*. Trends of stop-signal (red line) and response inhibition (blue line) publications over time, derived from PubMed entries. **B**. *Top*. Saccade countermanding task. Subjects initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target (a go cue) was presented simultaneously at one of two possible locations. On no-stop-signal trials, subjects shift gaze to the target. On stop-signal trials (typically $\sim$40% of trials), after the target appeared, the center of the fixation point re-illuminates after a variable stop-signal delay, which instructed the subject to cancel the saccade (stop cue). After holding fixation for a period of time, the subjects receive feedback. Stop-signal delay was adjusted such that monkeys successfully canceled the saccade in $\sim$50% of trials (staircase) or were randomly selected from a pre-determined set of stop-signal delays. If subject shifted gaze to the target on stop-trials, this resulted in non-canceled errors. *Bottom*. Manual response (button-press) countermanding task. Subjects initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. In its place, an arrow pointing to the left or right appeared. Subjects were required to press the corresponding button (i.e., left button press for left arrow). On stop trials, the stop-cue is presented through a change in the color of the arrow (i.e., from green to red), or with the presentation of an auditory cue. The sequence of events is otherwise similar to saccade countermanding. **C**. Independent race model. The independent race model considers behavior to result from a race between two processes: a GO process (green line) that drives the generation of an action after the presentation of a go cue, and a STOP process (red line) that aims to cancel any planned actions if a stop-signal appears. After their initiation following the relevant cue, activity of these processes starts racing toward a finish line, or threshold. These processes race independently and do not influence one another. If the STOP process reaches a finish line first, no response will occur, and the movement will be successfully canceled (depicted in the top schematic). Alternatively, the action will be executed if the GO process finishes first (depicted in the bottom schematic). Two key assumptions are made in this model. First (center left panel), reaction times are faster for errors (dashed line) compared to no-stop trials (solid line). Second (center right panel), there is increased probability of moving with an increasing delay between the go and stop cues. The relationship between going and stopping as a function of stop-signal delay is depicted through the inhibition function. These observations allow for the latency of the STOP process to be estimated – referred to as SSRT (right panel). SSRT can be calculated for a given SSD by first finding the point on the no-stop RT distribution where the area under the distribution is equal to the probability of responding at that stop-signal delay and subtracting this value away from the SSD. This is discussed further in text.

Takács, 2017). These observations suggest that these two tasks may have distinct cognitive mechanisms. Second, although overlapping in some regions (as described above), the two tasks also have distinctive patterns of task-related activity in the brain (Swick et al., 2011). Although neurons in FEF signal the location of relevant targets in both tasks, movement neurons only demonstrate preparatory activity to generate a saccade in the stop-signal task; this is absent in go/no-go tasks (Thompson et al., 1997; Schall, 1991a; Hanes et al., 1998). Furthermore, whilst go/no-go tasks were associated with increases in BOLD activity in the right medial frontal gyrus (rMFG) and inferior parietal lobe (IPL), no such activation was observed in the stop-signal task. Conversely, increased thalamic activity was noted in the stop-signal task but not in the go/no-go task. Distinctions were also made at the network level; the frontoparietal network is engaged to a greater extent when a movement is not made in the go/no-go task, whilst the cingulo-opercular network is more engaged when a movement is successfully inhibited in the stop-signal task. Third, recent work has also demonstrated distinct differences in the spatial-temporal pattern of EEG activity (Raud et al., 2020). In go/no-go tasks, the frontoparietal attentional system is activated early and is followed by activation of motor control components; in the stop-signal task, however, there is clear additional recruitment of frontal areas before the activation of attentional systems, followed by early changes in the recruitment of motor components. Finally, cross-species neuropharmacological studies have demonstrated contrasting effects of different neurotransmitters within this task (Eagle et al., 2008). Serotonin (5-HT) administration markedly impairs inhibitory control in the go/no-go task but has no significant impact on stop-signal task performance. Conversely, although less precise, noradrenaline (NA) administration notably improves stop-signal task performance but has little effect on go/no-go task performance.

Collectively these divergent observations between the go/no-go and stop-signal tasks support the claim that they may measure different features. This idea is not novel to this dissertation and has previously been proposed (Eagle et al., 2008; Raud et al., 2020). Instead, it is suggested that the go/no-go task measures a specific component of response inhibition labeled action restraint, and the stop-signal task instead measures another feature labeled action cancellation (Schachar et al., 2007; Eagle et al., 2008). As the go & stop cues are presented simultaneously in the go/no-go task, there is no active preparation of the motor plan to generate a movement, as it had not been cued and thus does not need to be inhibited. Resultantly, it is argued that this paradigm may be better suited to provide a measure of action restraint - a propensity to respond to external stimuli. Conversely, the stop-signal task has a temporal asynchrony between the presentation of the go and stop cue, allowing ample time for a movement to be prepared. As such, when the stop-signal appears, the design of this task means participants must actively inhibit a movement that is being prepared - referred to here as action cancellation.

I will focus on action cancellation in this dissertation. This feature of action inhibition is vital when

adapting to new cues in our ever-changing environment; in such contexts, we are likely already interacting with our world when this salient cue arrives. As such, it is crucial that we can identify the systems that interpret these salient cues and successfully stop previously prepared actions, which now may be irrelevant or detrimental to our goal.

1.4 Stop-signal task

Action cancellation is a form of response inhibition that has been classically examined using the stop-signal (or countermanding) task (Figure 1.1B). The first report of a stop-signal-like experiment probed the "psychological refractory period," in which it was observed that participants could inhibit an initial response when a second cue to stop was shown 50 ms after the initial cue to move (Vince, 1948). However, if this second cue was shown later ($\sim$100 ms and longer), stopping the original action rarely occurred. Similar experiments followed using similar methodologies, again demonstrating that stopping a response was more common when the instruction to stop was delivered soon after the instruction to respond (Lappin and Eriksen, 1966; Ollman, 1973; Slater-Hammel, 1960). Collectively, these studies demonstrated a critical period in which a planned response could be inhibited.

In the stop-signal task, participants must perform a standard reaction task in which they respond as quickly as possible to a go-cue. However, in a small proportion of trials, a second cue is presented at a variable time after the presentation of this initial go-cue instructing participants to stop the planned response. This period between the presentation of the go-cue and the appearance of the stop-signal is known as the stop-signal delay (SSD). Although this standard task remains the same, aspects of the task design have varied. Adaptations include changing the type of stimuli used to cue going & stopping (i.e. Cabel et al., 2000) and manipulating the difficulty of going and stopping (i.e. Middlebrooks et al., 2020), for example. Further adaptations have also seen the stop-signal task used to examine response inhibition across a range of effectors, including reaching (Mirabella et al., 2011), hand squeezing (De Jong et al., 1990), keyboard presses (Logan and Irwin, 2000), moving a joystick (Boucher et al., 2007; Scangos and Stuphorn, 2010), movements of the elbow (Kudo and Ohtsuki, 1998), precision gripping (De Jong et al., 1990), foot presses (Jong et al., 1995), or some a combination of the above (Brunamonti et al., 2012). Saccadic stop-signal tasks, in which subjects are required to stop ballistic eye movements to a target, have also been frequently used with humans (Asrress and Carpenter, 2001; Born et al., 2014; Cabel et al., 2000; Corneil and Elsley, 2005; Goonetilleke et al., 2012; Hanes and Carpenter, 1999; Joiner et al., 2007; Kornylo et al., 2003; Logan and Irwin, 2000; Morein-Zamir and Kingstone, 2006; Stevenson et al., 2009; Wessel et al., 2013) and macaques (Brown and Heathcote, 2008; Emeric et al., 2007, 2008, 2010; Godlove and Schall, 2016; Hanes et al., 1998; Hanes and Schall, 1995, 1996; Ito et al., 2003; Middlebrooks et al., 2020; Ogasawara et al., 2018; Pare and Hanes, 2003; Pouget et al.,

2011; Sajad et al., 2019; Stuphorn et al., 2000; Stuphorn and Schall, 2006; Stuphorn et al., 2010; Walton and Gandhi, 2006).

The stop-signal task has increased in popularity over time (Figure 1.1A) and is now a commonly used tool for experimentally testing response inhibition (Verbruggen et al., 2019b). However this increase in popularity has also led to increased heterogeneity in integral aspects of task design. Such heterogeneity can impact the validity of the stop-signal task as a measure of response inhibition and limits the extent to which findings can be compared across studies. To address this issue, a consensus was published outlining twelve recommendations to standardize and improve the overall quality of future stop-signal research (Verbruggen et al., 2019b). Whilst I will not review all these recommendations in this section in detail, I will summarize them briefly. First, these recommendations propose that stop-signal tasks should be designed to use an appropriate go-task and use a salient stop-signal on only a minority of trials ($\sim$25%). The timing of this stop-signal should be adjusted (or staircased) through a tracking procedure and include a sufficient range of stop-signal delays. Participants should also be provided with feedback on their performance. Second, the consensus states that an important metric derived from this task (stop-signal reaction time; see below) should only be calculated when assumptions and task performance criteria are met. Finally, it proposes that researchers should include sufficient trials in their analysis and provide clear detail about their methodology and behavioral observations. Ultimately, these recommendations aim to increase the quality of response inhibition research, significantly accelerating its progress across domains.

1.4.1 Independent race model

Further insight into the underlying mechanisms of response inhibition in the stop-signal task is afforded by the independent race model, which can explain behavior in this task (Figure 1.1C, Logan and Cowan, 1984). This model considers behavior to result from a race between two processes: a GO process that drives the generation of an action and a STOP process that aims to cancel any planned actions. It assumes that a GO process is initiated on the presentation of a go-cue (i.e., the appearance of a target). The activity of this process then starts accumulating toward a threshold. If and when a stop-signal occurs, a separate STOP process is initiated. Like the GO process, this STOP process also starts racing towards the same threshold. When active, these STOP and GO processes race independently and do not influence one another. The race is over when one process reaches the threshold. If the STOP process reaches a threshold first, no response will occur, and the movement will be successfully canceled. Alternatively, the action will be executed if the GO process reaches the threshold first. As such, the relative finishing times of these two processes can explain behavior on this task. Although the finishing time of the GO process can be observed through reaction time, the STOP process is unobservable. However, the latency of the STOP process can be estimated in a measure

known as stop-signal reaction time; this will be discussed later in this section.

Although the application of the race model to the stop-signal task is highly generalizable, as it does not require parameterization, it does have explicit assumptions about the underlying processes that must be met for its valid application. A fundamental assumption is that the finishing times for these STOP and GO processes are independent random variables. This independence is considered in two forms: context and stochastic. Context independence stipulates that the finishing time distribution of the GO process is the same whether or not a stop-signal is presented. This assumption is essential as it allows for the observed go-RT distribution to be used to estimate the distribution of go latencies on stop-signal trials. Stochastic independence, on the other hand, stipulates that the stop and GO processes do not interact or influence each other in any way. If these assumptions are met, the generalizability of the independent model allows for its application to any variation of the standard stop-signal task, such as those described in the previous sections. This includes different response modalities and effectors, populations, and species.

As alluded to earlier in this dissertation and demonstrated in previous studies, a key observation in this task is an increased probability of moving with an increasing delay between the go and stop cues. The relationship between going and stopping as a function of stop-signal delay is depicted through the inhibition function. The inhibition function classically shows an increased probability of generating a movement as a function of stop-signal delay and is a valuable representation of inhibitory control - they can be used to provide insights between different conditions or groups. Three factors influence this function: stop-signal delay, distribution of the GO response latencies, and stop response latencies. First, as implied above, the race model assumes that SSD will influence the relative finishing time of the STOP process. At longer stop-signal delays, the STOP process will become active later, relative to the GO process. As a result, the GO process has longer to accumulate at later stop-signal delays and thus be closer to the threshold before the STOP process becomes active. Alternatively, at early stop-signal delays, the go-process accumulation is lower when the STOP process starts and thus is more likely to be "caught up" or "beaten" in the race. Consequently, the probability that the GO process will finish before the STOP process increases with stop-signal delay, resulting in a monotonically increasing inhibition function. Second, as the inhibition function plots the relative finishing time of the GO and STOP processes, the respective latencies of these processes will influence the inhibition function. If the GO process becomes slower, this will decrease the probability of responding for each stop-signal delay; at each stop-signal delay, the STOP process is more likely to win, as the slower GO process is not as far into the race. This is depicted as a rightward shift in the inhibition function. Conversely, if the STOP process becomes slower, it will increase the probability of responding for each stop-signal delay; at each stop-signal delay, the STOP process is less likely to win as it will not have the speed required to catch the GO process within the race. This is depicted as a leftward shift in the inhibition

function. To allow for a fairer comparison between groups, these differences in RT can be accounted for by calculating the z-scored relative finishing time between the GO and STOP process (ZRFT).

As this approach accounts for differences in response latencies, if the ZRFT inhibition functions compared between groups are similar, we can consider that the same inhibitory processes can apply across them (Logan and Cowan, 1984). Such approaches have been helpful in comparing different effectors: although latencies were shorter for hand and eye movements, the same inhibitory principles apply to both (Logan and Irwin, 2000). However, suppose ZRFT functions are misaligned between groups. In that case, this can indicate that the mechanisms of inhibition vary between groups (Schachar and Logan, 1990), although caution has been urged about making such conclusions (Band et al., 2003).

Whilst the inhibition function has been a valuable tool for comparing inhibitory behaviors between groups, perhaps the most utilized affordance of the independent race model is that it allows for an estimate of the unobservable stop-latency, known as stop-signal reaction time. Whilst the finishing time of the GO process can be observed through the response time - the time at which a movement is generated relative to the time it is cued - there is no observable response if the STOP process wins the race.

1.4.2 Stop-signal reaction time

Stop-signal reaction time is the estimated latency at which the action of inhibiting a planned movement concludes (Logan and Cowan, 1984). This can also be considered as how long the STOP process requires to successfully finish the race. As such, SSRT is a key metric of inhibitory function, quantifying an individual's ability to inhibit responses; a faster SSRT is associated with greater inhibitory ability, whilst a slower SSRT is associated with poorer inhibition.

Stop-signal reaction time can be calculated with knowledge about the inhibition function (or the probability of responding given a stop-signal) and the distribution of no-stop reaction times. Briefly, SSRT can be calculated for a given SSD by first finding the point on the no-stop RT distribution where the area under (or integrating) distribution is equal to the probability of responding at that stop-signal delay. This value represents the point at which the internal response to the stop-signal occurred. We can then subtract the stop-signal delay from this value to get an estimate of stop-signal reaction time (Logan and Cowan, 1984). Mathematically, this process can be formalized in the following equation:

$$SSRT = [\int_{-\infty}^{r(SSD)} f(t)\,dt] - SSD$$

where SSD is the stop-signal delay, and $r(SSD)$ is defined as the point on the cumulative no-stop response time distribution that is equal to the probability of responding at the given stop-signal delay.

As we typically have more than one stop-signal delay within our experiments, we can then use this basic principle to estimate stop-signal reaction in multiple ways (Band et al., 2003; Verbruggen and Logan, 2009a). These methods differ in their assumptions and their measurements to estimate SSRT. Calculating SSRT at one SSD is unreliable; estimates of SSRT are shorter at later SSDs and longer at earlier SSDs (Logan and Cowan, 1984). As such, estimating SSRT based on an early or late SSD alone does not yield reliable estimates; this is because estimates for these SSDs are derived from the tails of the go-RT distribution, with less frequent observations. However, SSRT can be reliably estimated from SSDs in which the race between stopping and going is tied ($SSRT_{center}$; Logan et al., 1997). Such situations occur at the SSD in which participants respond on 50% of stop trials and inhibit on 50% of stop trials at the given SSD. Alternatively, biases in SSRT estimates can also be reduced by averaging estimates across all SSDs. These estimates can be weighted by the number of trials at each SSD ($SSRT_{weighted}$) or limited to SSDs in which the p(respond | stop-signal) is between 0.15 and 0.85 ($SSRT_{average}$). Finally, some experiments are designed so that the stop-signal task balances the race between going and stopping through a staircase procedure. As the stopping latency can be derived from the relative finishing time between STOP and GO processes, SSRT may be estimated by the difference between the mean SSD and the mean go-RT ($SSRT_{mean}$) or the median SSD and the mean go-RT ($SSRT_{median}$). Nevertheless, despite the slight variations in how they are calculated, these methods all provide viable estimates of estimating the latency of the STOP process.

Unfortunately, all these approaches to estimating SSRT provide only a summary measure of the latency of stopping. Whilst these singular estimates have been unquestionably useful, additional features of the data need to be considered. In simple reaction time tasks, not only is the mean RT informative, but the shape of the distribution also provides valuable insights into underlying cognitive processes (Heathcote et al., 1991; Matzke and Wagenmakers, 2009). Whilst some clinical populations may not differ in their mean RT, they may differ in the variability of the response latencies. For example, although similar in mean response latency, patients with Schizophrenia have response latencies that are more variable and have a markedly different distribution to that of controls (Belin and Rubin, 1995). As such, not considering other aspects of the SSRT distribution may lead to missing essential distinctions between groups or experimental manipulations.

Until recently, capturing such features of the data has proven troublesome. In the original description of the race model, Logan and Cowan (1984) showed that the variance of SSRT can be derived from the variance of the inhibition function and the variance of the GO response times. This is based on the observation that the variance is proportional to the slope of the cumulative distribution at the median. However, using this method, SSRT variance values were overestimated (Band et al., 2003). An alternative approach proposed that the survival distribution of SSRTs can be recovered using the distribution of no-stop RTs, non-canceled RTs, and the probability of responding for each SSD (Colonius, 1990). From this, descriptive features of

the distribution (i.e., median and interquartile range) can be derived. Further work demonstrated that for this approach to produce accurate and viable results, it requires many observations (Band et al., 2003). Additional critiques showed through simulations that this method also underestimates SSRT and overestimates its variability (Band et al., 2003).

Such problems are commonplace when using non-parametric methods that estimate distribution tails from data (Luce, 1986). However, the recent development of a novel Bayesian parametric approach to calculating SSRT distributions has begun to address this. This approach assumes that SSRTs follow an ex-Gaussian distribution - a distribution that results from the convolution of a Gaussian and an exponential distribution and is thought to capture the properties of RT distributions effectively. This distribution is described by three parameters that can be fit to the data: the mean of the normal distribution (μ), the standard deviation of the normal distribution (σ), and the exponential decay parameter ($\tau = 1/\lambda$). Although useful for describing reaction time distributions, the relationship between these parameters and specific cognitive constructs is debated (Matzke & Wagenmakers, 2009). In contrast to the classic frequentist approach, where these parameters are fixed, the Bayesian approach instead assigns probabilities to these parameters, allowing for an entire distribution to be estimated. This is achieved using Markov chain Monte Carlo (MCMC) sampling to calculate posterior distributions for the model parameters. Once these parameters have been estimated, we can then derive more features of the stop latency distribution.

Through the BPA approach, we can also address further issues that are considered to lead to biases in estimating SSRT. As described in the previous section, the race model assumes that the GO runner is triggered by the presentation of the go stimulus and the STOP runner by the presentation of the stop-signal. However, violations of this assumption are apparent within the literature. First, go trials with no response (labeled go omissions) are often observed; in such cases this would indicate the GO runner does not race despite being cued to start. Second, at very short stop-signal delays (including 0 ms, when go and stop cues are simultaneously presented), participants may still respond to the go-cue; this should not occur at early SSDs given the rapid nature of the STOP process relative to the GO process. Just as go-omissions happen when the go-cue does not trigger the runner, this finding indicates that the STOP runner may also not be triggered on all stop trials (labeled trigger failures). Indeed, effective response inhibition requires not just a fast STOP process but also requires this process to be reliably triggered. Recent literature has suggested that SSRT is a more effective measure when combined with knowledge about the reliability of triggering the STOP process (Skippen et al., 2019; Matzke et al., 2017a,b). Deficiencies in initiating this process may occur when stop-signals are not appropriately detected or attended to or when the stop-signal is not translated into an internal STOP process. Although traditional methods are unable to estimate the stop trigger failures unambiguously (Band et al., 2003), such estimates are afforded through the parameters derived from the BPA

model (Matzke et al., 2013a, 2017b).

Metrics of stopping have helped explain the development and age-related decline of other cognitive abilities (Diamond and Gilbert, 1989; Rey-Mermet and Gade, 2018). They have also provided numerous insights into many clinical disorders. For example, individuals with attention deficit hyperactivity disorder (ADHD) or obsessive-compulsive disorder (OCD) demonstrate significantly longer SSRTs and thus show impaired inhibitory control compared to healthy controls (Schachar and Logan, 1990; Penades et al., 2007). In patients with Parkinson's disease, slower SSRT is associated with less availability for dopamine binding in the amygdala and hippocampus (Mann et al., 2021). Patients with schizophrenia (SZ) also have considerable difficulties in inhibiting planned responses (Thakkar et al., 2011); poorer inhibitory ability was also significantly associated with higher composite scores of Negative Symptoms (i.e., anhedonia, avolition, social withdrawal) and was impaired to a greater magnitude in patients with worse occupational functioning. Recent work has also demonstrated that patients with SZ may also have a higher rate of trigger failures, which may represent deficits in the sensory or attentional processing of the stop-signal (Matzke et al., 2017a).

Collectively, stop-signal reaction time has proven itself a valuable metric for differentiating inhibitory control abilities between individuals, groups, and species. Further advances have been afforded in recent years with the development of a new model that allows for other features of stopping behavior to be derived, such as the variability in the stop latency and the reliability of triggering the STOP process. These advances require further consideration when investigating how response inhibition is achieved at the neural level.

1.5 Dual mechanisms of control

Adaptive behavior relies on the ability to meet and adjust to the demands of a changing, complex environment. Response inhibition allows us to stop actions that are inappropriate or may no longer be warranted, given new information about our environment. In defining response inhibition I have described it as a form of control that happens "in the moment" - an agent reacts to new information and immediately stops a planned process. However behavior can also be adapted in advance of potential situations where abrupt changes are anticipated ahead of time. These two forms of control have been formalized through the dual mechanisms of control framework (Braver et al., 2007; Braver, 2012). This framework stipulates that flexibility in allocating cognitive resources can be achieved by modulating how a particular control mechanism is deployed in response to changing task demands or internal goal states. The central hypothesis of this framework proposes that cognitive control processes, such as response inhibition, occur across two operating modes: reactive and proactive.

Reactive control can be conceptualized as a late-correction mechanism, engaging control only as needed and after detecting a high interference event. Such mechanisms can directly influence in-progress actions

and decisions to promote correct behavior through the rapid detection and resolution of interference. On the other hand, proactive control mechanisms are engaged before a cognitively demanding task, biasing attentional, perception, and motor systems in a goal-driven manner. Whilst reactive control is dependent on the rapid detection of interfering processes, proactive control can be guided by an agent's knowledge and expectations about the task. The development of such knowledge and expectations can guide inferences about the optimal time to deploy cognitive resources. These expectations that guide proactive control can be driven by two sources: explicit predictions provided through external cues and implicit predictions based on previous experiences within the task (Jiang et al., 2018). These sources of expectancy are learnt through an actor-critic loop, in which an agent monitors the outcome of their previous actions to learn about the current state of the environment. This information can then be used to prepare future behaviors proactively.

1.5.1 Reactive and proactive control of stopping

Over the past several decades, evidence has established that successful response inhibition relies on reactive and proactive control processes (Aron, 2011; Meyer and Bucci, 2016). Successful performance in the stop-signal paradigm not only requires the recruitment of control processes in reaction to the perception of a stop-signal but also involves monitoring the race between going and stopping and adjusting response strategies to balance the competing demands of the two processes. The stop-signal task, therefore, is uniquely suited to study the neuronal mechanisms of reactive and proactive control.

In the context of going, observable behaviors can be mapped to neural processes that occur around that behavior; for example, we can look at how the pattern of neural activity changes between trials in which a movement is made compared to when a movement is inhibited, and can also consider how the activity varies with changes in response latency. Similar approaches for stopping can be afforded through stop-signal reaction time. In the previous sections I explained how stop-signal reaction time provides us with a measure of how long it takes for a planned action to be stopped. As such, if the neural activity is reactive to the stop-signal and directly inhibits movements, then it must demonstrate a significant change in dynamics after the stop-signal appears and before the inhibition is completed (as indexed by SSRT).

These approaches for identifying activity that directly controls movement have been formalized (Hanes et al., 1998). To contribute to movement initiation the activity must demonstrate greater activity in trials in which a movement was made compared to activity in trials when no movement is made. On stop-trials this activity must be similar – activity that increases along a trajectory that would lead to the initiation of a movement. However, on the presentation of a stop-signal, this activity must then be suppressed and modulated back to a baseline level. Alternatively, suppose a brain area contributes to reactive inhibition. In that case, its activity must be greater in trials in which a movement was successfully inhibited compared to trials in which

a movement would have been canceled had a stop-signal appeared (referred to as latency-matched no-stop trials). In contrast to activity that generates movement, activity that contributes to movement inhibition must modulate after a cue to stop (i.e., a stop-signal) but before the completion of response inhibition (i.e., SSRT). If activity meets these criteria, we can consider it to reflect a reactive control process, mobilized only as needed and in a just-in-time manner. It is important to consider that as this control mechanism is engaged at short notice, it requires the ability to generate control signals at high speed that can influence ongoing motor activity even at a late stage of the movement preparation. Such activity has been well described in the frontal eye fields and superior colliculus and is discussed in later sections.

Processes supporting stopping can also be proactively recruited in anticipation of a stop-signal. Several studies have demonstrated that participants change their response strategies dependent on the likelihood of encountering a stop-signal (Ramautar et al., 2004; Dimoska and Johnstone, 2008; Emeric et al., 2007) and previous experiences in the task (Emeric et al., 2007). As performance in the task can be contextualized as a race between going and stopping, we can consider that success in the going process can imply failure in the stop task and vice versa: faster GO processes yield a higher probability of responding given a stop-signal, whilst slower GO processes result in a lower likelihood of incorrectly responding to the stop-signal. Furthermore, the stop-signal task is inherently temporal by design: a stop-signal appears in a relatively small window of time relative to a go cue. As such, although instructed to respond as quickly and accurately as possible, participants may learn to wait a period of time before initiating a movement to try and "catch" the stop-signal (Verbruggen et al., 2013).

With this, changes in response strategies can be characterized through a speed-accuracy trade-off, with participants demonstrating longer response latencies in instances in which the stop task is prioritized to try and improve accuracy. Several studies have shown that humans and monkeys slow their reaction time following non-canceled errors (Emeric et al., 2007; Sajad et al., 2019) – instances in which the GO process was too quick to be captured by a STOP process (Logan and Cowan, 1984). This mirrors descriptions of post-error slowing observed in choice reaction time tasks (Rabbitt, 1966). Interestingly, some studies have also documented similar increases in response latency following successfully inhibited trials (Emeric et al., 2007; Schachar et al., 2004; Verbruggen et al., 2008). As performance is correct in these trials, this observation is not reflective of error-correction mechanisms but still represents a shift in priority toward stopping following a stop-signal (Verbruggen and Logan, 2009b). This shift towards stopping, even on go trials, may be resultant of binding the action of stopping to the go stimuli which persists over several trials (Verbruggen and Logan, 2009b). Response latencies are also slower under instances in which the stop-signal is common (Emeric et al., 2007). Under such instances, the balance between the two tasks tilts away from a predisposition to go and towards a bias to stop or wait. Although where such a biasing single is generated is unknown, adjustments

in response latency are achieved by manipulating the onset of activity accumulation in motor areas (Pouget et al., 2011). Further research is needed to understand the mechanisms through which this is achieved.

Of course, performance adjustments need not be implemented through changes in the go response latency alone. In the stop-signal task, subjects may also adjust the time to inhibit a movement (SSRT). Unfortunately, insight into this is limited given that SSRT is calculated on a session-wise, rather than trial-wise, basis. Nevertheless, on a broader scale, we can observe that SSRT increases with stop-signal delay, which may reflect differences in the expectancy of the stop-signal (Logan and Cowan, 1984).

1.6 Generating and controlling action

1.6.1 Neural mechanisms of action

In this section, I will provide an overview of the literature describing the neural correlates of oculomotor and, to a lesser extent, given the focus of this dissertation, skeletomotor movements.

1.6.1.1 Oculomotor actions

Oculomotor movements are integral for allowing us to actively view the world around us. There are four basic types of eye movements: saccades, smooth pursuit movements, vergence movements, and vestibulo-ocular movements. In this dissertation, I will focus on saccades.

1.6.1.1.1 Defining saccadic eye movements

Saccades are rapid eye movements that result in the shift of gaze from one part of the visual field to another. This is typically done to foveate on objects of interest. Saccades can be horizontal, vertical, or oblique, and can be both voluntarily (e.g., looking towards a target) or involuntarily (e.g., the fast phase of nystagmus) executed. Importantly, saccades are ballistic - once the command to execute them is generated, they cannot be stopped or changed. The properties of saccades can be described through several key parameters: amplitude, duration, and velocity. Amplitude can be considered as the magnitude of the saccade and is typically measured in degrees or minutes of arc. A higher amplitude would reflect that the eyes have moved further in the saccade. Duration reflects the time taken to complete the saccade: from its initiation to reaching its endpoint. Finally, velocity reflects how fast the saccade is. These parameters have well-defined, reliable, and stereotyped relationships in healthy oculomotor systems. The saccades' duration and peak velocity increase as the saccades' amplitude increases (Yarbus, 1967; Robinson, 1964; Fuchs, 1967; Bahill et al., 1975). This increase is linear for short, low-amplitude saccades, while the speed asymptotically approaches a saturated value for large saccades. This relationship is referred to as the main sequence.

In this section, I describe the properties of subcortical and cortical areas considered key structures in the

generation of saccades. These areas are thought to form a hierarchy: from the frontal eye field performing the higher-order computations required for goal-directed visuomotor transformations to the brainstem saccade generator using neural signals to contract specific extraocular muscles. Furthermore, although these areas are discussed separately here, they are considered to form a more extensive network to facilitate and suppress goal-directed saccades. Finally, whilst valuable insights into these systems have also come from studies of other animal models, including humans (Leigh and Zee, 2015), the literature discussed here will primarily focus on findings from the macaque model. These findings are summarized visually in Figure 1.2.

1.6.1.1.2 Frontal Eye Field

The frontal eye field (FEF) is a cortical area considered to be a key node in the oculomotor network, important for translating sensory stimuli into motor plans (Bruce and Goldberg, 1985; Schall, 1991a). In the macaque, FEF corresponds to area 8 and is localized to the rostral bank of the arcuate sulcus. In humans, however, FEF is identified with area 6 and located at the intersection of the superior precentral sulcus and the superior frontal sulcus (Koyama et al., 2004; Amiez and Petrides, 2009).

FEF receives dense input from visual areas, including V4, lateral intraparietal sulcus (LIP), middle temporal area (MT), and area TE (Schall et al., 1995). With this dense input from lower visual areas, many FEF neurons respond strongly to visual stimuli. These visual neurons have short-latency responses - around 60 ms (Pouget et al., 2005) - and have large receptive fields (Bruce and Goldberg, 1985). These receptive fields are within the contralateral hemifield and topographically organized such that neurons in the lateral aspect of FEF represent the central visual field, and those in the medial aspect represent the periphery. Notably, visual activity FEF represents a saliency map: the combination of multiple feature maps (i.e., color, motion) to generate an overarching map that highlights locations for further processing. This ability may be afforded through topographic input from areas in both dorsal and ventral visual processing streams (Schall et al., 1995). In line with this proposal, a wealth of previous evidence has demonstrated that visual FEF neurons can exhibit a differential response to a target over a distractor within their receptive fields - a process referred to as target selection (Schall and Hanes, 1993; Thompson et al., 1996; Lowe and Schall, 2018; Monosov and Thompson, 2009). In addition, other studies have demonstrated that visual FEF neurons can contribute to discrimination in categorization tasks (Ferrera et al., 2009; Ding and Gold, 2012; Mante et al., 2013). Collectively, this ability to distinguish aspects of our visual field that require our attention provides evidence for FEFs proposed role in generating a saliency map.

Such signals are also vital in the visuomotor transformation process - once derived, saliency signals can be used to guide actions. In addition to its visual processing capacity, the frontal eye field is perhaps most notable for its role in producing saccadic eye movements. Early work using neurophysiological techniques

described neurons that discharge specifically before and during saccades (Bruce and Goldberg, 1985). These neurons are referred to as movement neurons, supporting observations of neurons that fire during saccadic eye movements, low-intensity micro-stimulation of FEF can elicit saccades (Bruce et al., 1985). Furthermore, reversible inactivation and ablation of FEF have been shown to impair, but not completely stop, saccade production (Dias et al., 1995; Dias and Segraves, 1999; Sommer and Tehovnik, 1997; Schiller and Chou, 1998; Schiller et al., 1980). Movement neurons are typically localized within the deeper cortical layers of FEF. Concordantly, anatomical studies have shown that neurons in FEF project directly to brainstem nuclei responsible for generating saccades and also terminate on neurons in the superior colliculus (Stanton et al., 1988; Cerkevich et al., 2014). Like the visual field map, saccade properties can also be mapped across FEF. Saccade amplitudes are also mapped along the medial-lateral axis of the arcuate sulcus with shorter saccades represented more laterally and longer saccades more medially (Bruce et al., 1985).

Contrasting these neurons that modulate around the saccade, other reports have described neurons in FEF that suppress their firing rate around the time of the saccade and have elevated firing rates during periods of fixation. These have been labeled as fixation neurons. Although less well described than movement neurons, several studies have directly examined the properties of this category of neurons. Using the countermanding task, Hanes et al. (1998) demonstrated that a population of neurons in FEF exhibited a declining discharge rate before saccade initiation. This observation corresponds with other descriptions of activity in FEF in which neurons are more active during periods of sustained fixation (Izawa et al., 2009). Interestingly, this group of neurons demonstrated a significant and rapid increase in firing rate when a planned saccade was inhibited. The implications of this observation in the context of response inhibition will be discussed further in the next section. In addition to these observations, causal manipulations of some sites in FEF through microstimulation have led to the suppression of visually guided saccade generation (Burman and Bruce, 1997; Izawa et al., 2004a,b). Fixation neurons are observed across FEF but are seemingly concentrated in a small area within the caudal part of the arcuate gyrus facing the inferior arcuate sulcus (Izawa et al., 2009). However, they appear to have a significantly lower incidence than movement neurons. One argument for this is that movement neurons are large pyramidal cells, while fixation neurons may have smaller cell bodies; as such, they may be less likely to be detected through the single-electrode neurophysiological techniques employed during the time they were first described. An alternative possibility is that fixation neurons may be inhibitory interneurons, acting through direct GABAergic suppression of the larger pyramidal neurons facilitating movement. These neurons are smaller and unlikely to be detected through classic neurophysiological techniques.

1.6.1.1.3 Basal Ganglia

The basal ganglia are a set of subcortical nuclei located at the base of the cortex. These nuclei are categorized into four principal volumes: the striatum, the internal and external segments of the globus pallidus, the subthalamic nucleus, and the substantia nigra.

Inputs to the basal ganglia are through the striatum, which consists of two structures: the caudate and the putamen. Cortical oculomotor systems, such as the frontal eye field, project mainly to the caudate, whilst the putamen receives input mostly from skeletomotor structures (Griggs et al., 2017; Parthasarathy et al., 1992). As this section focuses on oculomotor systems, I will focus here on the caudate's properties. Caudate can be compartmentalized into two sections: the head and the tail. Projections from FEF showed some weak topography into the caudate; the dorsal FEF almost exclusively projected to the head of the caudate nucleus, whilst projections from the ventral FEF terminated in both the caudate head and tail. Neurons from FEF projecting to the tail were found in layers III and V, whilst neurons projecting to the head of the caudate were primarily located in layer V (Griggs et al., 2017).

A long line of research has demonstrated that caudate activity is mediated by reward and value, and the dynamics of this activity are different between the head and the tail of the caudate. Neurons in the caudate head respond to different visual objects (Kim and Hikosaka, 2013). These neurons vary their activity based on the short-term association between the object and its value, increasing firing rates to cues that predict reward (Kawagoe et al., 1998). Changes in the magnitude of these neurons can follow changes in reward mapping; if a previously high-rewarded target changes to become a low-reward target, the neurons will alter their firing rate from high to low within a couple of repetitions (Kawagoe et al., 1998). Accordingly, these neurons are thought to flexibly guide saccadic eye movements to valuable stimuli within the environment. Neurons in the caudate tail also respond differentially to different visual objects (Caan et al., 1984; Yamamoto et al., 2012). This is most likely resultant of its dense input from areas in the inferior temporal cortex (Yeterian and Hoesen, 1978). In contrast to those in the head, neurons in the caudate tail are moderated by stable values, demonstrating higher firing rates for high-value objects that retain the same value over time (Kim et al., 2013). These observations have laid the foundations for the proposal that the head and tail of the caudate are considered to have different functional properties: neurons in the head are thought to elicit deliberate, planned saccades, whilst those in the tail are thought to elicit automatic responses.

The basal ganglia exert influence over the rest of the brain through two output nuclei: the internal segment of the globus pallidus (GPe) and the substantia nigra pars reticulata (SNpr). Through their projections, these areas exert influence through tonic neural inhibition: a process in which neural activity (such as spiking) is suppressed, blocked, or restricted. Although similar in their general function, these two areas are differenti-

ated in their efferent projections. The GPe sends projections directly to the thalamus, which can lead to the inhibition of skeletomotor movements. The SNpr instead sends projections to the superior colliculus, which can lead to the inhibition of oculomotor movements. Again, as this dissertation focuses on the brain systems in the context of saccadic eye movements, here I will present literature detailing the properties of SNpr during saccade generation. The principles presented here can also apply to skeletomotor commands.

Reports in macaques describe neurons in SNpr that decrease their neuronal discharge ahead of the saccade onset (Hikosaka and Wurtz, 1983a,b; Basso and Wurtz, 2002). Broadly, these results show SNpr neurons are more active during fixation periods and reduce their firing rate in the $\sim$50 ms before saccade initiation. These observations support the notion that SNpr acts as a gate for generating saccades through strong inhibitory inputs to the superior colliculus. Neurophysiological evidence has shown that microstimulation of the colliculus antidromically activates neurons in SNpr in macaques (Hikosaka and Wurtz, 1983c). Furthermore, reversible inactivation of SNpr through muscimol, a GABA agonist, produces saccades in monkeys (Hikosaka and Wurtz, 1985). As such, if the activity of SNpr is high, then activity in SC is inhibited. Alternatively, if activity in SNpr is low, then activity in SC can increase, allowing for the generation of saccades. Interestingly, more recent reports show that SNpr neurons may also show increases in their firing rate (Handel and Glimcher, 1999; Shin and Sommer, 2010). In the context of the above circuit mechanism, observations that SNpr neurons increase in activity of before saccades, are unexpected. One explanation for this observation is that SNpr may be involved in saccade selection (Basso and Wurtz, 2002); as such, SNpr may gate inhibition for saccades about to be made but also increases inhibition for saccades that would interfere with the goal. Further evidence for this hypothesis is required (Basso and Sommer, 2011).

Information can progress between the input and output nuclei of the basal ganglia through two pathways: direct and indirect. These pathways facilitate the generation and inhibition of actions, respectively. Through the direct pathway, the phasic firing of neurons within the caudate results in the direct inhibition of the SNpr. Relevant to the generation of eye movements, this inhibition of SNpr then releases inhibition over the superior colliculus. This allows for neurons in SC that can generate saccades to become active. In contrast, movement can be inhibited through the indirect pathway. Through this pathway, activity from the caudate can inhibit the activity of the external segment of the globus pallidus (GPe). As GPe has inhibitory projections to the subthalamic nucleus (STN), the inhibition of GPe neurons by the caudate results in the disinhibition of STN. This disinhibition allows for STN to increase its firing rate. As the projection from the STN to the SNpr is excitatory, the increase in activity in STN results in the rise of activity in SNpr, which in turn exerts inhibition over the superior colliculus. The balance between the direct and indirect pathways (to generate or stop movements, respectively) is tuned by dopaminergic inputs from the substantia nigra pars compacta (SNpc). Neurons in the direct pathway predominantly express D_1 dopaminergic receptors, whereas neurons

in the indirect pathway express D_2 dopaminergic receptors. Resultant of this dichotomy in receptor types between the pathways, dopaminergic input exerts a dual effect on this pathway: indirect pathway D_2-neurons are inhibited by dopamine, and direct pathway D_1-neurons are stimulated. A third pathway has also been proposed, referred to as the hyperdirect pathway (Monakow et al., 1978; Nambu et al., 1997). Much like the indirect pathway, activation of the hyperdirect pathway results in the inhibition of movement. This occurs through the direct activation of the STN by areas in the frontal cortex, which in turn drives inhibition over the SC via the SNpr. This pathway is typically discussed in the context of stopping skeletomotor commands (Aron et al., 2016) and will be discussed later in this dissertation.

1.6.1.1.4 Superior colliculus

The superior colliculus (SC) is a laminated structure in the dorsal mesencephalon and is considered another central structure within the oculomotor network. It is also well described for facilitating rapid sensorimotor transformations (including tactile, auditory, and visual) for skeletomotor movements (Jay and Sparks, 1987; Kozak and Corneil, 2021); these will not be discussed within this section, as the primary focus is on the generation of saccadic eye movements.

The superior colliculus receives input from many regions in the cortex, including the visual cortex, LIP, FEF, and basal ganglia (Fries, 1984). These inputs are all excitatory and are thought to facilitate the initiation of saccades. Much like FEF, SC has both visually responsive neurons and motor neurons. These classes of neurons are segregated into separate layers: superficial and deep. In the superficial layers, neurons are visually responsive and receive inputs directly from the retina in addition to other visual areas (Fries, 1984; Perry and Cowey, 1984). These neurons demonstrate a transient increase in firing rate within 40 to 70 ms following the presentation of visual stimuli in a specific part of the visual field (Marino et al., 2008; Bell et al., 2006). In deeper layers of the SC, there are motor neurons associated with saccadic movement (Munoz and Wurtz, 1995, 1993). Movement neurons in this layer increase their firing rate before and during a contralateral saccade in a particular direction and amplitude that defines the movement field of that neuron (Sparks et al., 1976; Wurtz and Goldberg, 1972). This movement field can be considered the motor equivalent of a visual receptive field: it defines the spatial point to which the eyes will orient once the motor command has been executed. Motor neurons in the superior colliculus are also organized topographically by their movement field: saccade amplitude is organized along the anterior-posterior axis, with smaller-amplitude saccades located anteriorly and larger-amplitude saccades at the posterior (Robinson, 1972; Stanford et al., 1996). The scaling of this map is logarithmic, with more space allocated for smaller saccades. Moving medial to lateral across SC, neurons are organized for saccades with upward and downward components, respectively (Robinson, 1972; Stanford et al., 1996). In contrast to the anterior-posterior amplitude map, the mediolateral

direction map is relatively linear (Ottes et al., 1986).

Whilst the role of SC in generating saccades to a location in the visual field is well described, there is debate about the nature of the fixation-related activity. One view is that there is a separate class of neurons housed in the rostral pole of SC, independent of movement neurons, that tonically discharge during visual fixation and pause during saccadic eye movements. Stimulation of this area increases the latency of large-amplitude saccades, whilst inactivating this area reduces the latency of large-amplitude saccades and disrupts fixation (Munoz and Wurtz, 1993). It is proposed these neurons facilitate fixation by suppressing the activity of movement neurons in the caudal aspect of SC, preventing saccade generation (Munoz and Istvan, 1998). Alternative proposals instead suggest that neurons in this rostral part of SC aren't directly related to fixation and instead contribute to very-small amplitude saccades (known as micro-saccades) near the point of fixation (Hafed et al., 2008; Hafed and Krauzlis, 2012; Krauzlis et al., 2017). This view proposes that successful fixation is achieved by the eyes producing these micro-saccades, which are resultant of an equilibrium of activation at the rostral-most aspect of the saccade map in SC. However, this view has been critiqued as evidence has demonstrated the prevalence of microsaccades is also reduced in trials in which larger saccadic eye movements are inhibited (Godlove and Schall, 2016).

1.6.1.1.5 Thalamus

The thalamus is a structure located in the dorsal aspect of the diencephalon, composed of several different nuclei. A collection of nuclei in the anterior thalamus is considered the oculomotor thalamus. These nuclei include the anterior group of the intralaminar nuclei and parts of the ventrolateral, ventroanterior, and mediodorsal nuclei that border the internal medullary lamina. These nuclei are densely and reciprocally connected with oculomotor centers within the cortex (such as FEF) and subcortex (such as the basal ganglia and superior colliculus). Given this dense connectivity with other oculomotor structures, neurons in these nuclei also share common patterns of saccade-related activity. Previous literature has demonstrated that neurons in these areas have activity related to saccades, fixation, and eye or gaze position in space (Schlag and Schlag-Rey, 1984, 1986; Tanibuchi and Goldman-Rakic, 2003; Wyder et al., 2003; Sommer and Wurtz, 2004b; Watanabe and Funahashi, 2004).

Of the oculomotor nuclei, the mediodorsal nucleus (MD) has received particular attention for its role in saccade generation. This is likely due to its dense connectivity to the frontal eye field and the superior colliculus (Benevento and Fallon, 1975; Goldman-Rakic and Porrino, 1985). Neurons within MD increase their firing rate during visual stimuli presentation and before a saccade is generated (Sommer and Wurtz, 2002, 2004a,b; Tanibuchi and Goldman-Rakic, 2003; Watanabe and Funahashi, 2004). This activity is greater for stimuli and saccades into the contralateral hemifield.

The mediodorsal nucleus, however, is most widely described for its role in corollary discharge (Sommer and Wurtz, 2002, 2004a,b, 2006). This is also referred to as an efference copy and has been well-reviewed previously (Subramanian et al., 2019) but is briefly discussed here. The central idea of corollary discharge is that a copy of a motor command is sent to other parts of the brain to inform it of the subsequent movement. This signal does not generate the movement and is thought to be used for planning following movements and for interpreting changes to the sensory signal that may result from the action. Notably, corollary discharge signals are believed to contribute to visual and motor sequence processing. This is demonstrated through the following two examples. First, as we navigate our visual fields using eye movements, we must account for this movement when interpreting the sensory signal. This is achieved through pre-saccadic remapping – a process immediately before a saccade is generated in which visual responses are shifted to the part of the visual field where the receptive field will land (Rao et al., 2016). Second, we can also gain insight from the double-step task; here, participants are presented with a sequence of two stimuli to which they are instructed to shift their gaze towards. In one condition of this task, participants must shift their gaze from the first target to the second target; without a signal conveying the properties of the first saccade, the oculomotor system could not adapt updated location and thus execute a misplaced saccade toward the previous retinal location of the second target. Interestingly, this approach has revealed deficits in corollary discharge in mental illnesses such as schizophrenia (Thakkar et al., 2015; Thakkar and Rolfs, 2019).

A system conveying the corollary discharge for saccadic eye movements has been described within a circuit including the MD, SC, and FEF. MD has projections to both SC and FEF (Benevento and Fallon, 1975; Goldman-Rakic and Porrino, 1985), making it well-situated to relay signals between the two areas. Through a series of studies, Sommer & Wurtz identified neurons within the MD that were orthodromically activated when SC was stimulated and antidromically activated when FEF was stimulated. Recordings at each step of this pathway showed neurons conveyed information about the temporal and spatial properties of the upcoming saccade. These signals began in SC, propagated through the MD, and terminated in FEF. Reversible inactivation of this pathway results in marked deficits in the corollary discharge and affects visual and motor sequence processing.

1.6.1.1.6 Brainstem Saccade Generator

The brainstem saccade generator circuit is the final stage for generating a saccade and is nested within the brainstem reticular formation (Bohlen et al., 2016; Scudder et al., 2002). Horizontal saccades are generated through circuits in the pons, and vertical saccades are generated through circuits in the midbrain. These areas are linked through the medial longitudinal fasciculus to produce saccades in all directions. Within this circuit are several classes of neurons essential for generating saccadic eye movements. Saccades are made when a

pulse of force is followed by a slide into step is applied to the extraocular muscles. This kinematic process is moderated through the activity of burst neurons, long-lead burst neurons, tonic neurons, and omnipause neurons.

Burst neurons (BNs) innervate the extraocular muscle motoneurons and provide the high-frequency burst of spikes required to move the eyes in a specific direction. This activity produces the initial pulse of force. The firing rate of these burst neurons is proportional to the velocity of the saccade, and during fixation, these neurons are silent. There are two types of burst neurons: excitatory burst neurons excite the on-direction motoneurons to contract the muscle; inhibitory burst neurons inhibit the antagonist motoneurons allowing the muscle to extend. Another group of neurons, long-lead burst neurons, also discharge a high-frequency burst of action potentials for saccades to the contralateral hemifield. In contrast to typical burst neurons, long-lead burst neurons also have a low-frequency buildup of activity before the burst occurs. It is thoughts that these neurons may be innervated by neurons from areas higher in the saccade hierarchy, such as the superior colliculus and frontal eye field.

To maintain gaze at an eccentric location (i.e., not parallel to the direction of the head), a step of force is necessary. This step of force provides the muscle tension needed to hold the eyes in position after the initial pulse. This is achieved through tonic neurons, which have firing rates proportional to the angle of the eye in orbit; when the eye is held at a larger eccentricity, tonic neurons have higher firing rates. These tonic neurons are housed in the nucleus prepositus hypoglossi and medial vestibular nucleus and, like burst neurons, also synapse onto extraocular muscle motoneurons. This set of neurons has also been referred to as a neural integrator; they can integrate the activity of burst neurons, reflecting velocity, to moderate their firing rate to produce an appropriate force step, holding the eye in place.

Finally, burst neurons are subject to inhibition from omnipause neurons (OPNs), which discharge tonically during all fixations and pauses. The omnipause neurons are located in the nucleus raphe interpositus and are considered to serve as a gating mechanism for saccade generation. When a saccade is generated, these neurons reduce their firing rate, becoming silent during the movement. This allows the appropriate pools of burst neurons to be activated. Once a saccade is completed, OPNs increase their firing rates and become active again, exerting inhibition over the burst neurons.

1.6.1.2 Skeletomotor actions

Whilst oculomotor movements allow us to view the world around us, skeletomotor movements are integral for allowing us to actively interact with the world around us.

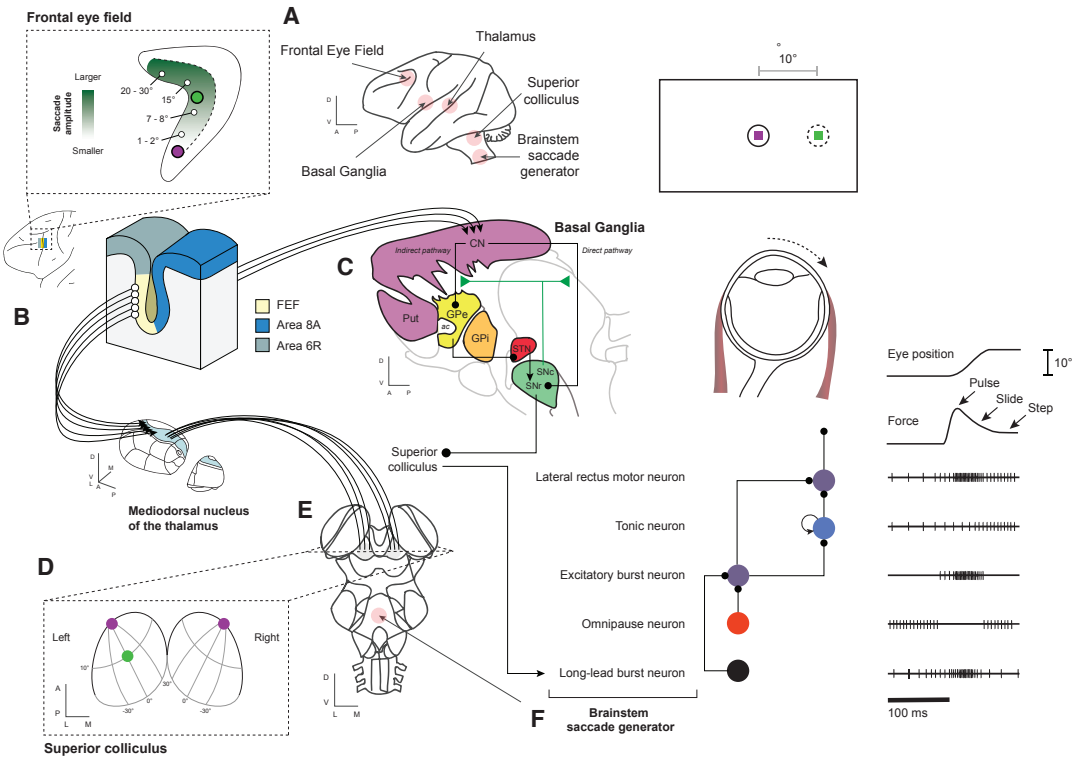

Figure 1.2: **Neural systems contributing to movement generation and control**. This figure represents several areas that have been demonstrated to contribute to the direct control of saccades, including the frontal eye field, basal ganglia, thalamus, superior colliculus, and brainstem saccade generator **(A)**. These areas are discussed in detail in the text. Frontal eye field is an area in the frontal cortex important for translating sensory stimuli into motor plans, located in the rostral bank of the arcuate sulcus **(B)**. It demonstrates dense connectivity to the basal ganglia **(C)** and the superior colliculus **(E)** through the mediodorsal nucleus of the thalamus **(D)**. Like FEF, the superior colliculus also neurons that can contribute to gaze-shifting and gaze-holding. Previous work has demonstrated that this activity can be modulated through the basal ganglia, which may act as a gate. The superior colliculus is directly connected to the saccade brainstem generator **(F)** which directly controls the extraocular muscles to move the eyes. Neurons in the brainstem saccade generator clearly map on to distinct kinematic processes during eye movements.

1.6.1.2.1 Defining skeletomotor movements

Whilst the macaque stop-signal literature has built upon the affordances offered by our knowledge of the saccade generation systems to probe the mechanisms of response inhibition, human stop-signal research typically requires participants to respond using manual movements (although see Scangos and Stuphorn (2010), Mirabella et al. (2011), Pani et al. (2022), Giamundo et al. (2021), Giarrocco et al. (2021), for examples in macaque). These movements could be a reaching movement, such as moving a joystick left or right, or a button press, where a participant uses their left or right index figure to press keys. In contrast to saccades, these actions involve moving multi-segmental body parts long enough to be modified and controlled throughout their generation to a target (Georgopoulos et al., 1981). Given the non-ballistic nature of these movements, there needs to be precise co-ordination for each movement segment to control goal-directed forelimb movements successfully (Jeannerod, 1984; Jeannerod et al., 1995). With this, unlike saccades, these movements can be stopped or adjusted during their execution (Battaglia-Mayer et al., 2014).

Classically, two cortical systems for driving actions have been proposed (Passingham, 2007; Haggard, 2008; Passingham et al., 2010): a parietal-premotor circuit that generates stimulus-driven actions and a medial-frontal circuit that generates voluntary actions. In this section, I will briefly describe the cortical and subcortical areas within these networks that contribute to the generation of skeletomotor movements. Just like the oculomotor network, these areas are thought to form a hierarchy: the posterior parietal cortex and supplementary motor cortex contribute to higher-order movement planning, which then guides the activation of motor plans in the motor cortex. From the primary motor cortex, motor commands from the brain are sent to the muscles through the corticospinal tract. The corticospinal tract is a descending white matter tract that extends from the motor cortex and progresses through the internal capsule to the medulla oblongata, where most fibers then decussate to the contralateral side. This bundle of fibers then goes down the spinal cord before synapsing with (lower) motor neurons in the anterior horn of the spinal cord. This lower motor neuron relays the command from the descending upper motor neuron, synapsing directly onto the effector's muscle. Activation of this neuron results in the contraction of the muscle, generating the observable movement. Like the oculomotor network section, I will primarily discuss findings from the macaque model here but will also provide further observations that have resulted from studies using the stop-signal task in humans.

1.6.1.2.2 Posterior Parietal Cortex

The posterior parietal cortex (PPC) is composed of several cytoarchitecturally distinct areas, including area 7a, area 7b, the medial superior temporal area (MST), the anterior intraparietal area (AIP), the lateral intraparietal area (LIP), and the ventral intraparietal area (VIP). Areas in the PPC receive dense connections from visual, auditory, somatosensory, and vestibular systems, and have projections to the frontal, premotor, and

motor cortices (Cavada and Goldman-Rakic, 1989a,b; Wang, 2020). Given the wide range of inputs from sensory areas, PPC is thought to combine the signals into a multimodal map of our position within 3D space that is used to guide spatial attention and spatial awareness. These signals can then be used to guide actions (Andersen and Buneo, 2002).

Neurons in areas of PPC also modulate with movement generation. Specifically, neurons within these areas modulate for more complex movements such as reaching and grasping (Mountcastle et al., 1975; Lynch et al., 1977). Such signals are helpful for removing the need to coordinate movements at a finer resolution. For example, neurons within the AIP encode entire hand shapes and configurations required for grasping objects (Murata et al., 2000). Other work has demonstrated that the movement-related activity of neurons in PPC can be moderated by direction (Eskandar and Assad, 1999; Grefkes and Fink, 2005) and is not lateralized (Chang et al., 2008).

A notable point, however, is that activity in PPC is elevated before the execution of a planned movement. In a series of studies, monkeys were cued to make an arm movement toward a target but to respond only after a second cue was presented. Neurons in the PPC demonstrated elevated activity during the delay period (Snyder et al., 1997). This sustained activity during a planned movement contrasts with other areas, such as M1, where movement-related activity is only observed immediately before, or during, the movement, and not observed during delay periods (Crammond and Kalaska, 2000). This observation led to the premise that PPC may reflect an intention map – a guide for executing complex goal-orientated movements.

1.6.1.2.3 Supplementary motor complex

The supplementary motor complex (SMC) is comprised of the supplementary motor area (SMA), pre-supplementary motor area (pre-SMA), and supplementary eye field (SEF). In humans, these areas are located in the dorso-medial frontal cortex, anterior to the primary motor cortex. Areas in the supplementary motor cortex have direct and substantial downstream connections to the basal ganglia and corticospinal tract (Dum and Strick, 1996; Takada et al., 1998). Connections to the descending movement systems are denser from SMA than pre-SMA (He et al., 1995; Wise, 1996; Dum and Strick, 1991), and neurons from SMA may have direct synapses onto motor neurons (Dum and Strick, 1996). Interestingly, whilst SMA has strong reciprocal connections with the primary motor cortex (M1), pre-SMA demonstrates denser connectivity to areas in the frontal cortex (Luppino et al., 1993; Lu et al., 1994).

Functionally, converging evidence from fMRI, EEG, and neurophysiology studies has shown that regions in the SMC are active before the generation of a movement. Single-unit recordings from macaques have demonstrated that neurons in SMC discharge prior to the execution of hand or foot movements (Brinkman and Porter, 1979; Tanji and Kurata, 1982). Similarly, microstimulation of SMA can elicit complex movements

and the urge to move (Penfield, 1950). In humans, such activity is also reflected through an event-related potential (ERP) known as the readiness potential (or *Bereitschaftspotential*) (Deecke and Kornhuber, 1978; Libet et al., 1983). This ERP is centered over the medial frontal cortex, superior to the SMC. The magnitude of this potential is greater prior to self-initiated movements compared to movements that were externally cued (Jahanshahi et al., 1995). This observation has also been observed through fMRI, in which BOLD activity is greater in the pre-SMA when participants are free to choose their actions than when they are instructed by external signals (Nachev et al., 2005; Deiber et al., 1999).

Although these observations suggest that the SMC may be integral in generating voluntary and intentional movements, alternative proposals indicate that the area may contribute to generating action sequences or learning instead. Indeed, neurophysiology studies have described neurons with activity patterns supporting the sequential generation of movement (Tanji and Shima, 1994; Shima and Tanji, 2000). For example, these neurons may only respond before a particular movement is made but only after another has been completed or change their firing rate depending on the position of a movement within an order. Moreover, the activity of these neurons can be modulated through learning these sequences (Nakamura et al., 1998; Sakai et al., 1999) or by learning stimulus-response mappings (Chen and Wise, 1995).

Interestingly, in addition to the literature describing its role in voluntary movement, the supplementary motor complex has also been widely described as a hub for monitoring movements and actions. This will be discussed in the following sections.

1.6.1.2.4 Primary motor complex

The primary motor cortex (M1) corresponds to area 4 as defined by Brodmann and is located on the precentral gyrus, rostral to the central sulcus. Although the caudal border of M1 is well defined at the fundus of the central sulcus, the rostral border is difficult to define as the transition from area 4 to 6 is gradual. Although initially proposed as a single homogenous region, literature has suggested that the anterior and posterior aspects of area 4 should be considered as separate sections of M1 (Geyer et al., 1996). Compared to other cortical regions, M1 is also classically defined by a high density of Betz cells; these cells are a type of giant pyramidal cell (up to 120 μm in diameter) located in layer V of M1 (Meyer, 1987). Anatomical tracer studies in the monkey have shown that M1 receives direct input from the parietal cortex, supplementary motor cortex, and thalamus (Matelli et al., 1998, 1986, 1989; Geyer et al., 2000). Specifically, in contrast to oculomotor projections, thalamic input to M1 originates from the ventrolateral (VL) and ventral anterior (VA) nuclei.

Functionally, most neurons in the motor cortex increase their activity with the generation of movement and scale their activity with the force of this movement (Georgopoulos et al., 1982; Cheney and Fetz, 1980). Like the generation of saccades, neurons in M1 may also demonstrate a rise-to-threshold mechanism (Lecas

et al., 1986; Riehle and Requin, 1989), although recent work proposes M1 activity may be best represented by dynamic system models (Shenoy et al., 2013). During arm movements, recordings from macaques showed that firing rates peak just before movement onset and return to baseline following the completion of the movement (Georgopoulos et al., 1982). Other work has also shown neurons in the macaque motor cortex are activated in a proximal to distal sequence: from muscles in the shoulder, to the elbow, to the wrist, to the finger (Murphy et al., 1985).

Movements may be stopped through increased neural inhibition over this motor network. This notion is supported by studies observing a reduction in corticospinal excitability (CSE) during inhibition – a quantity derived by applying transcranial magnetic stimulation to specific parts of M1 and measuring the magnitude of the deflection in the motor evoked potential. Lower CSE reflects a lower level of cortical activity, and higher values reflect a greater level of activity. When a planned movement is inhibited, CSE is reduced (van den Wildenberg et al., 2010; Greenhouse et al., 2012). This reduction also occurs when presented with surprising stimuli (Wessel and Aron, 2013; Dutra et al., 2018) and is also observed in non-task-relevant effectors (Wessel et al., 2013). Such evidence has been used to support the global-stopping/pause-then-cancel theories of stopping, discussed later in this chapter.

An alternative perspective proposes that areas within the motor cortex may be responsible for explicitly suppressing movements. In Penfield's earlier studies of motor cortex, he identified several regions that, when stimulated, led to movement inhibition and relaxation and reports of a sense of focalized paralysis and numbness (Jasper and Penfield, 1949). This observation has since been supported by numerous other studies defining "negative motor areas" (Luders et al., 1995; Filevich et al., 2012). The juxtaposition of these findings with others highlighting M1 in generating movement has been previously discussed (Ebbesen and Brecht, 2017).

1.6.1.2.5 Basal Ganglia

The main anatomy of the basal ganglia system has been described in the previous section in the context of oculomotor commands and, as such, will not be described again here. However, several key distinctions should be noted. First, whilst oculomotor input into the basal ganglia is most dense in the caudate, signals from cortical skeletomotor areas (such as those described below) are most dense in the putamen. Second, whilst the basal ganglia system exerts influence over the oculomotor network through the SNpc, influence over the skeletomotor network occurs through the GPe exerting inhibition over the thalamus, which removes the drive to the motor cortex.

Physiologically, movements are facilitated through the direct pathway and inhibited through the indirect pathway (Mink, 1996). Input into these pathways for skeletomotor movements is through the putamen. Puta-

men neurons typically have a low baseline firing rate which significantly increases with movement (Crutcher and DeLong, 1984; DeLong, 1973). The movement neurons are somatotopically organized in the middle and caudal aspect of the putamen (Flaherty and Graybiel, 1991): face representations are found more ventromedial, whilst leg movements are localized in the dorsolateral aspect. This observation matches the pattern of anatomical connections between the somatosensory and primary motor cortex (Flaherty and Graybiel, 1991). Functionally, putamen neurons have demonstrated variability in their pattern of activity: some show time-locked movement activity, some modulate following a cue to move but before a movement is executed, and some moderate their firing rate based on the context of prior movements and task conditions (Deffains et al., 2010). Further work has demonstrated that the firing rate of these neurons can be moderated by movement parameters, such as the direction and pattern of muscle activity (Crutcher and Alexander, 1990; Crutcher and DeLong, 1984). Interestingly, other studies have noted the onset of many putamen motor responses occurs late, following the execution of a movement (Crutcher and Alexander, 1990).

The Globus Pallidus is divided into two segments: the internal (GPi) and external (GPe). Through the indirect pathway, putamen neurons inhibit neurons in the external segment of the Globus Pallidus. Through the direct pathway, putamen neurons inhibit neurons in the internal segment of the Globus Pallidus. In contrast to putamen neurons, neurons in both segments of the Globus Pallidus exhibit a high, tonic, baseline firing rate (Brotchie et al., 1991; Elias et al., 2007). These segments receive similar patterns of inputs and have a similar somatotopy (Yoshida et al., 1993). Given this similarity in structure and input, it is unsurprising that the functional activity of neurons in the two structures is also similar; these functional properties are discussed later in this section.

In stopping a movement, signals from the external segment of the Globus Pallidus progress through the indirect pathway to the Subthalamic Nucleus (STN). The subthalamic nucleus has been highlighted as a critical node within the basal ganglia system for the rapid inhibition of a planned response (Aron and Poldrack, 2006; Bonnevie and Zaghloul, 2019). This proposal has been particularly well employed in human and rodent models but less commonly tested in macaque models. Models of response inhibition focusing on the basal ganglia will be discussed in later sections, but here I will describe insight into the properties of activity in the STN. Like GPe neurons, neurons in STN are also tonically active, although at a lower resting firing rate (DeLong et al., 1985; Georgopoulos et al., 1983; Matsumura et al., 1992; Wichmann et al., 1994). These neurons demonstrate significant, but non-effector-specific, increases in their discharge rate for movement (Matsumura et al., 1992). This modulation occurs approximately 50 ms prior to the execution of the movement (Georgopoulos et al., 1983). Many of these neurons demonstrate sensitivity to the movement direction (Georgopoulos et al., 1983), but it is unclear whether they encode other parameters of the movement. In addition to signals from the putamen (via the indirect pathway), STN also receives short-latency, excitatory

input from cortical areas (Monakow et al., 1978; Haynes and Haber, 2013; Chen et al., 2020) which appear to be somatotopically organized (Afsharpour, 1985; Monakow et al., 1978). This direct input to the STN can provide a mechanism through which cortical areas can bypass the longer latency, multi-synaptic input to the basal ganglia. This input can then drive a rapid inhibitory response to stop movement. This pathway is referred to as the hyperdirect pathway.

Several studies have used the stop-signal task (or a derivative) to probe this circuit further. As outlined in the previous section, this task affords clear criteria for assigning whether neural activity can contribute to generating or inhibiting a movement. Neurons in STN don't seem to meet these criteria. Work by Schmidt et al. (2013) in rats demonstrated that, although STN neurons do modulate in response to the stop-signal, this increase occurred in both canceled and non-canceled trials. This activity peaked exceedingly quickly ($\sim$15 ms) after the stop-signal presentation. Similar observations were also observed in later studies in macaques (Pasquereau and Turner, 2017). Collectively, these observations may reflect that STN may just reflect the sensory cue to stop rather than the process of response inhibition itself. However, this is difficult to reconcile with in-vivo human neurophysiology studies that have found STN neurons that are stopping-selective: they modulated only during successful inhibition (Bastin et al., 2014; Benis et al., 2016). This distinction warrants further consideration.

Recent work has broken down the contribution of the STN during stopping down further. Although the cytoarchitecture of the STN appears homogeneous, functional and anatomical evidence suggests it can be divided into its dorsal and ventral parts (Yelnik and Percheron, 1979; DeLong et al., 1985; Iwamuro et al., 2017). Whilst the dorsal STN receives inputs from the motor cortex, the ventral STN receives greater input from the frontal cortex (Haynes and Haber, 2013). Contrary to popular hypotheses from human neuroimaging studies, Mosher et al. (2021) demonstrated only a small proportion of individual neurons in STN demonstrated responsiveness to a stop-signal (11%), and an even smaller proportion (2%) modulated only on successfully inhibited trials. They also demonstrated that neurons in STN modulated for movements. These neurons increased their firing rate prior to a movement but did not peak until after the movement was initiated. Reflecting the topography of anatomical projections, stopping neurons were observed mainly in the ventral aspect of STN, which receives input from the frontal cortex, whilst movement neurons were located more dorsal, which receives more input from motor cortices.

Whilst GPi neurons are inhibited by neurons in the putamen through the direct pathway, they are also directly driven by neurons in the STN through the indirect or hyperdirect pathway. Neurons from GPi send direct GABAergic projections to the thalamus. Thus, excitation of GPi through the indirect pathway results in greater net inhibition over the thalamus and removal of the drive to generate movement; alternatively, inhibition of GPi through the indirect pathway results in the disinhibition of the thalamus, promoting movement. In

generating a movement, approximately 70% of GPi neurons increase their activity whilst the remaining 30% decrease their activity (Anderson and Horak, 1985; Brotchie et al., 1991; Georgopoulos et al., 1983; Mitchell et al., 1987). Such an observation is consistent with the notion that GPi maintains tonic inhibition over most of the thalamus but releases inhibition other another smaller part. Note, these observations are similar to those observed between the SNpr and SC. In comparison to other motor areas, GPi activity does not encode features of movement such as joint position, force, amplitude, or velocity (Brotchie et al., 1991). The onset of most activity in GPi occurs later than that observed in EMG but before the actual execution of a movement (Anderson and Horak, 1985; Brotchie et al., 1991; Georgopoulos et al., 1983; Mitchell et al., 1987). However, there is evidence that some neurons in GPi may also demonstrate activity that is not time-locked to the movement but modulate in preparation (Nambu et al., 1990).

Broadly, we can consider that skeletomotor actions can be initiated through the direct pathway and stopped through the indirect and hyperdirect pathways. Whilst oculomotor movements drive the basal ganglia through the striatum, skeletomotor movements drive the basal ganglia through the putamen. Influence over motor systems is exerted by the SNpr for oculomotor movements and the GPi for skeletomotor movements. These areas result in the inhibition or disinhibition of the relevant motor areas that execute the response. Interestingly, human literature has probed an additional pathway within the basal ganglia – the hyperdirect pathway – in which the frontal cortex can drive inhibition over planned movements. Whilst the properties of the pathway have received supporting evidence in humans, this is not well-tested in macaques.

1.6.2 Neurocognitive models of response inhibition

Whilst knowing the properties of neurons in a particular area is unquestionably essential for developing our understanding of the brain, the primary aim of cognitive neuroscience is to understand how complex cognitive processes can be instantiated through these neural circuits. The connection between these two separate features can be formalized through the concept of a linking proposition - "a claim that a particular mapping occurs, or particular mapping principle applies, between perceptual and physiological states" (Teller, 1984). These linking propositions identify the areas in which theory and neural data intersect, specifying the aspects of the data that are relevant to the theory (Teller, 1984).

1.6.2.1 Accumulator models

The generation of such linking propositions has been significantly advanced through developments within the mathematical psychology and neuroscience literature over the past few decades (Schall, 2019). In the context of the work presented within this dissertation, the development of evidence accumulator models has yielded numerous insights into the neurocognitive mechanisms of decision-making and action.

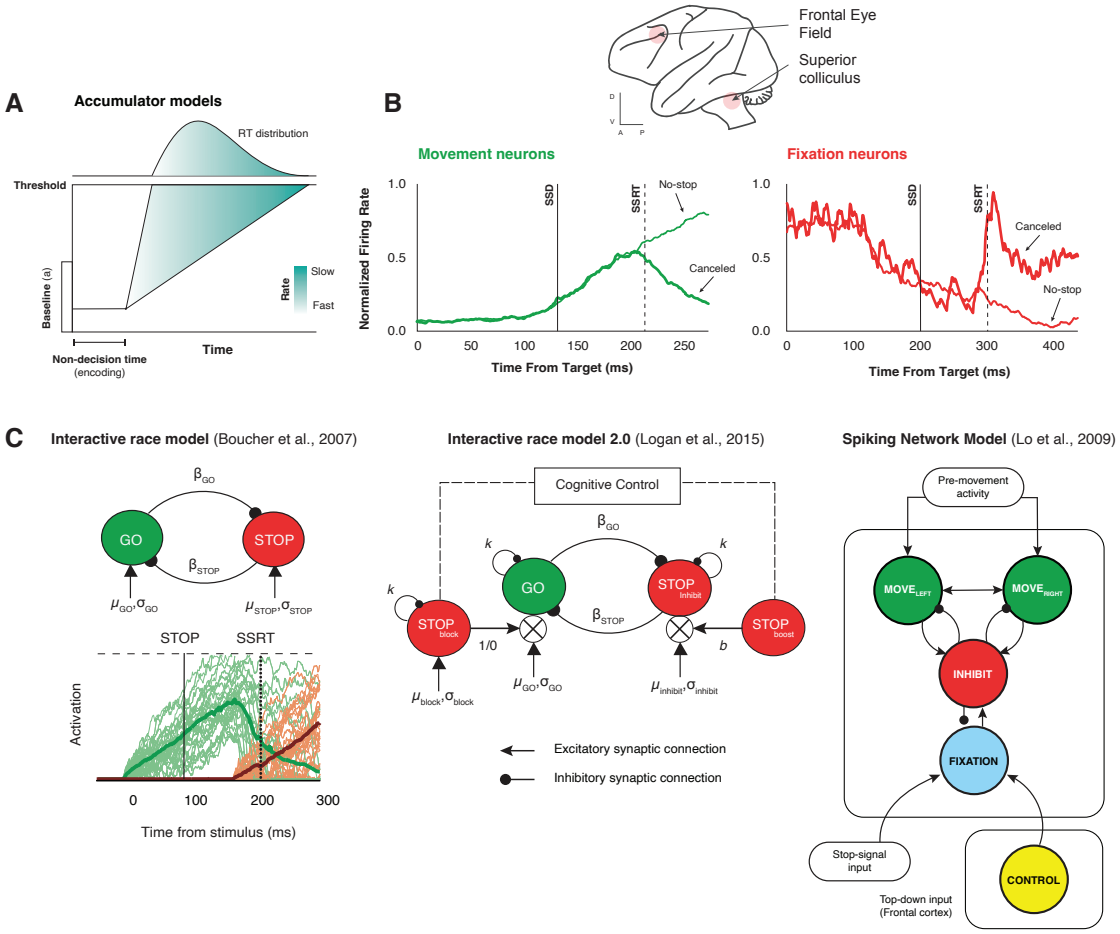

Figure 1.3: **Neurocognitive models of response inhibition**. **A**. Schematic of an accumulator models. Accumulator models integrate evidence for a decision over time. Most models are described by four parameters: baseline, threshold, onset, and accumulation rate. Here we show a fixed non-decision time during which the stimuli are encoded, and activity stays at the baseline level. After this time, activity accumulates towards a threshold. The time at which accumulation hits the threshold determines the response latency. Here we depict faster rates of accumulation results in quicker response latency (lighter cyan) whilst slow accumulation results in slower response latencies (dark cyan). **B**. Neural activity has accumulator-like dynamics. Movement neurons in oculomotor structures like FEF and SC demonstrate ramping towards a threshold following a period of baseline stability. If inhibited, this activity stops accumulating and returns to a baseline level. Fixation neurons demonstrate the opposite: they decrease their firing rate following target onset and demonstrate a sharp increase in activity when a movement is inhibited. **C**. Graphic representations of neurocognitive models describing stopping. These models instantiate the independent race model in a set of neurons that facilitate movement (green) or stop movement (red). These are described more in text.

Accumulator models (which include linear ballistic accumulators and drift-diffusion models, for example) assume evidence for a particular decision is integrated over time. Most models are described by four parameters: baseline, threshold, onset, and accumulation rat (Figure 1.3A). Accumulation starts from a baseline level of activity and finishes after reaching a defined threshold. Before this accumulation, it's proposed that there is a fixed non-decision time during which the stimuli are encoded, and the inhibition over the system is gently released to allow this accumulation to occur. During this period, no accumulation occurs. However, after this non-decision time, accumulation starts ramping from the baseline toward the threshold at a particular rate. The rate parameter determines the speed of the accumulation.

These parameters of accumulator models are thought to reflect different underlying cognitive processes (Purcell and Palmeri, 2017). For example, a speed-accuracy tradeoff may be achieved by manipulating the threshold at which activity must reach for a decision to be made (Brown and Heathcote, 2008); to contextualize this, if the threshold and rate are fixed, a higher baseline will result in shorter RT because the activity requires less time to accumulate up to the threshold. Alternatively, if the baseline and rate are fixed, a higher threshold will result in longer RT because the accumulation will require more time to accumulate up to the threshold. Furthermore, the accumulation rate may reflect confidence in a decision, as this parameter varies with prior experience and stimulus strength (Palmer et al., 2005). Finally, ambiguous or noisy sensory input prolongs the encoding time prior to the start of accumulation (Ratcliff and Smith, 2010).

In the following subsections, I will discuss several models (some of which are instantiations of such accumulator models) that attempt to explain how response inhibition can be achieved through neural circuits.

1.6.2.2 Interactive race model 1.0

Since the original inception of the independent race model (Logan and Cowan, 1984), numerous advancements in neurophysiology have occurred, allowing us to appreciate how particular brain areas change their activity during movement initiation and inhibition (Figure 1.3B). Studies have supported that GO and STOP processes can be instantiated through an interactive network of neurons that facilitate gaze-shifting (movement) or gaze-holding (fixation) (Hanes et al., 1998).

When looking at this neural data, it is clear to see that these two processes cannot remain independent: movement neurons reduce their activity when fixation neurons become active, as though they are directly inhibited (in the neural sense of the term). With this, and given the assumptions and observations of the independent race model, the identification of the GO and STOP process with movement and fixation neurons exposes a paradox: how can a network of gaze-shifting and gaze-holding neurons that interact with each other produce behavior that appears to be the outcome of a race between two independent processes? This paradox was addressed by the interactive race model 1.0 (Figure 1.3C, Boucher et al., 2007).

In this model, a simple network incorporating a go (movement) and a stop (fixation) unit was proposed. The GO unit is modeled as a stochastic accumulator that starts accumulating towards a threshold after a non-decision period. When this threshold is reached, a saccade is generated. The STOP unit is also modeled as a stochastic accumulator; this unit begins its accumulation after a non-decision period following the stop-signal delay. These two units can interact through lateral inhibition, in which activation of one unit can suppress the activity of the other. Importantly, these two units do not influence each other and are thus independent until very late in the STOP process. At this point, the STOP unit exerts robust and potent inhibition over the GO unit. If this STOP unit is active and strong enough, it can inhibit the GO unit from reaching the threshold, resulting in a canceled movement. However, if the STOP unit is too late or weak, the GO unit will reach the threshold, and an erroneous saccade will be triggered. In considering the neural mechanisms of response inhibition, the interactive race architecture prevails over the independent race architecture as it accounts for both behavioral and neural data.

1.6.2.3 Spiking neural model

Although the interactive race model provided a clear explanation for observed neural signals and behavior during stopping, it did not account for the interaction between movement and fixation units prior to stopping, and it was unclear how such a model could be implemented biophysically. Earlier work had demonstrated that biophysically plausible decision-making models could be instantiated through a recurrent network model of movement and fixation neurons (Lo and Wang, 2006). This approach was developed to explain response inhibition in the stop-signal task based on the same neurophysiological observations in the frontal eye field and superior colliculus of behaving monkeys (Figure 1.3C, Lo et al., 2009).

This spiking neural model to explain stop-signal behavior generated populations of movement neurons, fixation neurons, inhibitory interneurons and, importantly, also included a top-down control unit that turns fixation neurons on and off. This model accounts for fixation activity in the early trial period, the transition from fixation to movement, and the rise to threshold for movement activation. The model also produces a shift from fixation to movement by turning off a control unit that excites fixation units. This releases tonic inhibition on the movement neurons, allowing activity to rise to the threshold. Like other models, the time the threshold is reached produces the response latency. Compared qualitatively to the interactive race model, this model predicted the modulation of movement and fixation neurons and cancel time distributions equally well.

In addition to these stopping-related observations, the spiking neural circuit model also accounted for dynamics between movement and fixation neurons that occurred prior to the stop-signal. Specifically, it demonstrated variation in baseline activation of fixation and movement units affected the probability of suc-

cessful inhibition: successful inhibition was less likely when movement units were more active during the baseline period and more likely when fixation units were more active. This supports claims from previous literature that demonstrate reduced activity in motor areas prior to successful inhibition (Hu and Li, 2012).

Although biologically plausible, this model has two notable limitations. Firstly, the variability in RT almost entirely depends on the time the control signal is turned off. This timing of the control signal termination is determined by a Gaussian distribution whose parameters were adjusted to optimize the goodness of fit; currently, there are no explanations for the mechanisms producing this distribution. Secondly, although this model fits the data equally as well as the simpler interactive race model, it requires many more parameters and is not the most parsimonious explanation.

1.6.2.4 Interactive race model 2.0

Attempts to further reconcile the paradox between neural and behavioral observations and further separate interactions between GO and STOP processes resulted in a second iteration of the interactive race model (interactive race model 2.0; Figure 1.3C, Logan et al., 2015). Like the first interactive race model, the revised model proposed a movement and a fixation unit, both of which are stochastic accumulators. However, this revised model instead stipulated that these two units each have independent inputs (μ_{move} & μ_{fix}, respectively). Canceling a response is thought to occur in one-of-two ways: blocking inputs to a GO unit (blocked-input) or boosting inputs to a STOP unit (boosted-fixation).

This model equally fits the behavioral data and the interactive race model and provides a better description of physiological data. Importantly, this model also offered predictions about physiology and questioned the linking proposition that connected the STOP process with fixation neurons in FEF and SC, instead concluding that the STOP process must come from another node or operation outside the network. However, much like the spiking-neuron model, how and where this signal arises is unknown.

1.6.2.5 Cortico-basal-ganglia and pause-then-cancel models of stopping

As described in the previous section, parts of the basal ganglia have been implicated as crucial nodes in the going and stopping network across species. Briefly, going can be facilitated through the direct pathway, whilst stopping can be facilitated through the indirect pathway. More specifically, the standard model of basal ganglia function stipulates that movement can be initiated through the disinhibition of the thalamus via the GPi or SNpr and canceled by reinstating this inhibition. In addition to inputs from the striatum, this network can also be driven through the hyperdirect pathway. The hyperdirect pathway proposes a direct connection from the frontal cortex to the STN. Through this pathway, activity in the STN is driven by signals from the frontal cortex; this leads to a subsequent increase of activity in GPi and SNpr, exerting broad neural inhibition

over the thalamus.

The hyperdirect pathway was initially proposed as a way in which movement could be rapidly and broadly suppressed (Mink, 1996; Gillies and Willshaw, 1998; Nambu et al., 2002). In conjunction with this, recent work has proposed that suppression is not specific to a particular movement but is instead exerted over the whole motor system (global suppression). Evidence for this is demonstrated through a reduction in non-effector-related corticospinal excitability (CSE): stopping speech reduces hand CSE (Cai and Padoa-Schioppa, 2012; Wessel et al., 2016), stopping saccades reduces hand CSE (Wessel et al., 2013), and stopping hand movements reduce leg CSE (Badry et al., 2009; Majid et al., 2012). Additional studies have also demonstrated that stopping one movement can lead to considerable delays in continuing another (Coxon et al., 2007; MacDonald et al., 2014). However, this latter point is contrary to evidence from a double-step task in which latencies for compensated saccades (i.e., saccades made to the correct target following an incorrect saccade to the first-presented target) are significantly faster than saccades to no-step targets (Camalier et al., 2007).

Early applications modeling the basal ganglia circuitry for cognitive processes were described in the context of conflict - where two or more competing motor plans are simultaneously activated (Botvinick et al., 2001). Compared to contexts with no competing responses, conflicting situations require more evidence to plan and execute the correct response. With this, initial theories stipulated that this can be achieved by conflict detection signals in the frontal cortex activating the hyperdirect pathway, resulting in the inhibition of all actions. The basal ganglia system will require a greater drive from movement-related inputs into the striatum to overcome this inhibition. Through this process, would-have-been premature responses are withheld, allowing more time for correct response selection. Once this drive is greater than the inhibition exerted over the system, a movement can then be generated. Related to the accumulator models mentioned above, this can be considered as increasing the decision threshold: more evidence is required to execute a response. Collectively, these findings propose that the activation of this network is therefore thought to implement a "hold your horses" pause in executing an action (Frank, 2006).

This earlier model has since been adapted to explain how these basal ganglia systems can be driven by cortical structures to instantiate effective response inhibition (Wiecki and Frank, 2013). In go trials, saccades are generated by oculomotor structures such as the frontal eye field and superior colliculus. On stop-trials, these responses are inhibited after the right inferior frontal gyrus (rIFG) drives the subthalamic nucleus through the hyperdirect pathway, resulting in the excitation of the substantia nigra pars reticulata and subsequent inhibition of the superior colliculus. This model can reflect observed behavior and neural activity and can also correctly predict the effect of dopaminergic manipulations. However, this model is difficult to evaluate in comparison to more parsimonious models rooted in mathematical psychology, given the multiple, complex interactions between different nodes of the basal ganglia stopping network.

These observations have since developed into a model of stopping that more explicitly proposes response inhibition as a two-stage process: pause-then-cancel. This model was first developed in rodent data (Schmidt et al., 2013) but has recently been applied to human data (Diesburg and Wessel, 2021). Here, I focus on the human model, although the core concepts are similar. First, when unfamiliar, unexpected, or surprising events are detected in the environment, the rIFG drives activity in the STN. As described in the earlier Wiecki & Frank model, this increased STN activity results in greater inhibition over the basal ganglia, which requires greater input to generate a movement. This is reflected as a "pause" in which the threshold for deciding increases. This stage is not specific to stopping and is proposed as a response to any unexpected or surprising event in the environment. This is supported by evidence that rIFG is active in contexts that don't require stopping (including non-canceled trials) and when expectations are violated (Lenartowicz et al., 2011; Hampshire et al., 2010). This also unites the idea that rIFG is both an attentional circuit-breaker (Corbetta and Shulman, 2002) and an inhibitory node (Aron, 2007; Aron et al., 2014). It is proposed that this signal is communicated between rIFG and STN through beta oscillations (Fries, 2005); an increase in beta power is observed in both the rIFG and STN during stopping (Schaum et al., 2021; Swann et al., 2009, 2012; Wessel et al., 2016). However, direct evidence for information transmission between these areas is lacking (Chen et al., 2020; Mosher et al., 2021). Whilst this pause process increases the decision threshold for a movement to be executed, it's proposed a second process determines whether a movement is to be canceled. In such situations, a cancel signal is sent for situations requiring complete inhibition of an action. This is thought to be implemented by structures underlying the frontocentral P3. The frontocentral P3 is an event-related potential, classically considered as an index of motor inhibition (Kok et al., 2004; Wessel and Aron, 2015; Ramautar et al., 2004) but may also reflect unexpected events (Waller et al., 2021). Whilst unclear, the pre-SMA is one potential generator of the frontocentral P3 (Enriquez-Geppert et al., 2010; Swann et al., 2012), and thus is posed to provide the cancel signal. This signal removes the drive to move by inhibiting neurons in the striatum. If this cancel signal arrives in time, the active motor plan accumulating within the system will be halted. If it doesn't, then the plan can continue to accumulate.

1.7 Monitoring action

Although these neurocognitive models have provided mechanistic descriptions accounting for the reactive control of behavior and map onto distinct neural processes, they cannot explain proactive adjustments in behavior. As described previously, flexible goal-directed behavior requires a higher-order cognitive control system to organize and optimize information processing for the relevant context. These flexible behavior adjustments require continuous assessment and evaluation of our actions and environment. By monitoring and comparing our actions and their outcomes, we can evaluate the adequacy and success of performance. If

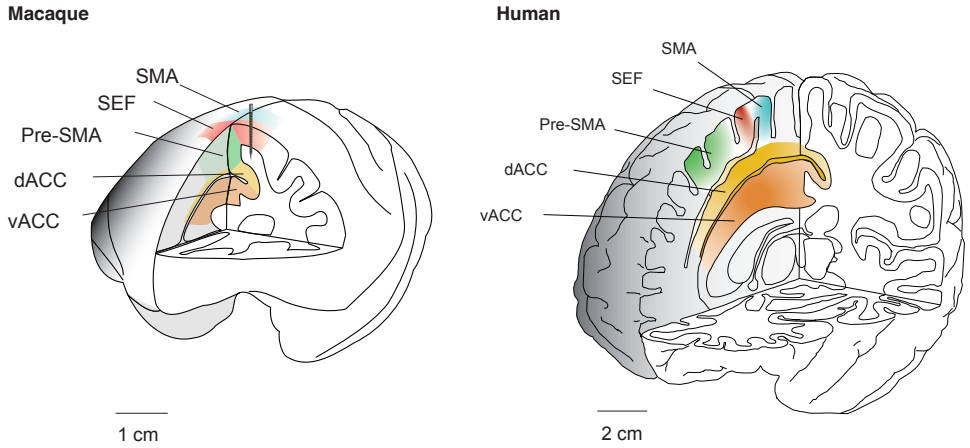

Figure 1.4: **Anatomy of the medial frontal cortex**. A cross-section of the macaque (left) and human (right) brain highlighting the relative locations of pre-SMA (green), SMA (blue), SEF (red), dorsal MCC (dMCC, yellow) and ventral MCC (vMCC, orange). Human brains are illustrated with the paracingulate sulcus present.

our performance is not meeting expectations or the desired outcomes, then adjustments to our behavior must be made.

1.7.1 Monitoring action in the Medial Frontal Cortex

Converging evidence from neuroimaging, electrophysiological recordings, and lesion studies in humans and non-human primates indicates that the medial frontal cortex is essential in monitoring our environment and eliciting cognitive control where necessary (Ridderinkhof et al., 2004; Rushworth et al., 2004; Klein-Flügge et al., 2022; Monosov et al., 2020). In this section of the introduction, I will first discuss the anatomy of the medial frontal cortex; here, I will describe how its structure and connectivity make it an ideal candidate for monitoring our actions within the world. I will then describe formalized theories of cognitive control, which highlight the role of the medial frontal cortex.

1.7.2 Anatomy of Medial Frontal Cortex

The medial frontal cortex is grossly considered as the part of the cortex lying dorsal and rostral to the corpus callosum. In the context of cognitive operations, two regions of the medial frontal cortex have been intensely studied: the supplementary motor complex and the anterior cingulate cortex. These areas in humans and macaques are depicted in Figure 1.4. Within this section, I will focus on findings from the macaque model. I will discuss the cross-species translatability of conclusions in the discussion of this dissertation.

1.7.2.1 Supplementary motor complex

The supplementary motor complex is a collection of cortical areas located within the superior frontal gyrus, anterior to the motor cortex, encompassing the supplementary motor area (SMA), pre-supplementary motor area (pre-SMA), and supplementary eye field (SEF). These areas are dorsal to the cingulate sulcus and gyrus and constitute the medial part of Brodmann's area 6aα and 6aβ (Matelli et al., 1991; Rizzolatti et al., 1998). As such, these areas are sometimes more broadly referred to as the dorsomedial frontal cortex in some literature (DMFC).

Functional properties of the areas comprising the Supplementary Motor Complex have been discussed in a previous section and in earlier reviews (Nachev et al., 2008); here I will instead focus on describing the differential contributions of pre-SMA, SEF, and SMA. Firstly, stimulation of these areas can evoke and inhibit movements – oculomotor movements in SEF (Schlag and Schlag-Rey, 1987; Stuphorn and Schall, 2006) and skeletomotor movements in SMA/pre-SMA (Zimnik et al., 2019). Second, SMA is somatotopically organized: movements of the lower extremities can be evoked by stimulation at caudal sites, and movements of the upper extremities (including orofacial) can be evoked by stimulation at more rostral sites (Mitz and Wise, 1987). Stimulating pre-SMA, on the other hand, is less well-defined and requires higher currents to elicit movements. Third, these areas differ in their pattern of projections – SMA has direct, reciprocal projections to motor systems, whilst pre-SMA and SEF instead demonstrate greater connectivity with regions in the frontal cortex (Luppino et al., 1993). Finally, whilst error and performance-related signals are commonly observed in pre-SMA and SEF, they are less widely reported in SMA. Collectively, these findings have been interpreted so that SMA can more directly contribute to the movement, whereas pre-SMA and SEF are more indirect, playing a more executive role.

As the research presented within this dissertation primarily focuses on the supplementary eye field, I will describe the important properties of this part of the supplementary motor complex in more detail.

1.7.2.1.1 Supplementary Eye Field

The supplementary eye field (SEF) is an oculomotor area in the dorsomedial frontal cortex, although the exact boundaries of the area have been well debated (Amiez and Petrides, 2009). Early descriptions of SEF deemed it a separate entity to the surrounding supplementary motor area, in which low currents (< 50 μA) could evoke saccades, and unit activity recorded through electrophysiological techniques demonstrated modulation to saccadic eye movements (Schlag and Schlag-Rey, 1987). In humans, SEF lies between the SMA and the pre-SMA, close to the paracentral sulcus (Grosbras et al., 1999). Although always located in the dorsomedial region of the frontal lobe, the exact location of the SEF has varied notably between studies; this may reflect that SEF is a functional area without a clearly defined anatomical landmark.

Cytoarchitecturally, the supplementary eye field is within area 6aβ of Vogt and Vogt (1919), area 6DR of Barbas and Pandya (1987), and corresponds to area F7 of Matelli et al. (1991). The area F7, as defined by Matelli et al., occupies the rostral part of the dorsal premotor cortex. This area is located 2 to 3 mm laterally from the midline to the fundus of the rostral part of the superior arcuate sulcus. Recent work has further divided F7 into three areas (dorsal, intermediate, and sulcal) based on the distribution of receptor density (Rapan et al., 2021); SEF appears to be restricted to the dorsal most aspect of F7, but this requires further validation.

SEF is considered an agranular area of the cortex – it lacks an inner granular layer (layer IV), has low neuron density, and has poorly defined layers. This is in stark contrast to areas such as the primary visual cortex, where a lot of previous work has contributed to a working hypothesis of a canonical cortical microcircuit (Ninomiya et al., 2015; Douglas et al., 1989; Beul and Hilgetag, 2014). Across the cortical layers of SEF, small and medium-sized excitatory pyramidal neurons with long apical dendrites are clustered in layer III. The cell bodies are slightly larger than those in adjacent areas, but the immunoreactivity is weaker. Small pyramids are also observed but loosely scattered in layer V. Inhibitory calretinin and calbindin-expressing neurons are most dense in layer II, and diminish towards deeper layers (Godlove et al., 2014). Neurons expressing parvalbumin, however, are more uniformly distributed across cortical layers (Godlove et al., 2014).

SEF is primarily innervated by sensory and motor-related brain regions. Two primary projections from the SEF are to the frontal eye field (FEF) and superior colliculus (SC); as previously discussed, both areas are major channels to the saccade generator in the brainstem. SEF is also reciprocally connected with the anterior and posterior cingulate cortex; injections into SEF result in considerable labeling in cingulate areas 24a, 24b, and 24c (Huerta and Kaas, 1990). Interestingly, these connections are not demonstrated by the FEF. Finally, SEF also shows strong reciprocal connections to the caudate nucleus and lateral prefrontal cortex (Huerta and Kaas, 1990; Griggs et al., 2017). Collectively, this pattern of connections is conducive to SEF as a region for monitoring action and exerting executive control over oculomotor behaviors.

1.7.2.2 Anterior Cingulate Cortex

A second area in medial frontal cortex frequently implicated in cognitive control is the cingulate cortex. Cingulate cortex is a collection of different regions supporting a range of functions, from emotion to motivation to cognition and motor control. Habitual descriptions of "anterior cingulate cortex" can be replaced by more precise nomenclature based on more refined anatomical and comparative studies, which distinguish the middle cingulate cortex (MCC) ventral to the supplementary motor cortex from the anterior cingulate cortex (ACC) around the genu of the corpus callosum (Vogt et al., 2003). MCC includes the cingulate motor areas with projections to the spinal cord whilst ACC, on the other hand, demonstrates greater connectivity to lim-

bic areas such as the amygdala and hypothalamus (Paus, 2001). Whereas the rostral ACC contributes more to emotion and autonomic regulation, the MCC contributes more to performance monitoring and executive control. Given this role, I will focus on the middle cingulate (or midcingulate) cortex in this dissertation.

In macaques, MCC is comprised of several separate anatomically distinct areas, extending ventrally from the dorsal and ventral banks of the anterior cingulate sulcus, along the medial wall, to the corpus callosum. Anatomical descriptions of the MCC in macaques note that the cytoarchitecture of the dorsal and ventral banks of the cingulate sulcus differ. Whilst the ventral bank of MCC is identified as area 24, and homologous to MCC in humans, the dorsal bank is recognized as an extension of area F6 caudally or area 9 rostrally. Given this evidence, some authorities argue the dorsal bank should not be considered part of the cingulate cortex proper, but others disagree. Nevertheless, despite these differences in cytoarchitecture, neurophysiological investigations have mainly sampled spiking activity in the dorsal bank (Heilbronner and Hayden, 2015; Procyk et al., 2016; Vogt et al., 2005). Little is known about how the functional organization of areas across both banks of the cingulate sulcus contributes to performance monitoring. To date, only a few studies to date have compared the neural signals in the two banks of the midcingulate cortex in macaques (Cai and Padoa-Schioppa, 2012; Procyk et al., 2016; Shidara and Richmond, 2002; Monosov, 2017). Understanding the differences between the dMCC, vMCC, and SEF is undoubtedly as relevant to our understanding of how performance monitoring signals function as the differences between V1 and V2 or MT and V4 are to our understanding of visual processing.

1.7.3 Theories of Medial Frontal Cortex function

A wealth of research across species and methodologies has highlighted that areas within the medial frontal cortex contribute to the monitoring of action, outcomes, and value (Ridderinkhof et al., 2004; Rushworth et al., 2004; Klein-Flügge et al., 2022; Monosov et al., 2020). With this, the medial frontal cortex has played a central role in many models of cognitive control. In this section, I will describe several key theories and hypotheses of cognitive control that are instantiated by areas in the medial frontal cortex.

1.7.3.1 Early studies of the functionality of the medial frontal cortex

The anterior cingulate cortex has been implicated as a critical area for cognitive control for almost a century. As early as Papez's theory of emotion (1937), it has been proposed that the ACC processes and broadcasts a context signal (in this case, emotional context) which biases processing in other brain areas. Although functional theories of ACC moved away from emotional to more cognitive-based ideas, the premise remained the same. The proposal of a broadcasted signal that influences other areas of the brain was also considered in later network theories of cognitive control (Mesulam, 1981; Posner and Dehaene, 1994). This was developed

further by attentional researchers, who stated that ACC was critical for the selection of goal-relevant actions and implementing the relevant control (Posner et al., 1988; Posner and Dehaene, 1994). In conjunction with these proposals, lesions of ACC produced marked attentional deficits (Janer and Pardo, 1991).

1.7.3.2 Conflict Monitoring

Whilst this early work provided vital insight into the function of the medial frontal cortex, this knowledge was driven further in the 1990s with the development of more nuanced neuroimaging techniques. Using these techniques, several studies demonstrated that ACC showed strong activation across multiple different tasks that required cognitive control (Carter, Mintun, & Cohen, 1995; Pardo et al., 1990, Corbetta et al., 1991). These observations led to one of the most popular early theories of cognitive control in the medial frontal cortex - the conflict monitoring theory (Botvinick et al., 2001).

The conflict monitoring theory proposes that cognitive control is essential when there are multiple competing and incompatible responses - referred to as response conflict. This model defines response conflict as the simultaneous activation of neuronal assemblies associated with incompatible behavioral responses. Computationally, the Hopfield energy formula (Hopfield, 1982) provides a simple formulation of this signal. Briefly, this formula proposes conflict can be quantified as the activity of one unit multiplied by the activity of the competing units, weighted by the connection between the two processes. This measure of conflict is not the only possible formulation (Yu et al., 2009), but it formalizes conflict intuitively and straightforwardly. After detecting this conflict, the model proposes that this information is used to signal downstream areas, adjusting the level of attentional control. For example, when high levels of conflict are experienced, control can be exerted to enhance the processing of task-relevant stimuli and response representations in subsequent trials. Increasing control after high-conflict trials can reduce conflict in subsequent trials (Gratton et al., 1992).

A range of neuroimaging evidence has shown increases in ACC BOLD activity during conflict (van Veen and Carter, 2002). In one such study, Kerns et al. (2004) demonstrated that ACC activity was greater during incongruent trials in the Stroop task; this activity was associated with an increase in dlPFC activity in the subsequent trial indexed slower response times. Despite this neuroimaging evidence in humans, surprisingly few single-unit studies support the hypothesis that ACC signals conflict (Ito et al., 2003; Nakamura et al., 2005). The reasons behind this are unclear and warrant further investigation.

1.7.3.3 Value-based decision-making & foraging

Although popular, the conflict monitoring model failed to account for other observations made in ACC. Specifically, a series of studies demonstrated that ACC exhibits pronounced modulation in response to out-

comes – such as errors and rewards (Gehring and Willoughby, 2002). This led to alternative proposals that ACC may be used to guide action selection based on previous outcomes (Rushworth et al., 2004). This was supported by other work in macaques that demonstrated that ACC lesions did not disrupt task-switching – a process that requires effective conflict detection - but did affect the evaluation of effort needed to get a reward (Rushworth et al., 2003). Furthermore, populations of ACC neurons also demonstrated activity that reflected action-outcome mappings required for sequence learning (Procyk et al., 2000).

1.7.3.4 Error-likelihood and PRO model

Attempting to integrate some of these differential findings, the error-likelihood model was developed (Brown and Braver, 2005). Specifically, this model aimed to reconcile the observed co-activation of the medial frontal cortex during errors and outcomes and during conflict. Consideration of this co-activity presents an important confound between these two concepts: when response conflict is high, errors are more likely, and reward is less likely. As such, it is unclear as to which cognitive process to attribute these signals to in ACC. Under the error-likelihood model, Brown & Braver framed the medial frontal cortex as a predictor, reflecting how likely an error is to occur. In situations in which there is an error executed, or there is high conflict, this model predicts that the medial frontal cortex exerts a higher level of cognitive control compared to cases in which conflict, and thus the likelihood of an error, is low.

While the error likelihood model accounts for several effects observed within ACC, initial replications of the error likelihood effect yielded null results (Nieuwenhuis et al., 2007). Furthermore, additional evidence highlighted features of ACC function that were not accounted for in the model. Most notably, the error-likelihood effect seemed to be influenced by surprise: unexpected errors elicited greater activity than less surprising errors (Holroyd and Krigolson, 2007). Resultantly, this model was later revised into the prediction response-outcome model (Alexander and Brown, 2011). Again, focusing on the medial frontal cortex, this model proposes that this area may have a broader role than simply detecting error-likelihood or conflict. Instead, it proposes the medial frontal cortex is central to both forming expectations about actions and detecting surprising outcomes and formalizes these features within a reinforcement learning framework. Through this framework, an agent (such as a monkey, human, or AI) can learn the anticipated values of multiple potential actions through response-outcome mappings.

1.7.3.5 Expected value of control

In addition to the development of the error-likelihood and PRO models, evidence for the foraging and value-based decision-making hypotheses drove another alternative and extensive revision of the conflict-monitoring hypothesis – the Expected Value of Control (EVC) hypothesis (Shenhav et al., 2014, 2016b). This hypothesis

incorporates the ideas conveyed in the earlier conflict-monitoring theory but offers a broader explanation for ACC function. This hypothesis predicts that control signals vary across two different dimensions: identity and judgment. Across the identity dimension, an agent must make a judgment about what the appropriate action and response may be, given the context. Across the judgment dimension, the agent determines how the appropriate response should be executed. Interestingly, these dimensions seem to parallel the two components of decision-making proposed by (Schall, 2001).

In the Expected Value of Control hypothesis, conflict is considered more broadly as the need to put forth effort; under high-conflict scenarios, more effort is required to suppress incorrect responses and execute the correct response. These signals then drive systems biasing decisions toward an appropriate response. Attempting to reconcile this theory with the foraging literature, the EVC considers that choice difficulty within foraging tasks can also be representative of effort – more difficult choices (i.e., whether to leave a patch or stay) require greater effort to resolve. This is supported by human neuroimaging studies that demonstrated that ACC activity peaked when decision difficulty was highest (Shenhav et al., 2014, 2016b).

By offering an alternative description of conflict as effort, EVC provides a more comprehensive explanation than that proposed in the original conflict-monitoring hypothesis. This extension also provides a way to account for divergent evidence of ACC function under one hypothesis.

1.8 Dissertation outline and summary

In the previous sections of this chapter, I have described how cognitive control is essential in ensuring we make the most appropriate action or decision given the relative context. Such control is vital in situations where there are multiple competing demands or when we have automatic or learned actions that may interfere with our goal. Here, these inappropriate actions must be stopped before their execution; this is achieved through a facet of cognitive control referred to as response inhibition. I explained how the stop-signal task is a paradigm in which response inhibition can be empirically tested. Performance in this task is rooted in a well-supported cognitive framework (independent race model) that allows for the latency of response inhibition to be estimated. This stop-signal reaction time has been used to determine whether neural activity can directly stop an action from occurring. After defining the neural systems that contribute directly to the generation or inhibition of an action, I then demonstrated how these findings have been synthesized into broader models of response inhibition that link neural activity to behavior.

Although these models can describe how control can be exerted in the moment, they typically fail to account for other observations, such as changes in behavior dependent on outcomes of previous trials and knowledge about the task. Indeed, flexible adjustments of behavior require the continuous assessment and evaluation of our actions and environment. Converging evidence from neuroimaging, electrophysiological

recordings, and lesion studies in humans and non-human primates has demonstrated that the medial frontal cortex is a key area contributing to performance monitoring. Given this role, I then reviewed key regions of medial frontal cortex, focusing on the midcingulate cortex and the supplementary eye field. I then described how these areas are considered to instantiate theories of cognitive control.

However, activity within the medial frontal cortex is complex, and signals vary markedly with no apparent pattern or adherence to a single cognitive theory or computational model. Furthermore, it is unclear how different areas within the medial frontal cortex interact with one another to promote cognitive control. Through this dissertation, I will describe five studies that investigate the mechanisms through which the medial frontal cortex contributes to cognitive control.

In chapter 2, I describe how executive control processes can influence seemingly automatic behaviors. Comparing observed and simulated data, I demonstrate that monkeys appear to modulate the proportion of short-latency saccades in conjunction with variations of higher-order features of the task environment, such as reward rate and temporal regularities.

In chapters 3 and 4, I build on recent literature that describes an increase in the incidence of beta-bursts over the medial frontal cortex during successful stopping in humans – a finding that doesn't correspond with evidence from single-unit recordings in areas of the supplementary motor complex and anterior cingulate cortex. In chapter 3, I first demonstrate that although the pattern of beta-bursts in macaque EEG mirror observations in humans, they occur too infrequently and don't represent stopping the behavior. In chapter 4, I develop this observation further to understand the underlying mechanism by describing the patterns of beta-burst activity within the cortex, their distribution across cortical layers, and their association with beta-bursts observed in a simultaneously recorded EEG signal on top of the skull. Through this work, I provide evidence that oscillatory activity in the medial frontal cortex lacks the causal efficacy to directly control behavior and stop eye movements reactively.

In chapters 5 and 6, I then describe how neural activity at the single-neuron level may reflect aspects of cognitive control. In chapter 5, I describe three groups of neurons in the Supplementary Eye Field that are modulated when a planned saccade is successfully inhibited: neurons signaling the level of conflict experienced, neurons reflecting the time of important task events, and neurons that maintain a representation of the task-goal. I also describe these neurons' biophysical and laminar properties and explain how they relate to two event-related potentials (N2 and P3) recorded in the EEG signal. In chapter 6, I extend this work to investigate and describe signals in the dorsal and ventral banks of the midcingulate cortex that contribute to cognitive control at the single-unit level.

In the final chapter of this dissertation, I will summarize and evaluate these findings, discussing their implication for theories of cognitive control and the medial frontal cortex.

CHAPTER 2

Executive Control of Response Latency in a Saccade Countermanding Task

2.1 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 learnt over time (Bibi and Edelman, 2009; Fischer et al., 1984; Jóhannesson et al., 2018; McPeek and Schiller, 1994; Pare 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, non-aging 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 ($\sim$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 $\sim$80 ms in macaques and $\sim$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; Kornylo 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, 2009b). 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 saccades are most common in gap task condition that encourages 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 contingences, 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.

2.2 Methodology

2.2.1 Animal Care

Data were collected from 3 male bonnet macaques (Macaca radiata, 6.9 to 8.8 kg), and 1 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).

2.2.2 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 ($\sim$48 x 38°, 70Hz) 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 x 400 pixel resolution, Reflective Computing, Olympia, WA). Stimuli had a luminance of 3.56 cd/m^2 (fixation point and stop-signals) or 2 cd/m^2 (targets) on a 0.01 cd/m^2 background.

2.2.3 Behavioral Task

Behavior and electrophysiological signals were recorded during the countermanding (i.e., stop-signal) task (Figure 2.1). 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 dataset were compared against data recorded several years later from monkey's Eu and X. Data from both of these monkeys were collected after both were well trained on the task, and neither monkey was rewarded for an express saccade at any point in their training.

Trials were initiated when monkeys fixated a centrally presented square which 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 & H) foreperiods were randomly sampled from a uniform distribution. For two other monkeys (Eu & X) foreperiods were randomly sampled from an approximately non-aging, exponentially decaying distribution. On no-stop trials (Figure 2.1, 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 re-illuminated after a variable delay providing a "stop-signal" which instructed the monkeys to cancel their

impending eye movements and maintain central fixation (Figure 2.1, 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 & 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 "non-canceled". If a saccade was initiated before the stop signal was scheduled to appear, the trial was classified as non-canceled based on the logic of the Logan race model (Logan and Cowan, 1984). No time outs were imposed for a non-canceled 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 SSD's typically ranged from 43 up to 443 ms, in steps of 16 ms. SSD was then manipulated using either an adaptive staircasing algorithm which adjusted stopping difficulty based on accuracy, or by randomly selecting one of the defined SSD's. 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 $\sim$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 minutes (IQR: 27 to 77 minutes), and for monkey H was 236 minutes (119 to 366 minutes). In later data recorded from monkey Eu and X, trial length was fixed. For these monkeys, session lengths ranged from 87 to 184 minutes (modal session time: 169 minutes) for monkey Eu, and from 62 to 193 minutes (modal session time: 120 minutes) for monkey X. For monkey F, the inter-trial interval was fixed ($\sim$1000 ms), resulting in a variable trial length. However, for monkey H, the inter-trial interval was varied. In all cases, trial length would be extended if a time-out ($\sim$500 ms for monkey F & H, $\sim$5000 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, timeouts were not issued for non-canceled trials; instead, the temporal progression of the trial would instead mirror that of a no-stop trial, but with reward omitted.

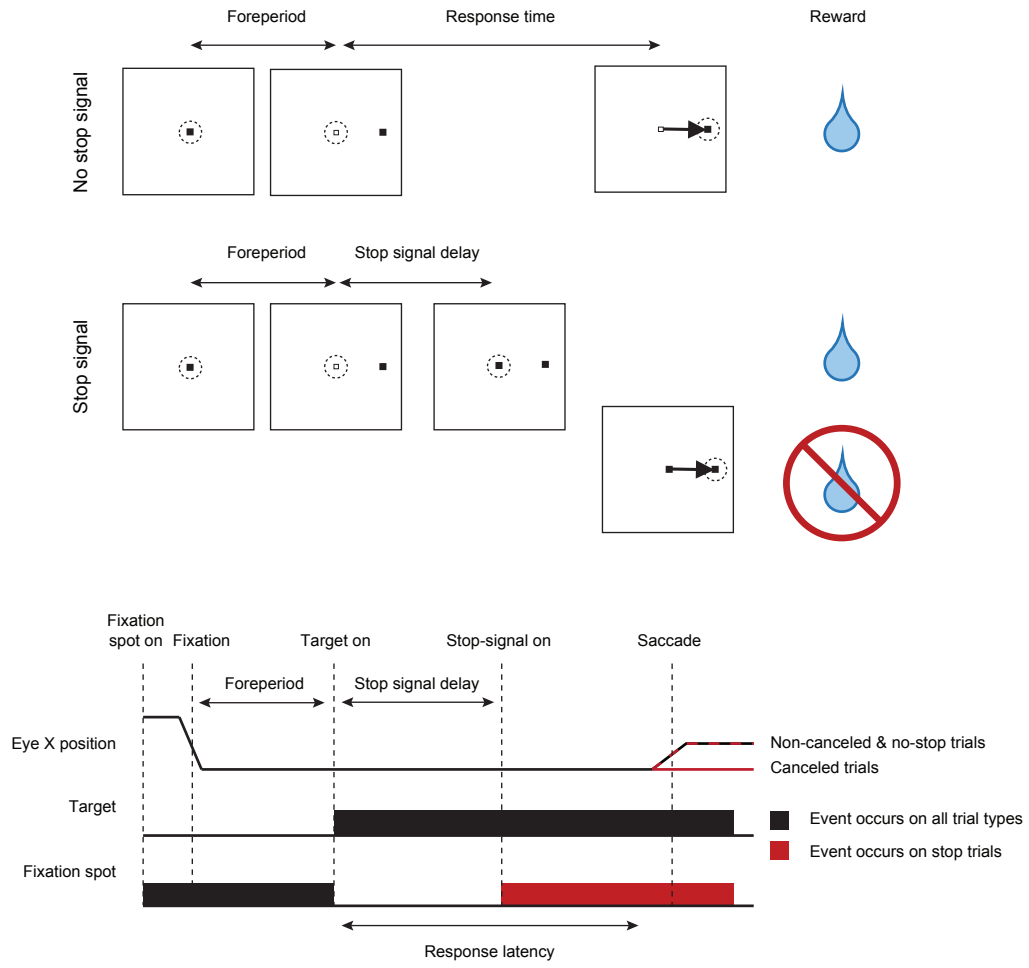

Figure 2.1: **Saccade countermanding task**. Monkeys initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target was presented simultaneously at one of two possible locations. On no-stop-signal trials, monkeys were required to shift gaze to the target, whereupon after a variable period of time, fluid reward was provided. On stop-signal trials ($\sim$40% of trials), after the target appeared, the center of the fixation point was re-illuminated 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 $\sim$50% of trials. In the remaining trials, monkeys made non-canceled errors, in which no reward was delivered. In non-staircased trials, stop-signal delays were randomly selected from a pre-determined set of stop-signal delays.

	Foreperiod	Ageing		Non-ageing	
Saccade type	**Trial type**	**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
	Non-canceled	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
	Non-canceled	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
	Non-canceled	71.0 ± 5.5	84.1 ± 4.1	71.2 ± 8.4	85.3 ± 13.1

Table 2.1: Average ± SD saccade latencies across trial types and data sets

2.2.4 Eye Tracking

Eye position data were acquired, calibrated, and streamed to the Plexon computer using the EyeLink 1000 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 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.

2.3 Results

2.3.1 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 multi-modal latency distributions. Evident in both single behavioral sessions and across sessions, saccades with latencies $\leq$ 100ms 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 datasets, split by trial type, are displayed in Table 2.1. 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 non-canceled). The response distribution was clearly multimodal (Figure 2.2A,

		Ageing Monkey F (n = 84)		Monkey H (n = 157)		Non-ageing Monkey Eu (n = 12)		Monkey X (n = 17)	
RT relationship	**Saccade type**	**RT (ms)**	**P(Express)**	**RT (ms)**	**P(Express)**	**RT (ms)**	**P(Express)**	**RT (ms)**	**P(Express)**
		n = 50 sessions		n = 83 sessions		n = 3 sessions		n = 0 sessions	
$RT_{Non\text{-}canceled} > RT_{No\text{-}stop}$	**No-stop**	225.8 ± 49.8	0.139 ± 0.119	261.4 ± 41.5	0.151 ± 0.084	303.7 ± 1.6	-	-	-
	Non-canceled	249.8 ± 50.3	0.042 ± 0.075	301.2 ± 51.9	0.110 ± 0.114	316.6 ± 17.7	0.003 ± 0.005	-	-
		n = 34 sessions		n = 74 sessions		n = 9 sessions		n = 17 sessions	
$RT_{Non\text{-}canceled} < RT_{No\text{-}stop}$	**No-stop**	237.8 ± 52.2	0.105 ± 0.138	267.0 ± 59.3	0.172 ± 0.1	314.7 ± 15.5	0.004 ± 0.004	265.8 ± 18.9	0.003 -0.003
	Non-canceled	221.0 ± 52.4	0.112 ± 0.177	240.3 ± 66.7	0.195 ± 0.21	270.9 ± 15.6	0.009 ± 0.009	232.4 ± 12.9	0.006 -0.006

Table 2.2: Average ± SD of saccade latencies and express saccades proportions across trial types and monkeys, split by race model violations.

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 ms – 92 ms). Monkey H generated saccades in 90,347 trials (73,750 no-stop, 16,597 non-canceled). Again, the response time distribution was multimodal (Figure 2.2A, 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 (Figure 2.2B).

2.3.2 Express saccade production & 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 non-canceled responses ($RT_{non\text{-}canceled}$) is consistently faster than the finishing time of responses when there is no stop signal ($RT_{no\text{-}stop}$). Saccade latencies and express saccade proportions, split by trial types and monkeys, are presented in Table 2.2, for sessions with and without violations of the race model. In these archived data, we noted a number of sessions in which the 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 sat-

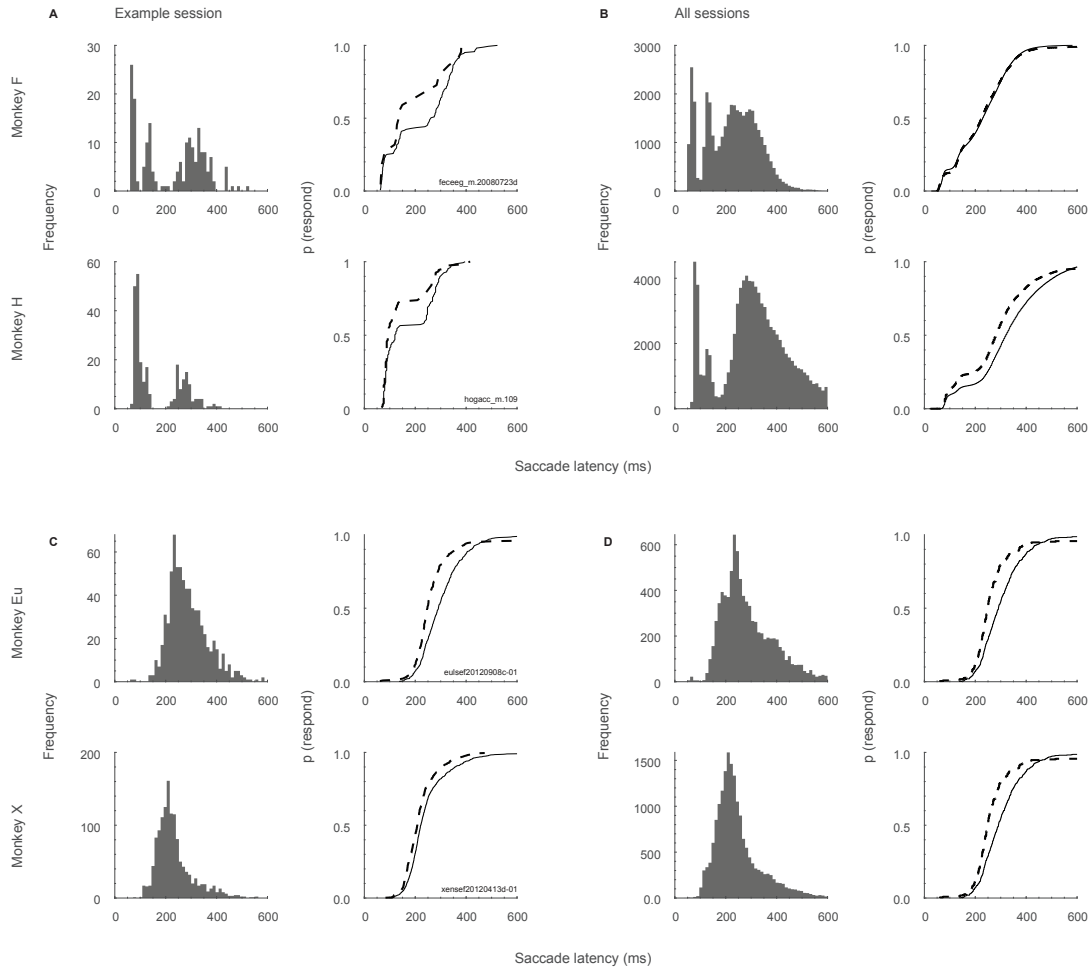

Figure 2.2: **Distributions of saccade latencies during saccade countermanding**. Histograms (left) and cumulative distribution functions (right) of saccade latencies in non-canceled (dashed lines) and no-stop trials (solid lines). **A**. Saccade latencies in example sessions for monkey's F & H. **B**. Saccade latencies collapsed across all sessions for monkey's F & H. **C**. Saccade latencies in example sessions for monkey's Eu & X. **D**. Saccade latencies collapsed across all sessions for monkey's Eu & X.

isfying the stochastic independence assumption of the race model. Monkey F violated the $RT_{non\text{-}canceled} < RT_{no\text{-}stop}$ 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 non-canceled 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 $RT_{non\text{-}canceled} < RT_{no\text{-}stop}$ 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 non-canceled 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 $RT_{non\text{-}canceled} > RT_{no\text{-}stop}$, but the proportion of express saccades did not differ across sessions or trial types. Monkey X had no sessions in which $RT_{non\text{-}canceled} > RT_{no\text{-}stop}$.

2.3.3 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 (Figure 2.3A). 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 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%; Figure 2.3B). 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.

2.3.3.1 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 in order 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) (Figure 2.3C, squares). However, when trial lengths are fixed, trial rates only slightly decrease and subjects can only initiate around 11 trials per minute (Figure 2.3C, 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 if these values varied as a proportion of stop trials in a given session (Figure 2.3E, upper panels). Interestingly, if a session had a lower proportion of stop-trials ($\sim$10%), then producing more express saccades led to only a small increase in the number of trials available. However, if a session had a greater proportion of stop-trials ($\sim$70%), then producing more express saccades led to a greater increase in the number of trials available compared to 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.

2.3.3.2 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 (Figure 2.3D). At typical stop-trial proportions ($\sim$40%), producing a small proportion of express saccades was not detrimental to performance, and led to slightly increased reward rates (Figure 2.3D). This relationship changed dependent on the proportion of stop-trials in a given session (Figure 2.3E, lower panels). If a session

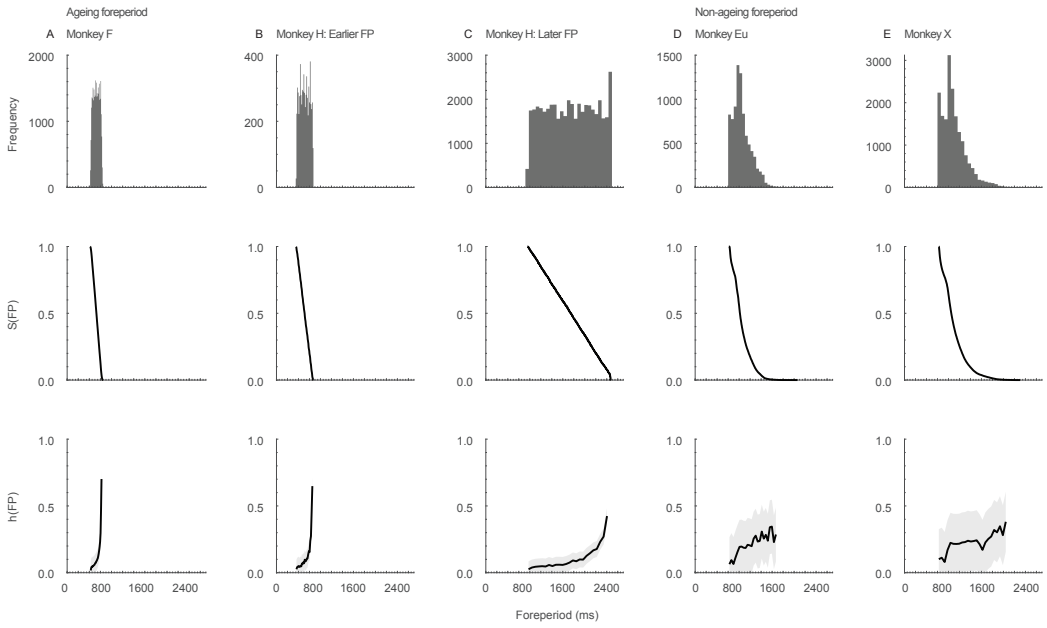

Figure 2.3: **Distributions of foreperiods**. Histograms (top), survivor distribution (middle), and hazard rate (bottom) of foreperiods used for monkey F (**A**), monkey H in sessions with earlier foreperiods (**B**), monkey H in sessions with longer foreperiods (**C**), monkey Eu (**D**), and monkey X (**E**). Monkey's F and H experienced aging foreperiods, and monkey's Eu and X experienced approximately non-aging foreperiods.

had a lower proportion of stop-trials (∼10%), then producing more express saccades lead to a higher reward rate (Figure 2.3E). However, if a session had a greater proportion of stop-trials (∼60%), then producing more express saccades was detrimental to performance (Figure 2.3E). 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.

2.3.4 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 10,000's of trials.

2.3.4.1 Express saccade production is learnt through training

For monkey F, we tracked the progression of express saccade production from the initial training session onwards. 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 ($F(4,171) = 7.76$, $p < 0.001$, Figure 2.4A), as did reward rate ($F(4,171) = 11.69$, $p < 0.001$, Figure 2.4B).

2.3.4.2 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 (Figure 2.4C). 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, χ^2 (1, N = 10465) = 405.648, $p < 0.001$ (Figure 2.4D). 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, $p < 0.001$, Figure 2.4E). Interestingly, average reward rate exhibited no trends (r = 0.183, p = 0.467, Figure 2.4F).

2.3.4.3 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 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 (Figure 2.2C & 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.

2.3.5 Temporal predictability can affect express saccade production

Express saccades occur more often when the timing of a target presentation is predictable (Pare 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 just 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 ageing function. This ageing 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 dataset. 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

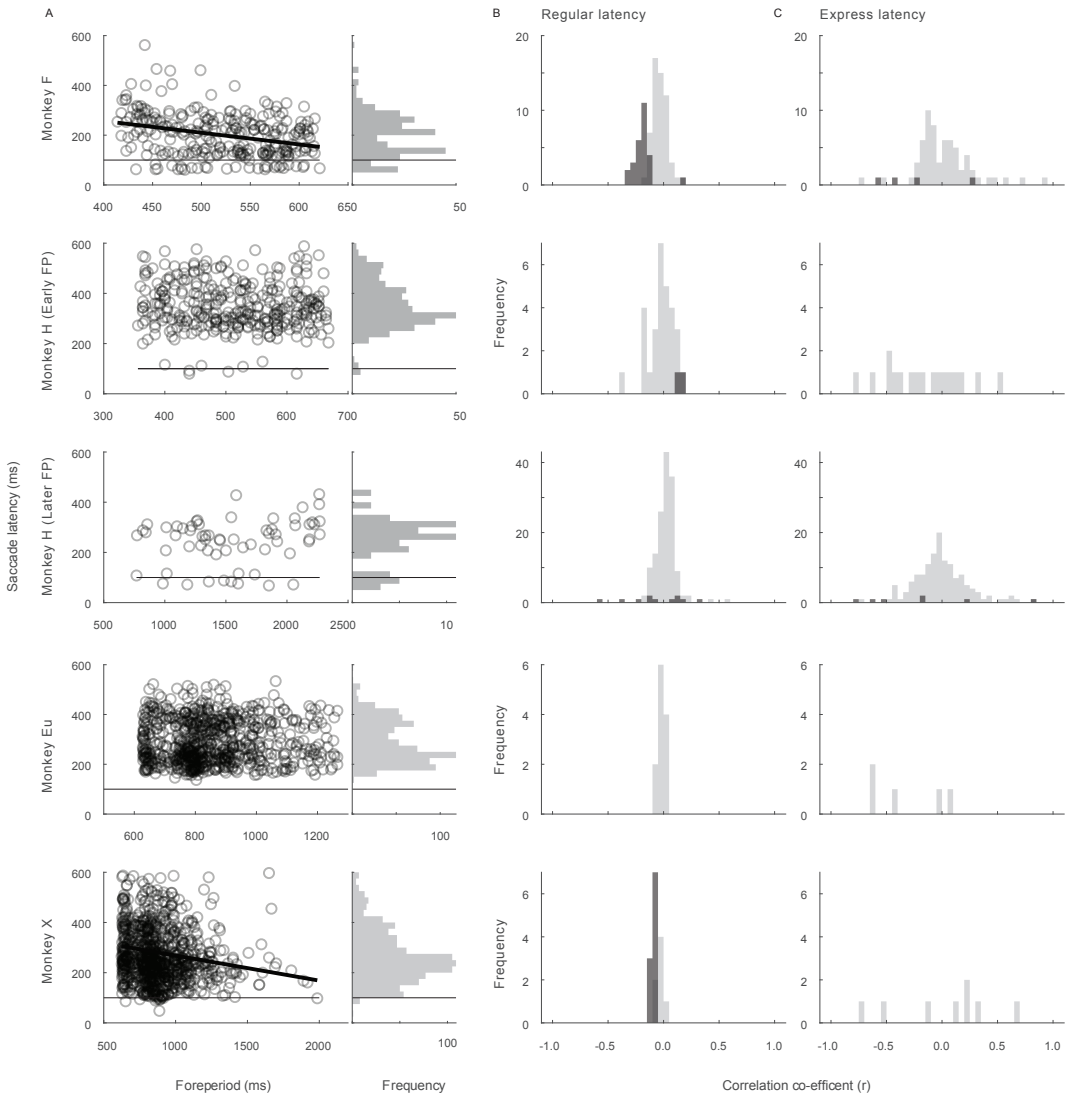

Figure 2.4: **Relation of saccade latency and foreperiod**. **A**. Scatter plots of saccade latency as a function of foreperiod with marginal distribution of saccade latency for representative sessions with monkey F (top), monkey H with early foreperiods (second), monkey H with later foreperiods (third), monkey Eu (fourth), and monkey X (bottom). Horizontal line indicates latency of 100 ms. The express saccades produced by monkeys F and H are evident. Significant correlations of saccade latency with foreperiod are indicated by black lines. **B**. Distributions of correlation coefficients for saccades with regular (>100 ms) latency with foreperiod for entire sample with each monkey. Monkey F and X produced some sessions with significant correlations ($p < 0.05$) (indicated by dark bars). **C**. Distributions of correlation coefficients for saccades with express latency with foreperiod for entire sample with each monkey. Express saccade latency was not systematically correlated with foreperiod.

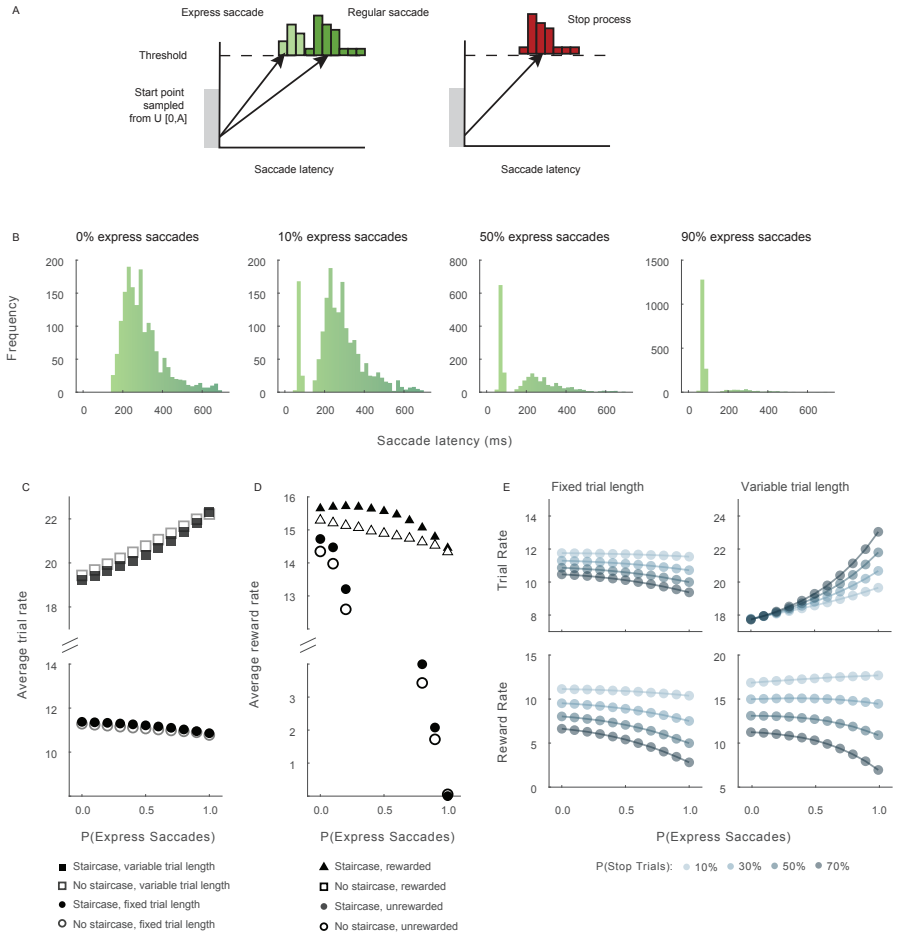

Figure 2.5: **Simulation of reward rate earned with variable fractions of express saccades**. **A**. Linear ballistic accumulators for GO process (left) and STOP process (right). The GO process consisted of two mean rates, one producing express saccade latencies, and the other producing regular saccade 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 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 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 (closed) and without (open) 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.

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 ∼630ms (Figure 2.5A). 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 ∼2275 ms (Figure 2.5B). 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 ($F(15, 2438) = 14.44$, $p < 10\text{-}5$).

Monkey Eu and monkey X experienced a common pattern of foreperiod distribution ranging from ∼600 to ∼ 2100 ms (n = 29/29 sessions). Compared to the foreperiod experienced by monkeys F & H, the survivor function of this distribution was exponentially decaying, resulting in a non-ageing foreperiod (Figure 2.5C & D). As expected with non-ageing 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, 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, $F(15, 206) = 2.26$, $p = 0.006$.

2.4 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 learnt 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.

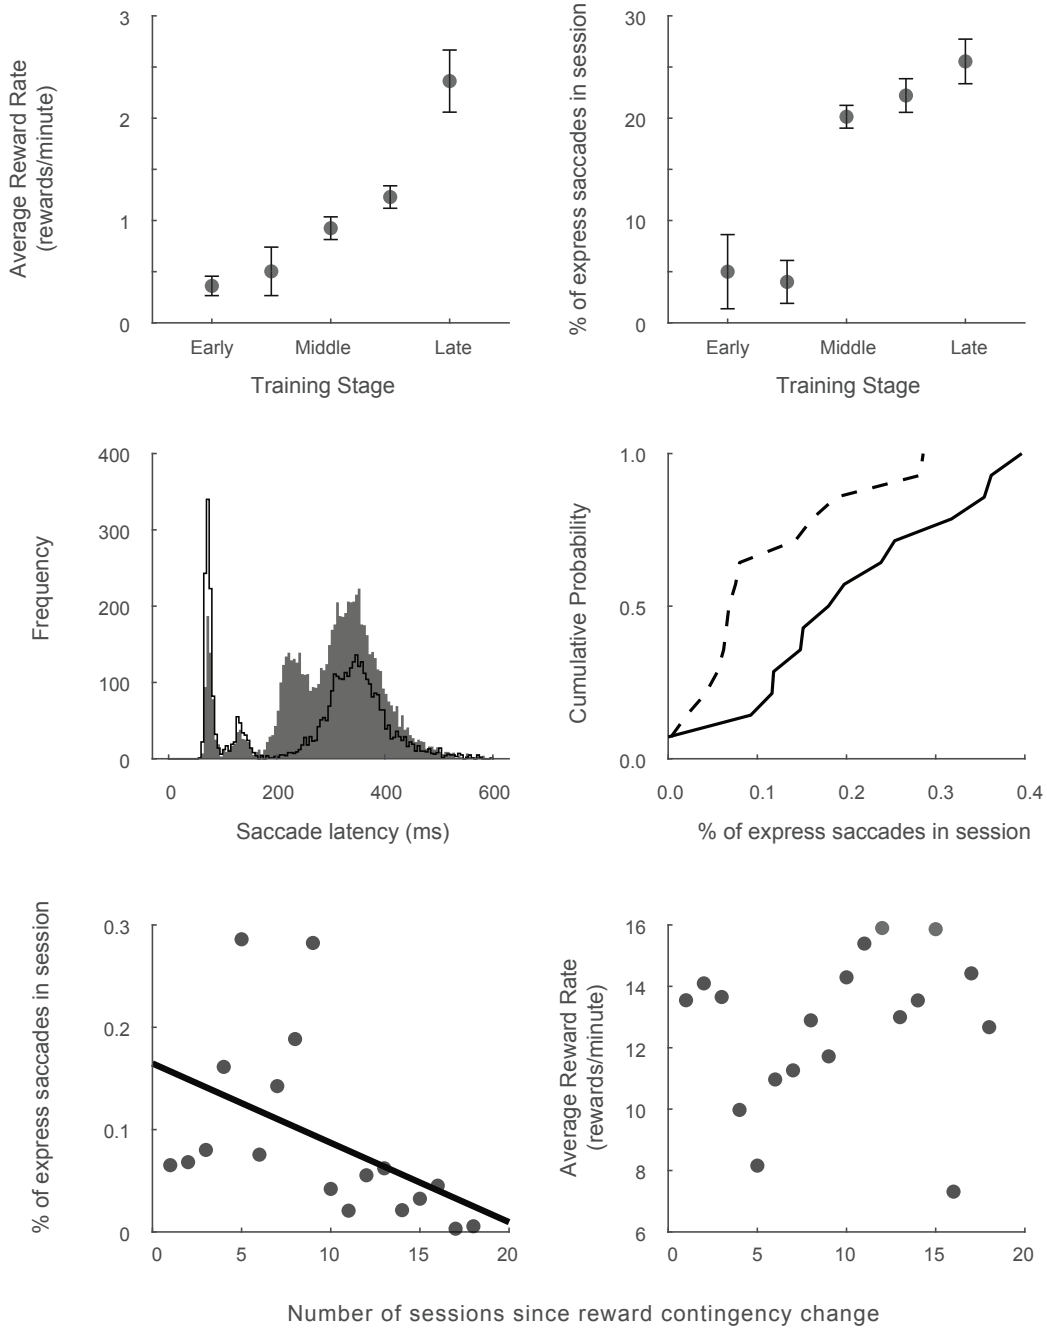

Figure 2.6: **Learning and unlearning express saccades**. Variation of performance of monkey F as a function of training and task contingencies. **A**. Average reward rate earned from early to late stages of training in 18 sessions when express saccades were rewarded just like regular saccades on no 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) and in subsequent sessions when only no stop trials with regular saccade latencies were rewarded (filled). When express saccades were not rewarded, the monkey produced more regular latency saccades to earn reward. **D**. Cumulative probability of express saccade production in sessions when they were rewarded (solid) and in later sessions when they were not (dashed). **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.

2.4.1 Express saccades production in atypical conditions

Consistent with classic reports of express saccades in monkeys, the earliest modes in the resulting distributions peak below 100ms 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 whilst 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 monkeys have some knowledge about the underlying contingencies in our experiment, and generated make predictions to adapt saccade latencies accordingly.

2.4.2 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.

2.4.3 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 in order 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.

Non-human primates are adept at optimizing reward rate by exploring different behavioral strategies and task parameters (Feng et al., 2009). Feng et al. (2009) 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-/anti-saccade visual search task where a vertical singleton cued that a pro-saccade should be made towards it, and a horizontal singleton cued an anti-saccade away from it. Close examination of the reward contingencies of the experiment found that shifting gaze towards 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.

2.4.4 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.

CHAPTER 3

Dissociation of Medial Frontal β-Bursts and Executive Control

3.1 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., 2019b). Recently, researchers have described a higher incidence of β-bursts in 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 (Jong et al., 1995; Kok et al., 2004; Stahl and Gibbons, 2007; Godlove et al., 2011; Reinhart et al., 2012b; Swann et al., 2012; Wessel and Aron, 2015) and neurophysiological (Hanes et al., 1998; Pare and Hanes, 2003; Brown and Heathcote, 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 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 (Pare 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 in spite of the stop signal with a neurometric function of the probability of neural activity exceeding a threshold (Brown and Heathcote, 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., 2011) and humans (Reinhart et al., 2012b) and identified a medial frontal source for the ERN (Sajad et al., 2019). Here, we seek to establish the homology of beta 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 (RT). We conclude that these beta 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.

3.2 Methodology

3.2.1 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 was collected from one male bonnet macaque (Monkey Eu, Macaca Radiata, 8.8 kg) and one female rhesus macaque (Monkey 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-hour 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.

3.2.2 Surgical Procedures

Surgical details have been described previously (Godlove et al., 2011). 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 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 (Crist Instruments). Chambers were implanted normal to the cortex (Monkey Eu: 17^o; Monkey X: 9^o; relative to stereotaxic vertical) centered on midline: 30 mm (Monkey Eu) and 28 mm (Monkey X) anterior to the interaural line.

3.2.3 Data Acquisition

An identical daily recording protocol across monkeys and sessions was carried out. 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/m^2). The monitor had a refresh rate of 70 Hz, and the screen subtended 46^o x 36^o of the visual angle. Eye position data was collected at 1 kHz using an EyeLink 1000 infrared eye-tracking system (SR Research, Kanata, Ontario, Canada). All data were streamed to a single data acquisition system (MAP, Plexon, Dallas, TX). Time stamps of trial events were recorded at 500 Hz.

3.2.4 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 $\sim$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 bandpass filtered between 0.7 and 170 Hz. All data were streamed to a data acquisition system (MAP, Plexon, Dallas, TX).

3.2.5 Saccade Stop-Signal (Countermanding) Task

The saccade stop-signal task utilized in this study has been widely used previously (Hanes and Schall, 1995; Hanes and Carpenter, 1999; Cabel et al., 2000; Colonius et al., 2001; Kornylo 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., 2019b). 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., 2019b). 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 to 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). Secondly, although we employed a staircase procedure, this stepped by one to three stop-signal delays. 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 to 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 (stop-signal delay; SSDs). An initial set of SSDs, separated by either 40 or 60 ms, were 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 approximately 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 600ms later, and followed by a high or low fluid reward, depending on the stimulus-response mapping.

3.2.6 Experimental Design & Statistical Analysis

3.2.6.1 Bayesian modelling 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. Stop-signal reaction time was estimated using a Bayesian parametric approach (Matzke et al., 2013a, 2017b). Compared to 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 response time (RT) distribution. Furthermore, this model also allows for the estimation of the probability of trigger failures for a given session (Matzke et al., 2017b). 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.

3.2.6.2 EEG processing and β-burst detection

For each session, raw data was extracted from the electrode. This signal was then bandpass filtered between 15 and 29 Hz. This signal was then epoched from -1000 ms 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:

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

Time-frequency power estimates were extracted by calculating the squared magnitude of the complex

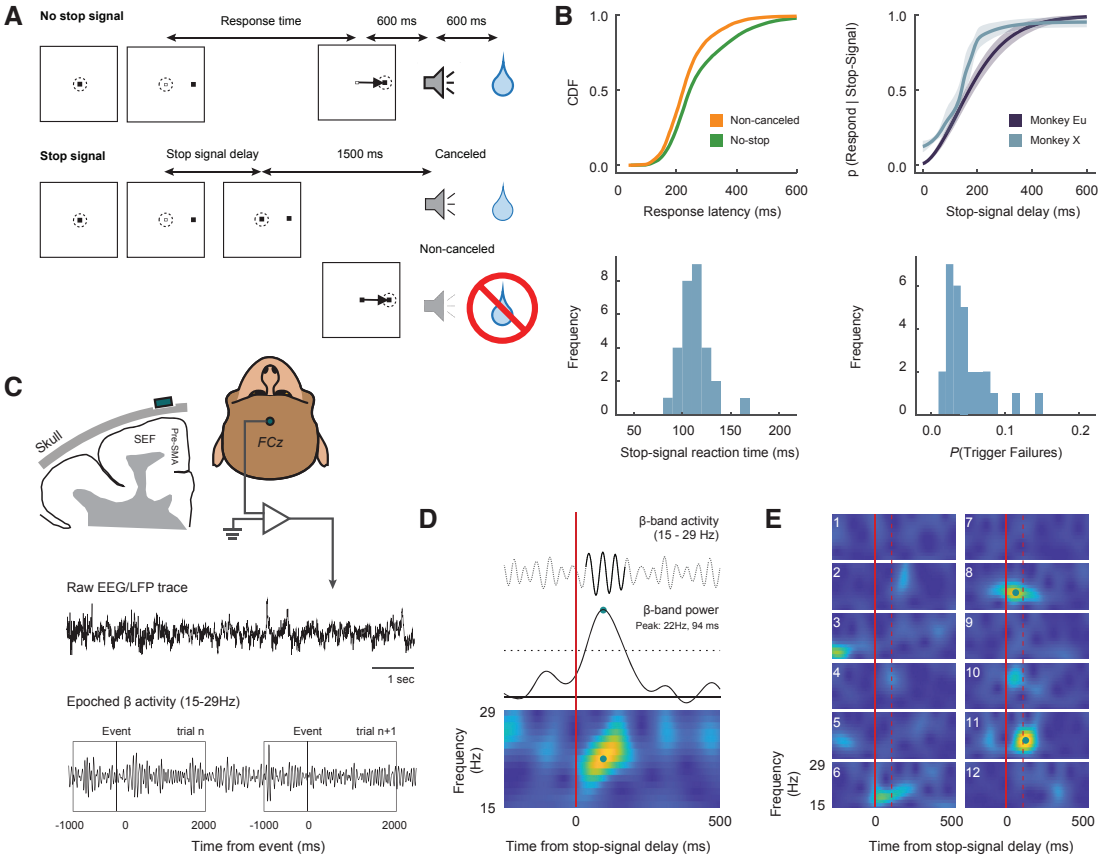

Figure 3.1: **Experimental procedures**. **A**. Saccade-countermanding task. Monkeys initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target was presented simultaneously at one of two possible locations. On no-stop-signal 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 (∼40% of trials), after the target appeared the center of the fixation point was re-illuminated after a variable stop-signal delay, which instructed the monkey to cancel the saccade in which case the same high-pitched tone was presented after a 1,500 ± 0 ms hold time followed, after 600 ± 0 ms by fluid reward. Stop-signal delay was adjusted such that monkeys successfully canceled the saccade in ∼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 stop-signal delays. 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 stop-signal reaction times across sessions. Bottom right: distribution of the proportion of trigger failures across sessions. **C**. LFP processing. Electroencephalogram (EEG) was recorded with leads placed on the cranial surface over the medial frontal cortex at the location analogous to FCz in humans. For each session, raw data was extracted. After bandpass filtering between 15 and 29 Hz, this signal was epoched from -1000 ms 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 6-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 stop-signal delay (solid red), with the corresponding SSRT (dashed red).

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:

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

where μ = 00.0 ms, and σ = 22.5 ms. These values represent the time required for half a cycle at the median β-frequency.

3.2.6.3 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 stop-signal reaction time mean and standard deviation 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-Geiser corrections were applied when assumptions of sphericity were violated.

3.2.6.4 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 stop-signal delay onset and stop-signal reaction time. Whilst stop trials by definition had a predefined stop-signal delay associated with them, no-stop trials were assigned a stop-signal delay 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 prior to the target onset, to -200 ms minus the mean stop-signal reaction time on the given session. For each trial type and time period (baseline & 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 timepoint at which the 95% confidence intervals for the β-burst density functions no longer overlapped. Greenhouse-Geiser corrections were applied when assumptions of sphericity were violated.

3.2.6.5 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 stop-signal delays 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.

3.2.6.6 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 stop-signal delay 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-Geiser corrections were applied when assumptions of sphericity were violated.

3.2.6.7 Linking β-bursts to post-error response time 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.

$$\text{Post-error slowing index} = \frac{RT^{No-stop}_{Non-canceled}}{RT^{No-stop}_{No-stop}}, \text{Post-canceled slowing index} = \frac{RT^{No-stop}_{Canceled}}{RT^{No-stop}_{No-stop}}$$

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 if 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 prior to the target appearing. We then compared if the proportion of β-bursts during this baseline period differed between different trial types using a one-way repeated measures ANOVA.

3.3 Results

3.3.1 Countermanding performance

We acquired 33,816 trials across 29 sessions from two macaques (Eu: 11,583; X: 22,233) performing the saccade stop-signal (countermanding) task (Figure 3.1A). Both monkeys exhibited typical sensitivity to the stop-signal. Summary measures of performance are in Table 3.1. First, response latencies on non-canceled (error) trials were faster than those on no-stop trials (Figure 3.1B, left). Secondly, the probability of failing to cancel and executing an erroneous saccade was greater at longer stop-signal delays (Figure 3.1B, right). These two observations validated the assumptions of the independent race model (Logan and Cowan, 1984), allowing us to estimate the stop-signal reaction time (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; Matzke et al., 2013b), 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: $F(1, 27) = 18.458$, $p < 0.001$, $BF_{10} = 114.778$).

3.3.2 β-bursts and response inhibition

The monkeys' electroencephalogram (EEG) was recorded with a lead placed on the cranial surface over the medial frontal cortex at location analogous to FCz in humans (Figure 3.1C). At the individual trial level, β-band activity was characterized by obvious, burst-like events, rather than by steady changes in modulations (Figure 3.1D). As observed in human studies (Jana et al., 2020; Wessel, 2020), the overall prevalence of

	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

Table 3.1: Stop-signal task performance (mean ± SEM) for both monkeys across all sessions

	No-stop	Non-canceled	Canceled
Baseline	11.6 ± 0.6%	11.7 ± 0.7%	11.2 ± 0.7%
Stopping Period	11.8 ± 0.8%	10.5 ± 0.7%	13.1 ± 0.8 %

Table 3.2: Percentage of trials (mean ± SEM) with β-bursts during a baseline and stopping period, for all trial types.

these bursts was low during both baseline (-400 to -200 ms pre-target, $\sim$12.6 ± 3.8% across all sessions) and task-relevant (0 to 200 ms post-target, $\sim$16.1 ± 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, Figure 3.2A, Table 3.2). 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 to 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 to non-canceled trials (post-hoc test, adjusted $p < 0.001$) but not compared to 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 (Figure 3.2B). A β-burst density function was determined by convolving this discretized array with a gaussian function over time since the stop-signal (Figure 3.2C). 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 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 discharge rates of neurons causally involved in movement initiation and inhibition (Middlebrooks et al., 2020)(Figure 3.2D, bottom panel), β-burst density functions were so small and noisy that no time distinguished canceled from no stop trials in individual sessions (Figure 3.2D, 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 (Figure 3.1B). 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 stop-signal delay. The neurometric function derived from the single neuron discharges of movement-related neurons parallels the inhibition function (Brown and Heathcote, 2008). Although this relationship has only been demonstrated for neural activity that initiates responses, the relationship with neural activity that instantiate 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 stop-signal delay, 50 ms before SSRT, using the conventional threshold of 6x median amplitude plus lower (2x median) and higher (10x median) thresholds. This β-burst neurometric function was compared to the probability of inhibiting a response at each stop-signal delay (Figure 3.3A) 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 6x median threshold was significantly different from zero (one-sample t-test: $t(28) = 6.70$, $p < 0.001$, $BF_{10} = 56120.13$, Figure 3.3B). 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, Figure 3.3B).

3.3.3 β-bursts, error monitoring, & 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 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 ± 0.7%) compared to correct trials (6.7 ± 0.5%) during this period (one-way repeated measures ANOVA: F (1, 28) = 55.103, p $<$ 10-3, BF_{10} = 9068.665, Figure 3.4A).

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 if 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 (Figure 3.4B). 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, Figure 3.4C).

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).

3.4 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 SSD's where inhibition is more successful. This relationship between neurometric and psychometric measures of response inhibition has been observed in the

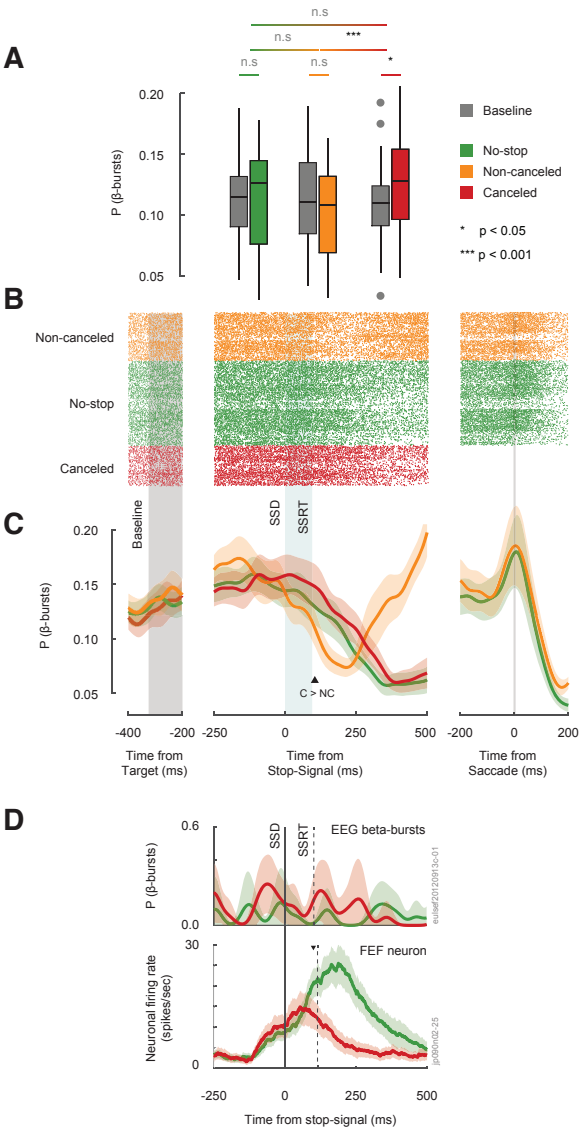

Figure 3.2: **β-bursts during stopping A**. Boxplots showing the incidence of β-bursts observed during the STOP process interval (from scheduled SSD to SSRT) during no-stop (green), non-canceled (orange), canceled (red), and during an equivalent period of time before target presentation (grey). β-bursts are observed in $\sim$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 pre-target 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. Raters 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.

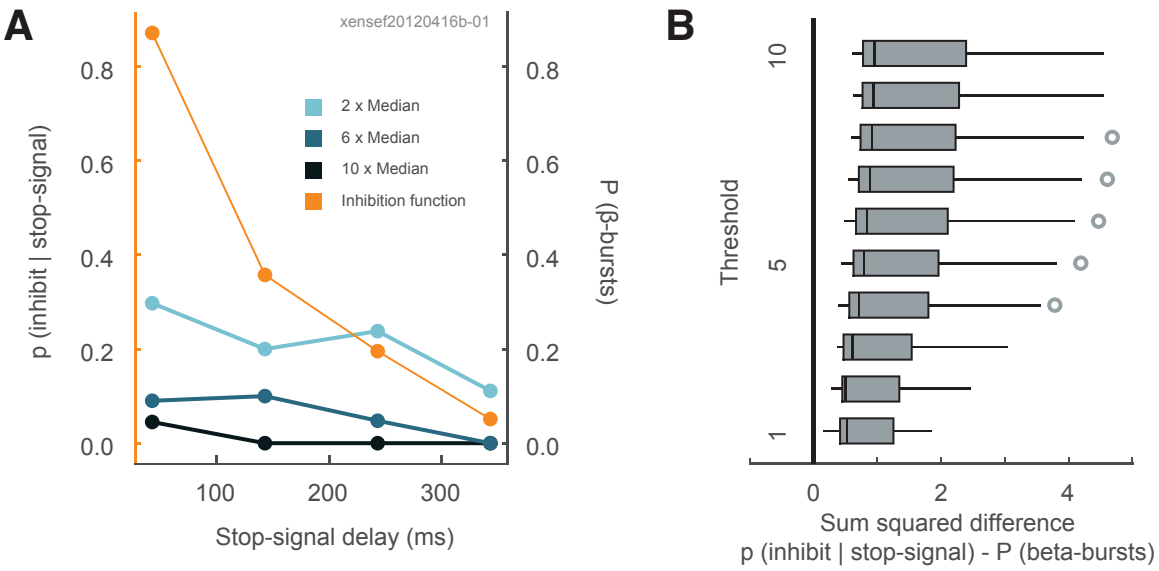

Figure 3.3: **No relationship between β-bursts and response inhibition**. **A**. Primary ordinate, average probability for a representative session of inhibiting on stop signal trials as a function of stop signal delay (orange line). Secondary ordinate, average probability of β-bursts determined with amplitudes exceeding the median threshold by 2 (light blue), 6 (blue), and 10 (dark blue) 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 1 to 10 times the median threshold. Open circles indicate outlier values. From generous to severe measurement thresholds, the incidence of β-bursts did not account for probability of cancellation with stop signal delay.

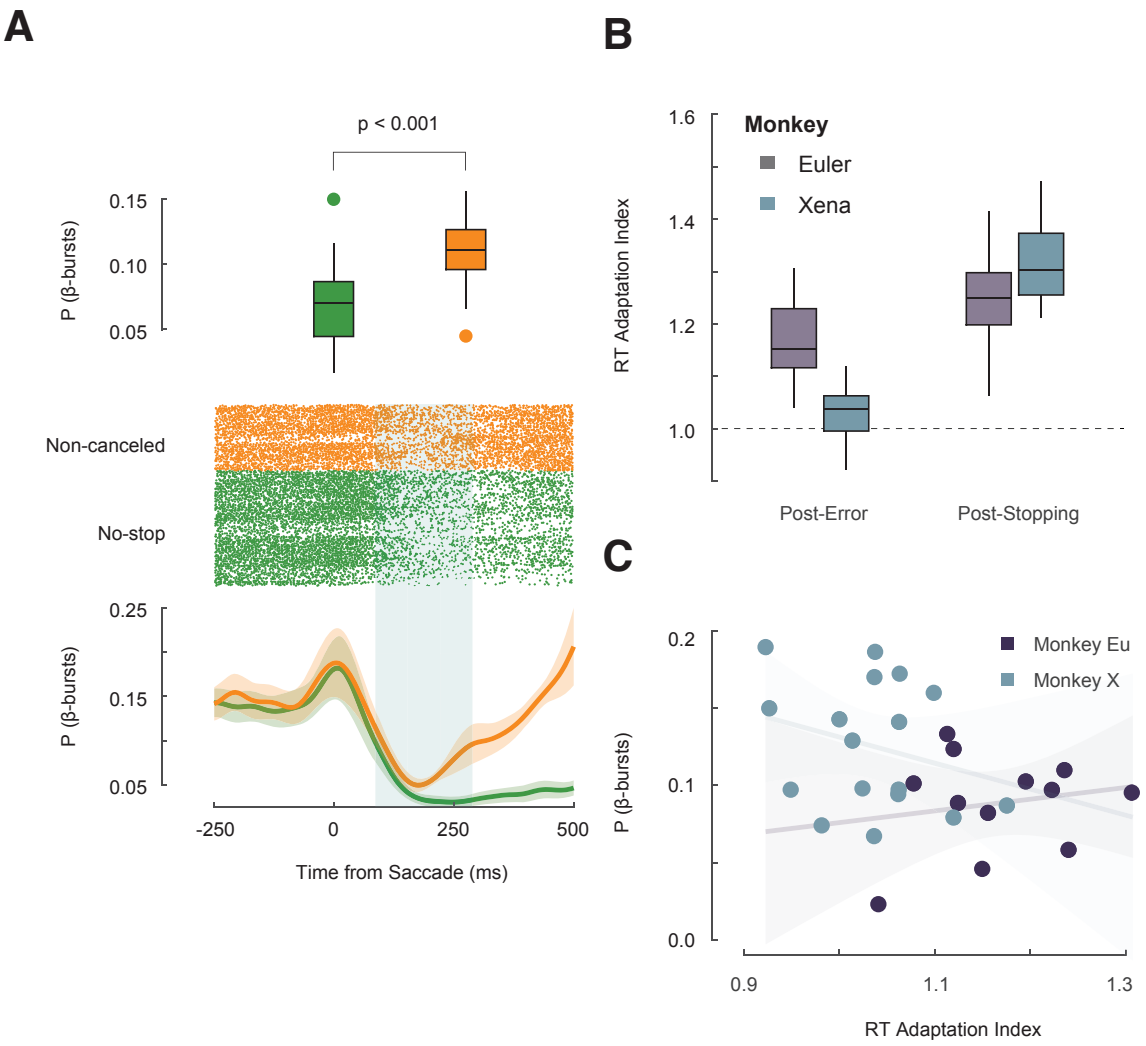

Figure 3.4: **Relationship between β-bursts and performance monitoring**. **A**. β-bursts during error and correct responses. Top: Boxplot showing significantly greater incidence of β-bursts observed 100 to 300 ms following non-canceled (orange) compared to correct saccades (green). Middle: Raster plot 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$150 ms after noncancelled error saccades. textbfB. Boxplots of RT adaptation following non-canceled errors and successful cancellations across sessions for Eu (purple) and X (cyan). Values greater than one represent slowing. As observed previously, both monkeys tend to delay responses somewhat after errors and more after successful stopping. **C**. β-bursts are unrelated to RT adaptation. The incidence of β-bursts observed after errors did not vary as a function of the post-error RT adaptation index across sessions for Eu (purple) or X (cyan). Non-significant regression lines ($\pm$ 95% CI) include 0 slope.

discharges of movement-related neurons in frontal eye fields (Brown and Heathcote, 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 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 pre-supplementary motor area (pre-SMA) contribute to reactive control through the hyper-direct pathway (Aron and Poldrack, 2006; Aron et al., 2007). Whilst the rIFG has been implicated in the attentional capture of the stop-signal & 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 focused on literature that highlights the role of pre-SMA. These previous studies used tasks 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 to 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$5%) compared to what has been previously observed in human ($\sim$20%) studies (Skippen et al., 2019, 2020). However, although response times and SSRT's 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).

Whilst 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 theta (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 to 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-Pallars et al., 2008; Beyer et al., 2012). Interestingly, the magnitude of this β-activity correlated with the theta-activity underlying the ERN (Marco-Pallars et al., 2008).

Whilst 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; Castellar et al., 2010; Godlove et al., 2011; Reinhart et al., 2012b; 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 theta, 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 to between-subjects. Supporting this, a recent large-scale study of $\sim$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-Pallars et al., 2008; Jha et al., 2015), but not with the latency of corrective responses (Marco-Pallars 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 to 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.

CHAPTER 4

Cortical Contributions to Medial Frontal β-Bursts during Executive Control

4.1 Introduction

Cognitive control describes a set of processes, associated with the frontal cortex and related brain networks, that are important when automatic or learned actions are detrimental to achieving a goal (Logan and Cowan, 1984). Inappropriate actions must be stopped before execution to avoid interference with task goals. This feature of cognitive control is referred to as response inhibition. Recent work in humans has proposed that β-bursts in the frontal cortex enable response inhibition (Enz et al., 2021; Hannah et al., 2020; Jana et al., 2020; Wessel, 2020). However, we recently showed that EEG β-bursts occurred too infrequently, did not parallel stopping behavior, and multiplexed other cognitive functions (Errington et al., 2020).

Several explanations have been offered for these divergent observations. First, the rarity of β-bursts and associated lack of causal efficacy for stopping has been attributed to the low signal-to-noise ratio of scalp EEG (Hannah et al., 2020; Jana et al., 2020; Wessel, 2020). Second, although conclusions drawn from human studies have emphasized evidence that medial frontal areas contribute to the interruption of response preparation (Diesburg and Wessel, 2021; Rae et al., 2015; Sharp et al., 2010), the conclusions are refuted by neurophysiological evidence obtained in the medial frontal cortex of macaques (Scangos and Stuphorn, 2010; Stuphorn et al., 2010). Finally, although β-bursts have been observed in EEG contacts directly over areas in the medial frontal cortex, EEG signals reflect complex contributions from multiple areas (Næss et al., 2021).

Although the relationship between intracortical LFP and EEG signals is uncertain, we know that EEG signals measured over the scalp are generated by microcircuits within the cortex (e.g., Herrera, 2020, 2022) such that associations between the activity of particular neurons and EEG signatures can be observed (e.g., Sajad et al., 2019, 2022). Biophysical modeling has suggested that cortical β-bursts can be produced via a transient dipole reversal caused by the interaction between a prolonged excitatory synaptic drive in the basal dendrites of pyramidal neurons that interacts with a transient pulse of excitatory drive in the apical dendrites (Sherman et al., 2016; Shin et al., 2017). However, it remains to be tested whether such layer-specific activity is observed in the medial frontal cortex, and how this contributes to EEG signals recorded from the overlying cranial surface.

We addressed these issues by simultaneously recording LFP across all cortical layers of the Supplementary Eye Field and from EEG electrodes placed on top of the skull in monkeys performing a saccade countermanding task. This task has a strong theoretical framework through which neural activity related to reactive or proactive control can be distinguished (Boucher et al., 2007; Lo et al., 2009; Logan and Cowan,

1984; Logan et al., 2015; Wiecki and Frank, 2013).

Despite the increase in signal-to-noise ratio, we observed intracortical β-bursts as infrequently as those in the EEG. LFP β-bursts were more common during stopping but lacked causal efficacy. LFP β-bursts were also prevalent after execution of the response and during anticipation of reward. LFP β-bursts were most prominent in the upper cortical layers but were not synchronized across the cortical column. Although SEF exhibits the functional, biophysical, and anatomical properties required to contribute to EEG signals (Sajad et al., 2019; Herrera et al., 2020), we found no evidence that EEG β-bursts could be caused by LFP β-bursts. These findings challenge a prevailing hypothesis and constrain models of the functional microcircuitry of β-burst generation during cognitive control.

4.2 Methodology

4.2.1 Experimental models and subject details

Data was 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 countermanding task (Godlove et al., 2014; Hanes and Schall, 1995). All procedures were approved by the Vanderbilt Institutional Animal Care and Use Committee in accordance with the United States Department of Agriculture and Public Health Service Policy on Humane Care and Use of Laboratory Animals.

4.2.2 Animal care and surgical procedures

Surgical details have been described previously (Godlove et al., 2011). 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 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 (Crist Instruments). Chambers were implanted normal to the cortex (Monkey Eu: 17º; Monkey X: 9º; relative to stereotaxic vertical) centered on midline, 30 mm (Monkey Eu) and 28mm (Monkey X) anterior to the interaural line.

4.2.3 Cortical mapping and electrode placement

Chambers implanted over the medial frontal cortex were mapped using tungsten microelectrodes (2-4 MΩ, FHC, Bowdoin, ME) to apply 200 ms trains of biphasic micro-stimulation (333 Hz, 200 μs pulse width). The SEF was identified as the area from which saccades could be elicited using < 50 μA of current (Martinez-Trujillo et al., 2004; Schall, 1991b). In both monkeys, the SEF chamber was placed over the left hemisphere.

A total of five penetrations were made into the cortex; two in monkey Eu, and three in monkey X. Three

of these penetrations were perpendicular to the cortex. In monkey Eu, the perpendicular penetrations sampled activity at site P1, located 5 mm lateral to the midline and 31 mm anterior to the interaural line. In monkey X, the perpendicular penetrations sampled activity at site P2 and P3, located 5 mm lateral to the midline and 29 and 30 mm anterior to the interaural line, respectively. However, during the mapping of the bank of the cortical medial wall, we noted both monkeys had chambers place $\sim$1 mm to the right respective to the midline of the brain. This was confirmed through co-registered CT/MRI data. Subsequently, the stereotaxic estimate placed the electrodes at 4mm lateral to the cortical midline opposed to the skull-based stereotaxic midline.

4.2.4 Data acquisition

Spiking activity and local field potentials were recorded from five sites within the SEF using a 24-channel Plexon U-probe (Dallas, TX) with 150 μm interelectrode spacing allowing sampling from all layers. Penetrations in three of these sites were perpendicular to the cortex. The U-probes were 100 mm in length with 30 mm reinforced tubing, 210 μm probe diameter, 30° tip angle, and had 500 μm 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 dura and entered 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 μm 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). All data were streamed to a single data acquisition system (MAP, Plexon, Dallas, TX). Time stamps of trial events were recorded at 500 Hz. Eye position data were streamed to the Plexon computer at 1 kHz using an EyeLink 1000 infrared eye-tracking system (SR Research, Kanata, Ontario, Canada).

4.2.5 Stop-signal task

The saccade stop-signal (countermanding) task utilized in this study has been widely used previously (Cabel et al., 2000; Colonius et al., 2001; Godlove and Schall, 2016; Hanes and Carpenter, 1999; Hanes and Schall, 1995; Kornylo et al., 2003; Morein-Zamir and Kingstone, 2006; Thakkar et al., 2011; Verbruggen et al., 2019b; Walton and Gandhi, 2006; Wattiez et al., 2016). Briefly, trials were initiated when monkeys fixated at 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 proportion 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 to 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 (stop-signal delay; SSDs). An initial set of SSDs, separated by either 40 or 60 ms, were 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 approximately 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 600ms later, and followed by a high or low fluid reward, depending on the stimulus-response mapping.

4.2.6 Data collection protocol

An identical daily recording protocol across monkeys and sessions was carried out. In each session, the monkey sat in an enclosed primate chair with their head restrained 45cm from a CRT monitor (Dell P1130,

background luminance of 0.10 cd/m^2). The monitor had a refresh rate of 70 Hz, and the screen subtended 46 deg x 36 deg of the visual angle. Eye position data was 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 amalgamated with other behavioral and neurophysiological data. After advancing the electrode array to the desired depth, they were left for 3 to 4 hours until recordings stabilized across contacts. This led to consistently stable recordings. 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 then performed approximately 2000 to 3000 trials of the saccade countermanding (stop-signal) task.

4.2.7 Quantification and Statistical Analysis

4.2.7.1 Cortical depth assignment

The retrospective depth of the electrode array relative to grey 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 pulsated 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 layers (Bastos et al., 2018; Godlove et al., 2014; Maier et al., 2010; Ninomiya et al., 2015; Smith and Sommer, 2013; Westerberg et al., 2019; Xing et al., 2012). Thirdly, an automated depth alignment procedure was employed which maximized the similarity of CSD profiles evoked by passive visual stimulation between sessions (Godlove et al., 2014).

4.2.7.2 LFP processing and β-burst detection

For each session, raw data was extracted from the electrode. This signal was then bandpass filtered between 15 and 29 Hz. This signal was then epoched from -1000 ms 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:

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

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-

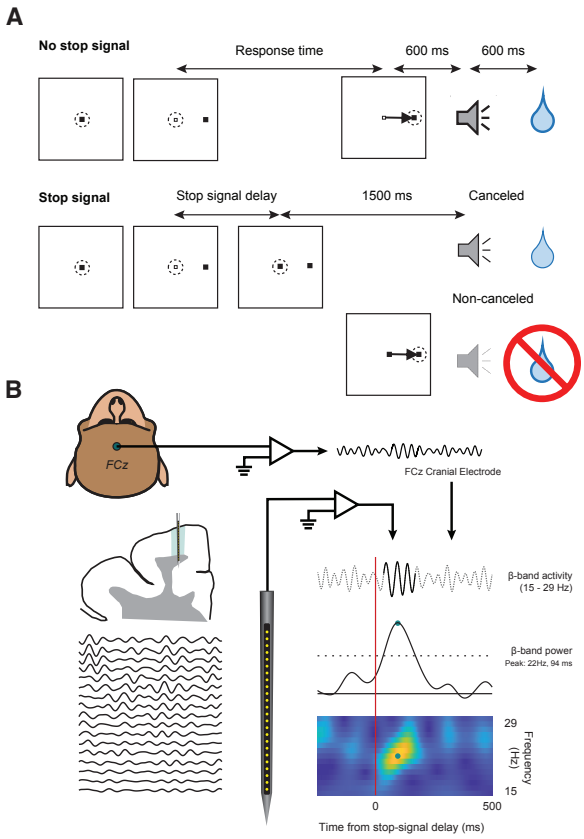

Figure 4.1: **Experimental approach**. **A**. Saccade-countermanding task. Monkeys initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target was presented simultaneously at one of two possible locations. On no-stop-signal trials, monkeys were required to shift their 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 (∼40% of trials), after the target appeared the center of the fixation point was re-illuminated after a variable stop-signal delay, which instructed the monkey to cancel the saccade in which case the same high-pitched tone was presented after a 1,500 ± 0 ms hold time followed, after 600 ± 0 ms by fluid reward. The stop-signal delay was adjusted such that monkeys successfully canceled the saccade in ∼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**. Neurophysiology methods. Electroencephalogram (EEG) was recorded with leads placed on the cranial surface over the medial frontal cortex at the location analogous to FCz in humans. Local field potentials were recorded from all channels from an electrode placed in the Supplementary Eye Field. For each session, raw data was bandpass filtered between 15 and 29 Hz and sampled from -1000 to +2500 ms relative to target presentation, saccade initiation, and stop-signal presentation. To detect β-bursts the signal in each interval 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 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. An example burst is shown in the time-frequency plot.

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:

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

where μ = 00.0 ms, and σ = 22.5 ms. These values represent the time required for half a cycle at the median β-frequency.

4.2.7.3 General analysis considerations

All statistical comparisons were conducted in JASP 0.16.2. When conducting ANOVA's, Mauchly's test of sphericity was used to assess sphericity. Greenhouse-Geisser corrections were employed when this assumption was violated. When multiple comparisons were made, p-values were adjusted using the Holm method. When binning data, we simply recorded whether one or more bursts occurred within the bin; and did not code the quantity of bursts within a given bin. For all analyses, except neurometric, we defined bursts as periods in which the LFP power was greater than 6-times the median power of the entire electrode. This parameter was varied between 1 x and 10 x the median power in the neurometric approach.

4.2.7.4 Bayesian modelling 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. Stop-signal reaction time was estimated using a Bayesian parametric approach (Matzke et al., 2013a,b). Compared to classical methods of calculating SSRT (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 response time (RT) distribution. Furthermore, this model also allows for the estimation of the probability of trigger failures for a given session (Matzke et al., 2017b). 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.

4.2.7.5 Examining the functional properties of medial frontal β-bursts

We used a repeated-measures ANOVA to compare bursts during the stop period (from SSD to SSRT) and baseline period (-200-SSRT to -200 ms, pre-target) across all canceled, non-canceled, and no-stop trials for each individual contact within the cortex (n = 509). We observed the assumption of sphericity was violated (Mauchly's test of sphericity, $p < .05$), so used Greenhouse-Geisser corrections when considering the outcome of our ANOVA analysis. P-values for our post-hoc comparisons were Holm adjusted for comparing a family of 15 (3 trial types, 2 periods).

We also considered how the properties of β-bursts may relate to behavioral measures of stopping. For this we determined several properties of bursts that occurred in the STOP period: from SSD to SSRT. For no-stop and non-canceled trials, we assigned a pseudo-SSD based on the SSD of the previous trial. P(Burst) was calculated as the proportion of trials in which a burst was observed during the period of interest. Onset and offset were determined as the time at which the activity crossed the defined power threshold, relative to the time of the identified peak. The duration was determined as time from onset to offset. Frequency was identified as the frequency at which the β-burst peak was located. We also determined the range of frequencies a burst occupied, by finding the frequencies closest to the burst peak at which the power crossed the defined threshold. Volume was calculated as the duration of the burst, multiplied by the amplitude, multiplied by the range of frequencies that it covered. As the behavioral measure was estimated on a session-by-session basis, we averaged the properties of all bursts observed within a session. To examine the relationship between behavioral variables of stopping (such as the estimated mean SSRT, estimated variance in SSRT, and proportion of trigger failures) and metrics of the burst, we fit a generalized linear model using the MATLAB statistical toolbox, using the burst metric as a predictor.

We also consider bursts more broadly across a wider range of epochs. For this approach compared the proportion of β-bursts that occurred during various epochs between canceled, non-canceled, and no-stop trials. These epochs included: fixation (-400 to -200 ms, pre-target), early post-action (100 to 300 ms, post-SSRT on canceled trials, or 100 to 300 ms, post saccade on no-stop and non-canceled trials), late post-action (-300 to -100 ms, pre-tone on canceled trials, or 400 to 600 ms, post-saccade on no-stop and non-canceled trials), and post-tone (100 to 300 ms, post-tone on all trial types). The p(bursts) was estimated at the individual contact level (n = 509). We compared activity between trial type and epoch using a two-way repeated-measures ANOVA.

4.2.7.6 Neurometric approach

For the neurometric analysis, we calculated the proportion of bursts in the last 50 ms of the stop period on canceled trials (-50 to 0 ms, relative to SSRT). For each SSD, we subtracted the proportion of trials with a

β-burst (neurometric) from the probability of canceling a saccade (psychometric measure). For the shuffled condition, we subtracted the proportion of trials with a β-burst at another SSD other than the one of interest. We then used these differences between the psychometric and neurometric values to calculate the sum of squared differences for both the shuffled and observed condition. These values were then compared using an independent groups t-test, adjusting the p-value for multiple comparisons. We repeated this from a 1x to 10 x threshold of the median baseline power for the given session.

4.2.7.7 Relating β-bursts to proactive control

To probe how baseline p(bursts) varied with the outcome of the previous trial, we compared p(β-bursts) occurring in the same fixation period (-400 to -200 ms) in no-stop trials following no-stop, canceled, and non-canceled trials. We compared activity between preceding trial types using a one-way repeated-measures ANOVA. Finally, we look to see if β-bursts may influence the RT on a given trial. We took trials in which a movement was generated (correctly or incorrectly; no-stop and non-canceled trials, respectively), and then vincentized these trials into six quantiles (i.e., from fast RTs to slow RTs, binned in 20% bins from 0 to 100% of RTs). We then calculated the p(bursts) that occurred within the baseline (-400 to -200 ms, target) and immediately prior to the saccade (-200 to 0 ms, saccade) of the trials within each quantile. We then fit a generalized linear model using the MATLAB statistical toolbox, using the vincentized response time as a predictor.

4.2.7.8 Comparing β-bursts during errors

We compared the proportion of bursts between trials in which a movement was correctly generated (no-stop) and those in which one was incorrectly generated (non-canceled error). We did this for two separate periods following the saccade: an early period (100 to 300 ms, post-saccade), and a late period (400 to 600 ms, post-saccade), for each individual contact within the cortex (n = 509). We compared activity between trial type and early/late error periods using a two-way repeated-measures ANOVA.

4.2.7.9 Relating error related β-burst activity to RT adaptation

We first identified no-stop trials following errors. We then divided these trials into two groups: ones in which the error trial had a burst within the 400 to 600 ms post-error period, and those that didn't. We then identified the degree of post-error slowing by subtracting the RT on the no-stop trial from that on the proceeding non-canceled trial. We then compared whether this change in RT was significantly different between those trials in which a burst occurred, and those that didn't, using an independent groups t-test. For this analysis, the alpha value was set at 0.05/509 (the number of contacts) to account for multiple comparisons. Similar results

were observed at an alpha level of 0.05.

4.2.7.10 Comparing β-power across layers

We computed the power spectral density to determine relative power of the β oscillatory activity across cortical depth (Bastos et al., 2018; Maier et al., 2010; Westerberg et al., 2019). LFP was filtered between 1 and 200 Hz with a 4th-order bidirectional Butterworth filter and full-wave rectified to estimate power at each frequency. Power was normalized relative to the mean of a given column's power at each frequency. We also computed the average β-power during the early post-action period (100 to 300 ms, post-SSRT) for each cortical contact within our laminar penetrations (n = 249 contacts), using the bandpower function of the MATLAB signal processing toolbox. For each session, we then normalized the power by denoting the power for each contact(/depth) as a function of the maximal β-power in the penetration. We then averaged power across all contacts within a given cortical layer. This normalized activity was then compared across sessions, and trial types. We compared activity between trial type and cortical layer using a two-way mixed-measures ANOVA, with layer as an independent factor, and trial type as a repeated factor.

4.2.7.11 Comparing β-bursts across layers

We computed the p(bursts) observed in canceled trials across several epochs for each cortical contact within our laminar penetrations (n = 249 contacts). We focused on four epochs: fixation (-400 to -200 ms, pre-target), early post-action (100 to 300 ms, post-SSRT), late post-action (-300 to -100 ms, pre-tone), and post-tone (100 to 300 ms, post-tone). We then averaged p(bursts) across all contacts within a given cortical layer. We used a two-way mixed-measures ANOVA to compare activity across sessions, with layer as an independent factor, and trial type as a repeated factor.

We then looked at the spatiotemporal pattern of β-bursts by considering the BBDF across layers and times. We averaged the BBDF on canceled trials for each cortical contact within our laminar penetrations (n = 249 contacts), for each epoch of interest. Each contact was assigned to the corresponding cortical depth at which it was recorded. We then averaged this depth aligned BBDF across sessions. Finally, we then normalized the time-depth BBDF by calculating the Z-score value for each depth-time point, relative to the mean and SD of the entire array for the epoch.

4.2.7.12 Linking intracranial β-bursts to EEG β-bursts

We first looked at the p(burst) varied across epochs, dependent on the spatial scale of measuring bursts. We considered bursts in the EEG signal and three levels of the intracortical signal: the probability of observing a burst on individual intracortical contact, the probability of observing a burst on any of the intracortical con-

tacts in each laminar layer (upper or lower), and the probability of observing a burst on any of the intracortical channels. We focused on six epochs: fixation (-400 to -200 ms, pre-target, canceled), target (0 to 200 ms, post-target, canceled), stopping (0 to 200 ms, post stop-signal-delay, canceled), early post-action (100 to 300 ms, post-saccade, non-canceled), late post-action (400 to 600 ms, post-saccade, non-canceled) and post-tone (100 to 300 ms, post-tone).

We then found trials in which bursts occurred in both the EEG and cortex, for each individual contact in our perpendicular penetrations (n = 249 contacts). Once we identified these trials, we determined the temporal difference between the time of the burst peak in the cortex, relative to the time of the EEG burst. We generated a shuffle condition by randomizing the trials from which the LFP bursts were sampled from. Negative values represent bursts in the cortex prior to the EEG burst, and positive values represent bursts in the cortex following the EEG burst. These latency differences were then into 10 ms bins between -250 to 250 ms, across contacts in the upper and lower cortical layers. We then subtracted the shuffled from the observed p(bursts) in each bin.

To probe whether bursts in the cortex led those in the EEG, we calculated the area under the curve by multiplying the bin width (10 ms) by the observed-shuffled p(burst | bin). We then summed bins in the -50 to 0 ms, and 0 to 50 ms period relative to the EEG burst. We whether this area was greater than zero using a one-sample t-test, for both the upper and lower layers.

We then calculated the odds ratio to quantify the relationship between the observation of β-bursts in the EEG and LFP signals. LFP signals were considered at the individual channel level. To achieve this, we counted the number of trials in which bursts were observed in each window for: (a) both the LFP and EEG signal, (b) neither the LFP or EEG signal, (c) just the LFP signal, and (d) just the EEG signal. Based on our analysis, we used 50 ms bins. These bins ranged from -600 to 200 ms relative to the target onset on canceled trials, 0 to 800 ms following the saccade on non-canceled trials, -200 to 600 ms relative to the stop-signal on canceled trials, and -600 to 200 ms relative to the feedback tone on canceled trials. We then conducted a Fisher's exact test to determine whether the bursts were significantly co-incidental. To achieve this, we conducted a right-tailed test reflecting an alternative hypothesis that the odds ratio is greater than 1; this would represent greater co-incidence of a burst in cortex and EEG. This analysis used an alpha level set at 0.003 (0.05/17 bins). We present results as an average of the difference between observed and shuffled bursts, across all contacts in our perpendicular penetration (n = 249 contacts).

Finally, we repeated this analysis comparing the intracortical signals. Across the same windows and epochs as above, we counted the number of trials in which bursts were observed in each window for: (a) both the channels, (b) neither channel, (c) just channel one, and (d) just channel two. We then align the relevant channels to their respective depth in the cortex and, for each epoch, average this activity across all

time bins and all sessions. To determine intralaminar (i.e., within a layer) and interlaminar (i.e. between layers) differences, we averaged the odds ratios within these conditions. We then compared these odds ratios between conditions using a two-way independent-group ANOVA (epoch and laminar-congruence).

4.3 Results

4.3.1 Countermanding performance and neural sampling

We acquired 33,816 trials across 29 sessions from two macaques (Eu: 11,583 trials; X: 22,233 trials) performing the saccade countermanding task (Figure 4.1A). Both monkeys exhibited typical sensitivity to the stop-signal (Supplementary Figure 4.1A). First, response latencies on non-canceled (error) trials were faster than those on no-stop trials. Secondly, the probability of failing to inhibit a saccade was greater at longer stop-signal delays. These two observations validated the assumptions of the independent race model (Logan & Cowan, 1984), allowing us to estimate the stop-signal reaction time (SSRT). Monkeys had indistinguishable mean SSRT (one-way independent measure ANOVA: $F(1, 27) = 0.108$, $p = 0.745$) and variance of SSRT (one-way independent measures ANOVA: $F(1, 27) = 0.819$, $p = 0.819$). However, trigger failures were significantly more common for monkey Eu relative to monkey X (one-way independent measures ANOVA: $F(1, 27) = 18.458$, $p < 0.001$, BF10 = 114.778).

Cortical local field potentials were sampled across 509 contacts (Eu: 217 contacts, X: 292 contacts), across 29 sessions (Eu: 12 sessions; X: 17 sessions), from electrodes placed in Supplementary Eye Field. In 16 of these sessions (Eu: 6; X: 10), penetrations were made perpendicular to the cortical surface. This allowed for the assignment of activity to specific cortical layers. Simultaneous to these recordings within the cortex, electroencephalogram (EEG) voltage was recorded with a lead placed on the cranial surface over the medial frontal cortex (Figure 4.1B). The placement of these electrodes was analogous to FCz in humans. Like previous observations in EEG, the overall prevalence of cortical β-bursts was low during both baseline (-400 to -200 ms pre-target, 21.9 ± 0.6% across all sessions) and task-relevant (0 to 200 ms post-target, 22.8 ± 0.6% across all sessions) periods. Our first original finding is that the proportion of β-bursts observed within the cortex (21.9 ± 0.6%) mirrored those observed in EEG recorded on the cranial surface (baseline: 21.5 ± 0.9%, target: 27.9 ± 1.4%) during the baseline period of valid trials.

4.3.2 Functional properties of intracortical β-bursts during executive control

Previous work has postulated that an increased incidence of β-bursts during the stopping period on canceled trials suggests a possible contribution to action-stopping. However, their role in error monitoring, overall sparsity, and prevalence in trials in which movements were executed warrants wider consideration of alternative functional contributions. Here, we more widely consider the pattern of medial frontal β-bursts across several

task epochs reflecting different aspects of cognitive control (Figure 4.2A, Supplementary Figure 4.2A). We observe that the incidence of β-bursts clearly fluctuates over different task epochs and trial types, demonstrating notable variation in their functional contributions (Two-way repeated measures ANOVA, trial type x epoch: F (3.212, 1631.471) = 243.221, $p < 0.001$).

We first aimed to resolve the extent to which MFC β-bursts contributed to the active and direct inhibition of a planned response. Intracortical β-bursts were significantly less common during the STOP process compared to a pre-target baseline (Figure 4.2B; two-way repeated measures ANOVA, main-effect of epoch: F (1, 508) = 226.52, $p < 0.001$). In conjunction with this, although observed in equal incidence during the baseline, β-bursts were significantly reduced during the potential stopping interval on trials in which a movement was generated, but not in trials in which a movement was inhibited (two-way repeated measures ANOVA, interaction between epoch and trial type: F (1.794, 911.151) = 31.812, $p < 0.001$). Mirroring previous observations in EEG (Enz et al., 2021; Hannah et al., 2020; Jana et al., 2020; Wessel, 2020), we found β-bursts were significantly more common during the STOP period when a planned saccade was inhibited (canceled trials), compared to when an erroneous movement was generated (non-canceled, adjusted $p < 0.001$), and to a comparable period in correct trials (no-stop, adjusted $p < 0.001$) trials. Whilst β-bursts were more common when stopping was successful, the relationship between the occurrence of β-bursts and behavioral measures of stopping was weak and unreliable across monkeys (Supplementary Table 1).

Although the proportion of β-bursts was significantly different between trial types, there was no reliable relationship across monkeys between features of the β-burst (onset, offset, duration, frequency, volume) and trial types (Supplementary Table 2, Supplementary Figure 4.2B). In addition, the prevalence of β-bursts was unchanging across stop-signal delays and did not reflect changes in the probability of inhibiting a saccade. The response inhibition function shows the fraction of non-canceled trials increases as a function of stop signal delay, as movements become less likely to be canceled as movement preparation progresses (Figure 4.2C, black line). In conjunction with the psychometric function, a neurometric function can also be derived. Previous work has achieved this by inspecting the single neuron discharges of movement-related neurons

[1]**Figure 4.2.** **A**. Mean ± SEM of the probability of LFP β-bursts across different epochs on canceled (red), non-canceled (yellow), and no-stop (green) trials. LFP β-bursts were more common in stop-signal trials and most common after errors. **B**. Mean ± SEM of LFP β-bursts observed during the STOP process interval between actual or scheduled SSD and SSRT during no-stop, non-canceled, canceled trials compared with an equivalent interval during foreperiod before target presentation (unsaturated). LFP β-bursts are more common in the foreperiod but observed in <15% of trials with slightly but significantly higher incidence when saccades are inhibited. **C**. Mismatch between p(inhibition) across SSD (black, left ordinate scale) for a representative session and neurometric functions across SSD (right ordinate scale) of p(β-burst) with median detection thresholds of 2x (darkest), 4x (intermediate), or 6x (lightest) baseline value. In no sessions were LFP β-bursts observed commonly enough to account for response inhibition. **D**. Mean ± SEM of p(β-bursts) during the foreperiod in no-stop trials following canceled, non-canceled and no-stop trials. LFP β-bursts were more common after error trials. **E**. Relationship between mean p(β-burst) in 200 ms preceding saccade initiation and means of session-wise response time quantiles, collapsed across all sessions. Regression line with CI demonstrates significant negative association. **F**. Mean ± SEM of bootstrapped effect sizes comparing the proportion of bursts in no-stop trials to canceled and non-canceled trials during fixation, stopping, early post-action, and late post-action epochs. LFP β-bursts were most common after responses, and evidence was weak for a higher incidence of LFP β-bursts during response inhibition. **G**. Histogram of change in RT between successive non-canceled and no-stop trials when an LFP β-burst did or did not occur during the error period. Slowing or speeding of RT after errors was unrelated to the occurrence of LFP β-burst after errors.

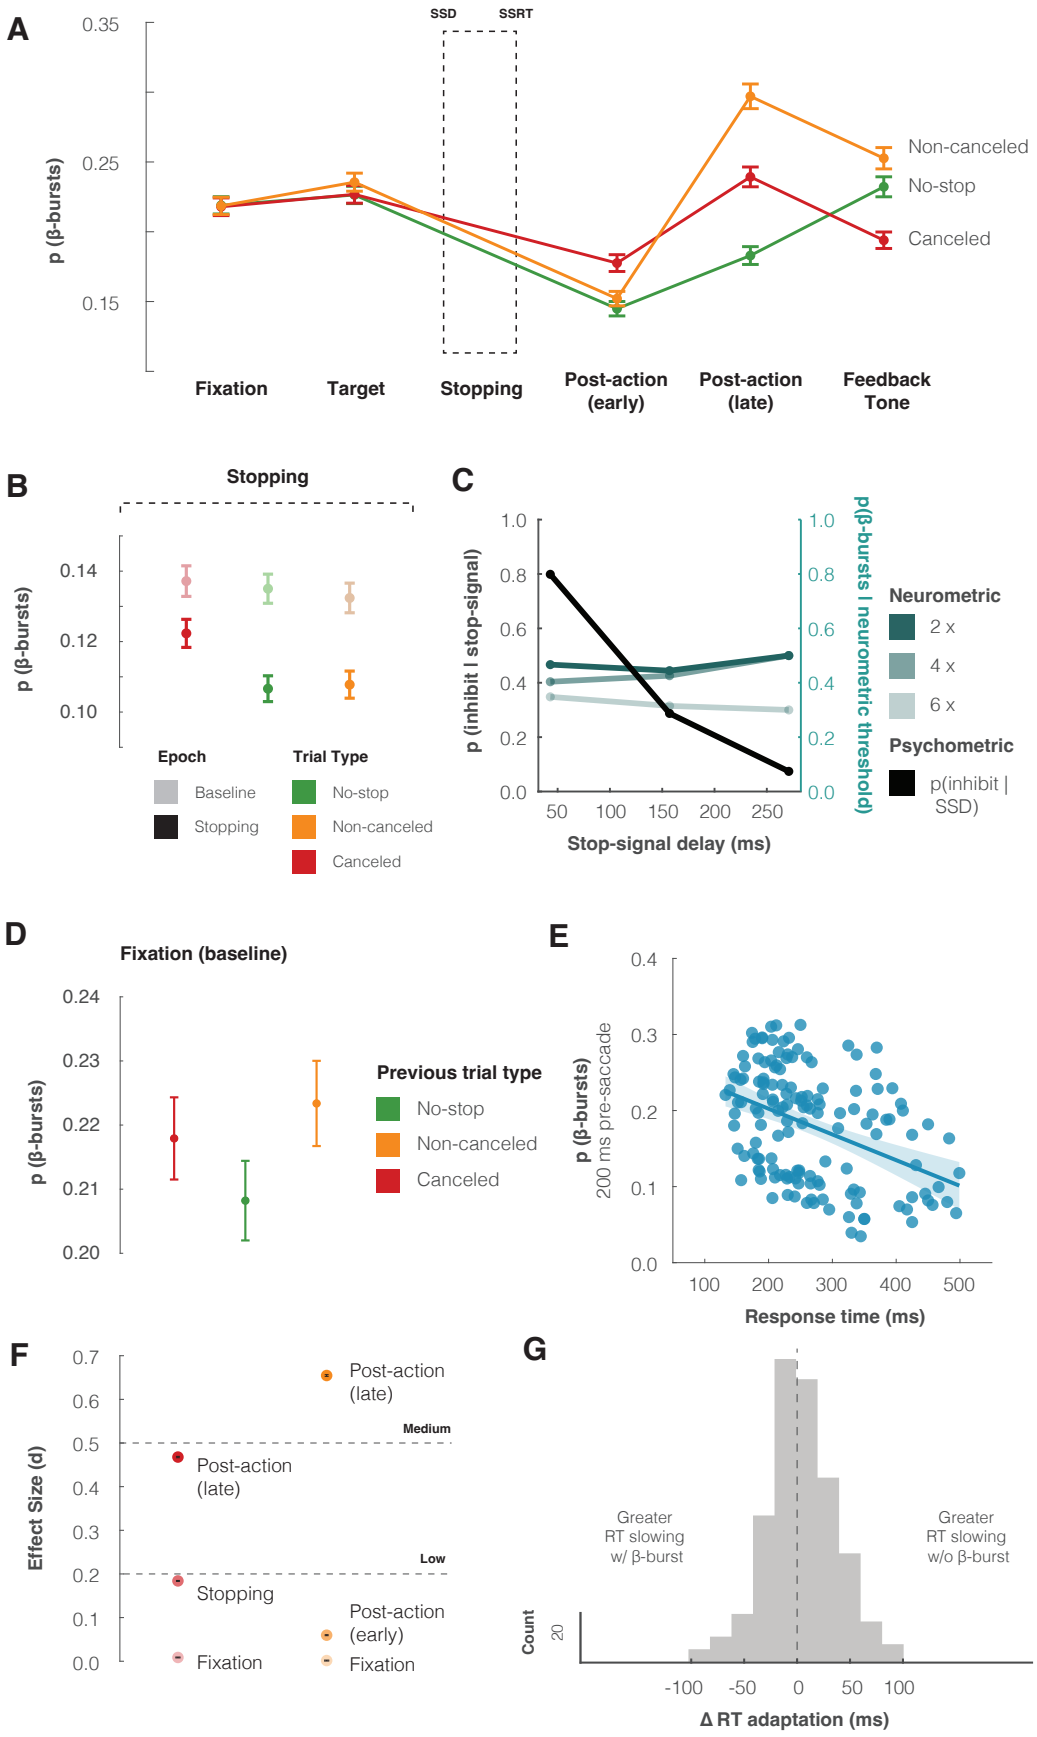

Figure 4.2: **Functional properties of intracranial β-bursts**. See footnote [1]

in motor structures, such as the frontal eye fields (FEF), and determining the probability of single neurons modulating within the stopping period. Just like the psychometric function, this neurometric function also increases as a function of stop-signal delay (Brown et al., 2008). To examine if this relationship occurs with β-bursts, we measured the number of β-bursts observed in the last 50 ms of SSRT at several thresholds for each channel, on canceled trials at each stop-signal delay. We then calculated the sum of their squared differences relative to the probability of inhibiting a response at each stop-signal delay (Figure 4.2C). To generate a valid null distribution, we repeated this process, but this time shuffling the p(β-bursts) observed at each SSD, before calculating the sum of squared differences. Across all channels, the summed squared differences were not significantly different from the shuffled condition at the 6x median threshold (two-sample t-test: t (1016) = 0.101, p = 0.920). This conclusion did not depend on the β-burst measurement threshold (Supplementary Figure 4.2C, two-sample t-test at each threshold, $p > 0.05$ for all thresholds after corrections for multiple comparisons). Thus, β-bursting during the stopping period is unrelated to stopping behavior.

Second, we then considered that bursts may instead reflect another function supporting the ability to successfully stop, rather than the process of response-inhibition itself. We approached this from two perspectives: (1) activity before stopping that may influence or bias the stop process, and (2) activity following response inhibition that maintains inhibition until the response is validated and the goal is achieved (until the feedback tone is elicited).

To address the first perspective, we looked at how baseline activity could proactively influence the stop process. To do this, we compared the proportion of β-bursts that occurred during the fixation period between those trials which would ultimately progress to become inhibited, those which would lead to the generation of a correct saccade, and those which resulted in the generation of an erroneous saccade. β-bursts were not distinguishable in the fixation period of canceled, non-canceled, and no-stop trials (Figure 4.2A, fixation, Holm post-hoc test, adjusted $p > 0.999$ for all comparisons). This observation was reliable across monkeys (Supplementary Figure 4.2A). Second, we then looked at how baseline p(bursts) varied with the outcome of the previous trial. We compared p(β-bursts) occurring in the pre-target period (-400 to -200 ms) in no-stop trials following no-stop, canceled, and non-canceled trials. We found a small, but significant, reduction in p(β-bursts) following no-stop trials, compared to non-canceled and canceled trials (Fig 2D, one-way repeated measure ANOVA, F (1.836, 932.657) = 14.043, $p < 0.001$). Finally, we considered the proposition that β-bursts may influence the go process (Muralidharan et al., 2022) by looking at the relationship between β-bursts in saccade trials (both no-stop and non-canceled) and their relationship with RT. To do this, we looked at the proportion of β-bursts that occurred in two periods: (1) baseline/fixation and (2) pre-saccade. Although we found no relationship between baseline p(β-bursts) and RT (R2 = 0.012, p = 0.161), we did observe that the p(β-bursts) within the 200 ms pre-saccade window were increasingly less common as RT's became longer

(Figure 4.2E; R2 = 0.169, p $<$ 10-5). These findings were consistent across monkeys (Supplementary Figure 4.2D).

Concerning the second perspective, we then examined patterns of β-burst activity following the period of reactive inhibition (between SSD and SSRT). To successfully complete the trial and receive a juice reward, monkeys had to maintain fixation on the central stop-signal following successful inhibition until a tone was sounded. During this period on canceled trials, we observed an increase in β-burst incidence following successful inhibition, compared to when a movement was correctly generated (Figure 4.2A, early post-action, Holm post-hoc test, adjusted p $<$ 0.001). This elevated bursting rate was maintained until the period immediately preceding the feedback tone (Figure 4.2A, late post-action, Holm post-hoc test, adjusted p $<$ 0.001). To determine whether this activity was stopping specific, we compared p(β-bursts) in this period to that in the fixation period. Here, the behavior is the same as that following successful stopping: the monkey was maintaining fixation at the same central point and waiting for the presentation of the next task cue – the presentation of a target for the fixation period, and the sounding of a tone for the stopping period. Interestingly, we observed a differential pattern of activity. Whilst β-bursts were observed in equal proportions on canceled and no-stop trials during the fixation period (Figure 4.2A; fixation, adjusted p $>$ 0.999), they were significantly more common on canceled trials in the early and late post-action periods (Figure 4.2A, early post-action, adjusted p $<$ 0.001; late post-action).

We quantified the magnitude of the difference in the p(bursts) between canceled and no-stop trials by calculating the effect size at three key epochs: during fixation, during stopping, and during the late post-action period. Although there were no differences between trial types during fixation, this developed into a low effect size during the stopping period, and a medium effect size during the late post-action period (Figure 4.2F, two-way repeated-measures ANOVA, epoch x control signal, F(1.924, 190.468) = 7349.496, p $<$ 0.001). Interestingly, β-bursts were less common on canceled trials and more common on no-stop trials following the feedback tone (Figure 4.2A, Holm post-hoc test, adjusted p $<$ 0.001). The increase of β-bursts in the early and late post-action period was consistent across monkeys (Supplementary Figure 4.2A).

In addition to response inhibition, the stop-signal task is useful for exploring performance monitoring as, by design, errors occur in $\sim$50% of stop-signal trials. Using this task, we previously reported that EEG β-bursts increase following errors (Errington et al., 2020). These findings mirrored previous work demonstrating increased spiking activity in SEF, occurring after errors (Purcell et al., 2012b; Sajad et al., 2019; Stuphorn et al., 2010). Interestingly, contrasting our previous findings from EEG, we did not observe a reliable difference in the incidence of β-bursts in the 300 ms period following the execution of an erroneous or correct saccade. Instead, we observed a clear increase in β-bursts on error trials in the late post-action period leading to the auditory feedback tone (Figure 4.2A). Congruent with this, we found no reliable difference in

the proportion of β-bursts between error and correct saccade trials in the early post-action period between monkeys (Supplementary Figure 4.2A, early post-action period, adjusted p = 0.116), but found β-bursts were significantly more common for both following errors in late post-action period (Supplementary Figure 4.2A, late post-action period, adjusted $p < 0.001$).

Much like we did for trials in which inhibition was successful, we quantified the magnitude of the difference in the p(bursts) between non-canceled and no-stop trials by calculating the effect size at three key epochs: during the fixation, early post-action, and late post-action periods. There were no notable differences between error and correct trials during both fixation and early-action periods. This developed into a medium-to-large effect size during the late post-action period (Figure 4.2F, two-way repeated-measures ANOVA, epoch x control signal, F (1.924, 190.468) = 7349.496, $p < 0.001$). Interestingly, in contrast to canceled trials, β-bursts were equally as common on non-canceled trials and no-stop trials following the feedback tone (Figure 4.2A, Holm post-hoc test, adjusted $p < 0.001$).

To examine if these signals may be used to recruit control in future trials, we determined whether the presence of a β-burst in the late post-action period on error trials (trial n) influenced the RT on a following no-stop trial (trial n+1). For each channel, we compared the change in RT from non-canceled error trials to no-stop trials, between error trials with a β-burst present and those without a β-burst present. Here, negative values represent post-error slowing and positive represents an increase in RT's following errors. We then calculated the difference in the RT adaptation between those trials with a burst against those without a burst. This provides a measure of the difference in RT adaptation between burst/no-burst trials. Here positive values represent greater post-error slowing when a burst was observed. In our sample, no channels demonstrated a relationship between bursting and RT adaptation (Figure 4.2G). Thus, β-bursting during the error period is unrelated to RT adaptation.

4.3.3 Laminar patterns of β-band activity

To investigate the cortical mechanisms contributing to β-bursts, we examined their profile across cortical layers of SEF from 16 sessions for which we had confidence in the layer assignments (Monkey Eu, 6 sessions; Monkey X, 10 sessions). The details of this procedure have been previously described (Godlove et al., 2014; Ninomiya et al., 2015).

We first used power spectral density (PSD) to look at the relative power of frequencies across all trials, and across all cortical depths. Qualitatively, we replicate previous findings that β-activity was most powerful in the middle-to-lower cortical layers, around the L3/5 divide in SEF (Supplementary Figure 4.1B). This was observed across both monkeys (Supplementary Figure 4.1B). To look at functional differences, we then compared β-power across all cortical layers on canceled trials relative to latency-matched no-stop trials during the

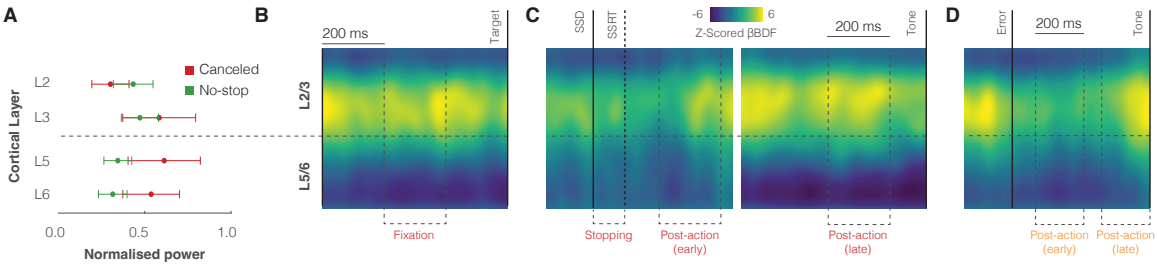

Figure 4.3: **Laminar properties of intracranial β-bursts.** **A**. Normalized power in each cortical layer during response inhibition for canceled and latency-matched no-stop trials. The elevation of LFP β power during response inhibition happened in the deep layers. Conventions as in Figure 4.1. **B**. Heat maps of LFP β-burst incidence across time and cortical layers on canceled trials during foreperiod (left), stopping and early post-action (left-center), late post-action period (right-center), and error (right) epochs. Yellow represents a greater proportion of (β-bursts) at a given time and depth, relative to other times and depths.

STOP process. We found that normalized β-power was significantly greater in L5 & L6 on canceled trials, but not in L2 & L3 (Figure 4.3A, two-way mixed-measures ANOVA, trial type x cortical layer interaction: $F(1, 247) = 6.633$, $p = 0.011$). However, we note this effect was driven by one monkey (Supplementary Figure 4.3A).

We then progressed to look at the pattern of β-bursts, rather than power, through cortical layers across different functional epochs within our task. We compared the proportion of β-bursts observed in each layer during fixation, stopping, early, and late post-action periods on canceled trials, and the late post-action period on noncanceled trials. Although variable at the individual layer level between monkeys, we observed bursting was reliably more prevalent in L3 for both monkeys (Supplementary Figure 4.3B, two-way mixed-measures ANOVA, main effect of cortical layer: $F(3, 245) = 27.734$, $p < 0.001$).

To examine the temporal profile of this activity, we generated time-depth β-burst density functions which allow us to see the pattern of β-bursts across cortical layers relative to different task events. Given the higher incidence of beta-bursts during these epochs, we examined the laminar pattern of bursts during fixation, stopping, and in the late post-action period on canceled trials, and the early and late post-action period on noncanceled error trials. For each session, this resulted in a β-burst density function over time and depth which was then z-scored relative to the mean value across all depths and during the entire epoch; higher Z-scores represent a greater proportion of β-bursts observed relative to other times and depths. Using this approach, we observed that β-bursts were qualitatively more prevalent in upper layers compared to lower layers during fixation, stopping, and errors (Figure 4.3B). This effect was observed in both monkeys (Supplementary Figure 4.3C).

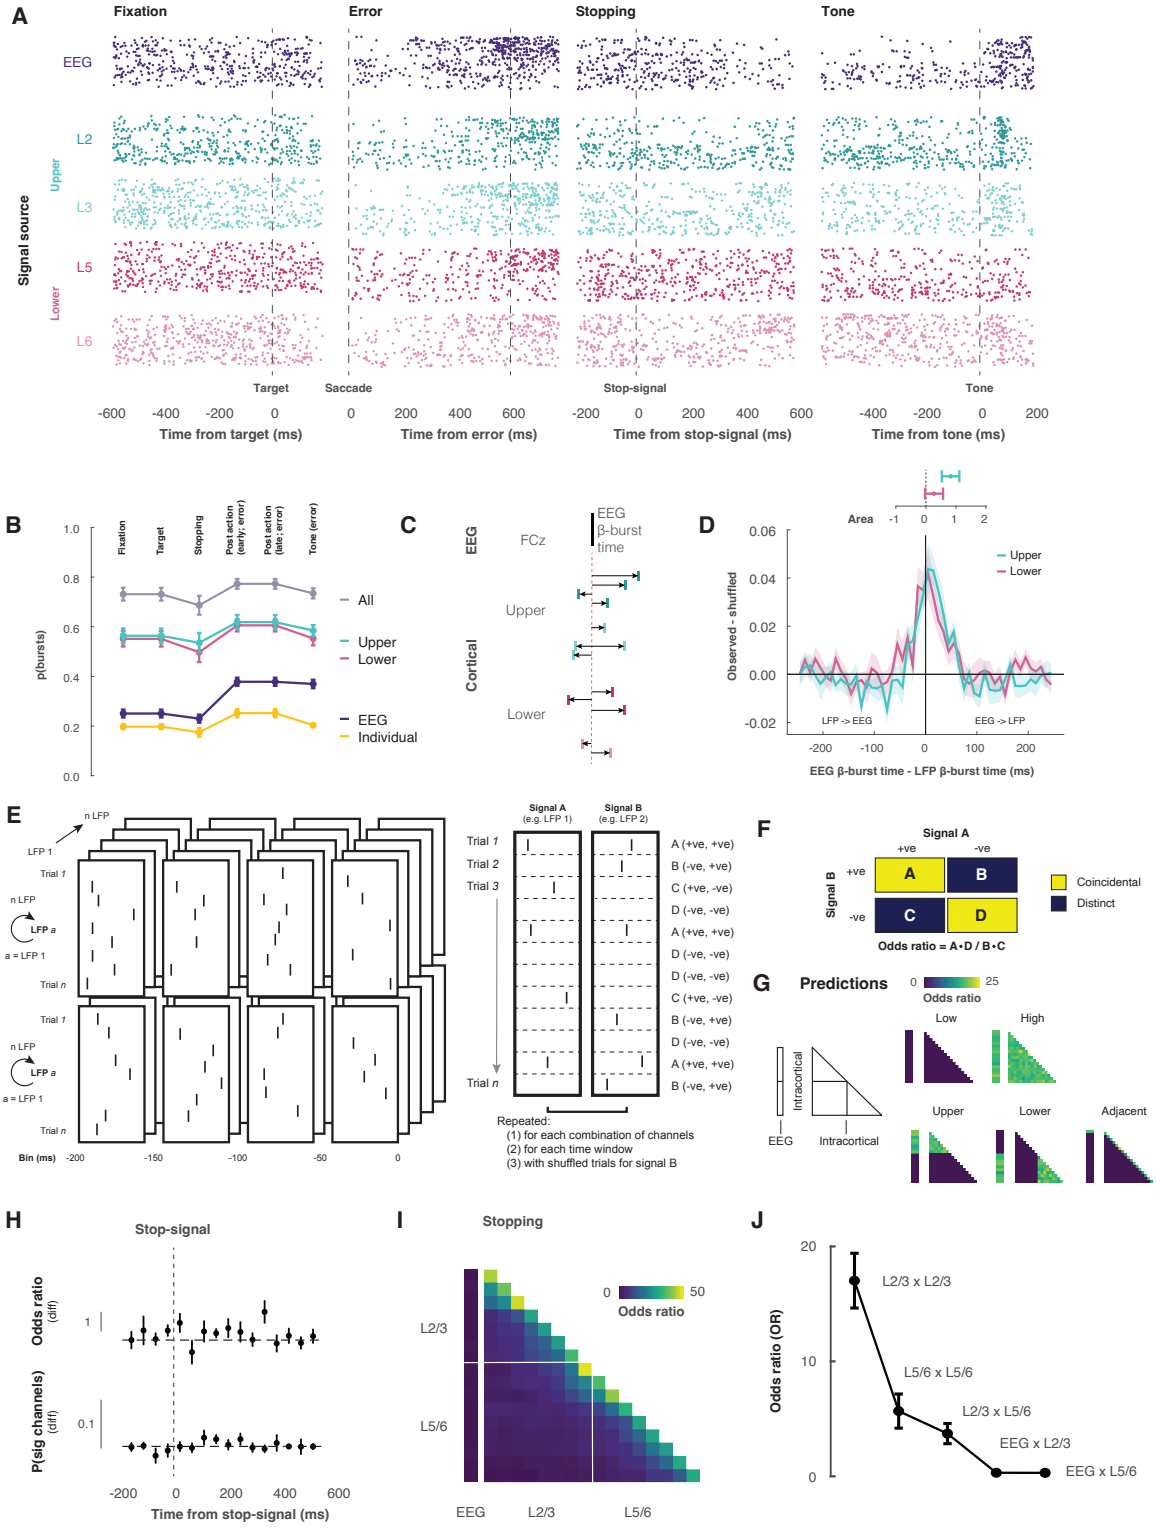

Figure 4.4: **Comparing EEG and intracranial β-bursts.** See footnote [2]

4.3.4 Weak association between intracortical and cranial β-bursts

After establishing the patterns of β-bursts activity within SEF, we then determined the extent to which this activity could contribute to β-bursts recorded on the cranial surface through several approaches. We first compared the pattern of bursts observed in the cortex with the pattern of bursts observed in the simultaneously recorded EEG signal (Figure 4.4A). We quantified the proportion of bursts in each epoch in each electrode channel, across all channels in the upper and lower layers, and across all channels in the cortex. In the limit, EEG β-bursts could be generated by β-bursts observed simultaneously in all sampling channels across all cortical layers. Alternatively, EEG β-bursts could be generated by β-bursts observed simultaneously in some fraction of sampling channels across cortical layers. In the other limit, EEG β-bursts could appear sporadically with no relationship to β-bursts arising randomly in sampling channels across cortical layers.

By simply counting the number of β-bursts in the EEG electrode and each LFP channel in each task epoch, we found that on average they were infrequent on any individual LFP channel and only slightly more frequent in the EEG after saccades (Figure 4.4B; Supplementary Figure 4.4A). When summed within L2/3 and L5/6 or across all layers, the incidence of β-bursts was naturally higher. The sum of β-bursts within L2/3 and L5/6 or across all LFP channels exceeds by 3 or 4 times, the number of β-bursts in the EEG channel. This reveals unexpected complexity in the biophysical relationship between LFP and EEG β-bursts.

We then quantified the temporal relationship between LFP and EEG β-bursts (Figure 4.4C) and found that β-bursts were significantly more likely in both L2/3 and L5/6 in the ±50 ms interval around the EEG β-burst (Figure 4.4D, main). LFP β-bursts were more likely to occur after EEG bursts, than before (Figure 4.4D, inset; one-sample t-test (two-tailed), $t(15) = 2.167$, $p = 0.047$). Although we observed between-monkey differences (Supplementary Figure 4.4B), we reliably observe this finding for bursts in upper relative to lower cortical layers (Lower: one-sample t-test (two-tailed), $t(15) = 0.904$, $p = 0.380$; Upper: one-sample-test (two-

[2]**Figure 4.4. A**. With each row representing a trial from a representative session, each point in the raster marks the time EEG (purple) the LFP β-bursts in L2 (dark teal), L3 (light teal), L5 (dark magenta), L6 (light magenta) during fixation (left), error (left-center), stopping (right-center), and canceled late post-action (right) epochs. Note the lack of columnar synchrony with density after errors and with feedback. **B**. Mean ± SEM of p(β-burst) counted within epochs in individual LFP (gold) and EEG (purple) contacts and summed across contacts in L2/3 (teal), in L5/6 (magenta), and from L2 to L6 (grey). The number of LFP β-bursts within a cortical column exceeds the number of EEG β-bursts. **C**. Diagram of how the coincidence of EEG and LFP β-bursts was measured. **D**. Distribution after shuffle subtraction of shortest times between EEG β-bursts and LFP β-bursts detected in L2/3 (teal) and L5/6 (magenta) plotted with positive (negative) values for LFP β-bursts following (preceding) EEG β-bursts. Only $\sim$5% of EEG β-bursts were associated more than chance with LFP β-bursts, and this association was observed only within ±50 ms of the EEG β-burst. Mean ± SEM area under the distributions across sessions, plotted above, indicates a weak tendency for LFP β-bursts in L2/3 but not L5/6 to follow EEG β-bursts. **E**. Diagram of how odds ratio (OR) of presence or absence of EEG and LFP β-bursts in non-overlapping 50 ms intervals surrounding task events. The presence or absence of β-bursts was counted pairwise across every EEG and LFP channel for each interval of all trials. **F**. Odds ratios were calculated from the counts of the four possible combinations of presence or absence of β-bursts in each pair of channels were tabulated: β-burst present in both (A), in one channel but not the other (B, C), and β-burst absent in both (D). **G**. Predicted odds ratios were calculated in the matrix of EEG and LFP sampling contacts for simulations of coincident presence or absence of β-bursts occurring rarely (low) or frequently (high) across all channels or occurring frequently only in the upper, lower, or adjacent channels. **H**. Mean ± SEM of the shuffle-corrected odds ratio of detecting an LFP β-burst ±50 ms from the occurrence of an EEG β-burst (above) and of the probability of detected significant odds ratio > 1 (below) during response inhibition. Coincident presence or absence of β-bursts across EEG and LFP channels was vanishingly rare. **I**. Color map of odds ratios of LFP and EEG β-bursts shows association in adjacent but not distant contacts. **J**. Mean ± SEM odds ratios of coincidences of β-bursts occurring within L2/3 LFP channels, within L5/6 LFP channels, across L2/3 and L5/6 LFP, and across EEG and LFP in L2/3 and L5/6. The odds ratio scaled with prevalence and was effectively nil across EEG and LFP channels.

tailed), t(15) = 2.821, p = 0.012).

Finally, having determined the ±50 ms window of association of EEG and LFP β-bursts, we quantified the coincident presence or absence of EEG and LFP β-bursts. This was achieved by identifying bursts trial-by-trial in non-overlapping 50 ms windows around the event of interest for every LFP contact within the cortex and the EEG signal (Figure 4.4E. left). For each possible pair of contacts trial-by-trial we counted how often a β-burst co-occurred, occurred in one channel or the other, or neither channel (Figure 4.4E, right) and calculated the odds ratio (Figure 4.4F). This odds ratio provides a measure of association between observing a burst in one channel and another burst in a separate channel. To test alternative hypotheses about β-burst production, we simulated low and high coincidence between all contacts, only between contacts in the upper or lower layers, and only between adjacent contacts (Figure 4.4G).

Relative to shuffled control, no periods of significant association between the proportion of EEG and LFP β-bursts were observed in any task epoch (Figure 4.4H; Supplementary Figure 4.4C, D). We also considered the intra-contact similarity to probe the coincidence of bursts between and within cortical layers (Figure 4.4I; Supplementary Figure 4.4E). β-bursts coincided significantly more on adjacent contacts within layers than across layers (Figure 4.4J; Supplementary Figure 4.4F, two-way independent-groups ANOVA, main effect of laminar comparison: F (3, 300) = 103.151, $p < 0.001$). This relationship was more pronounced in upper cortical layers compared to lower cortical layers (adjusted $p < 0.001$) and did not vary between epochs (interaction between epoch and laminar comparison: F (3, 300) = 0.102, $p > 0.999$). Finally, the OR of the association between LFP and EEG β-bursts was significantly smaller than the OR of the association between β-bursts within the cortex (Figure 4.4J, adjusted $p < 0.001$ for all comparisons). These results demonstrate that β-bursts are not produced concomitantly within a cortical column but instead arise at low and stochastic rates relatively independently across cortical depth.

4.4 Discussion

4.4.1 Functional properties of intracranial medial frontal β-bursts

Previous work in humans (Wessel, 2020) and macaques (Errington et al., 2020) has described a higher prevalence of EEG β-bursts recorded over the medial frontal cortex (MFC) during response inhibition. Like other reports of EEG β-bursts in the frontal cortex (Jana et al., 2020; Hannah et al., 2020; Enz et al., 2021), the incidence of these bursts was low and their association with stopping behaviors was weak. This lack of causal efficacy is important as it limits the potential theoretical and practical applicability of such signals.

We proposed several reasons why EEG β-bursts may only weakly index stopping. First, the rarity of β-bursts may result from the low signal-to-noise ratio (SNR) of EEG recordings. We addressed this directly by recording directly from the cortex, adjacent to where these signals are recorded from, using microelectrode

arrays. As we record directly from the cortex, we reduce any filtering of neural activity by the skull, muscles, and other tissue, effectively allowing us to increase the SNR. In our intracortical sample, we observed that β-bursts became less common as a saccade was being prepared, consistent with interpretations proposing β-activity acts as a brake on movement or decisions (Muralidharan et al., 2022). Although we observed that β-bursts were less common prior to a longer-latency saccade being generated, we note that this may be resultant of contamination of a marked visual transient β-burst response at earlier response latencies. Importantly, although β-bursts were more prevalent during the period in which a saccade was inhibited, they were infrequent, and their properties did not account for the likelihood of canceling a planned response. As such, despite improving the SNR, we did not see any increase in the incidence of β-bursts or any improvement in their relation to stopping behaviors.

This finding demonstrates that β-bursts within the Supplementary Eye Field cannot act as a causal mechanism for reactive response inhibition. Instead, we consider that β-bursts may best reflect proactive executive control processes. First, we observed an increased likelihood of β-bursts following errors. Although these bursts were not associated with post-error adjustments in behavior, these findings mirror previous observations of error-related activity in the medial frontal cortex (Emeric et al., 2010; Purcell et al., 2012b; Sajad et al., 2019; Stuphorn et al., 2000). Interestingly, β-bursts were more frequent well after the error, just before the feedback tone. Second, we observed a prolonged, higher incidence of β-bursts after successful stopping. After successfully inhibiting a plan to generate a movement, our task required monkeys to maintain fixation until a feedback tone sounded. Hence, β-bursts during this post-stopping period could be interpreted in the context of previous findings implicating β-activity in "clearing out" working memory (Lundqvist et al., 2016) and blocking distractions (Castiglione et al., 2019; Hanslmayr et al., 2014). These processes may allow the maintenance of task goals (e.g., sustain unblinking fixation) that will result in reward (Sajad et al., 2022). Collectively, these observations both provide evidence that SEF signals are most prominent following an action, supporting previous theories of medial frontal cortex function (Bonini et al., 2014; Passingham et al., 2010; Schall et al., 2002).

Interestingly, these observations also highlight a clear distinction between human and macaque models of stopping, with particular regard to the role of the medial frontal cortex. Whilst human studies emphasize that the pre-SMA is active during response inhibition and can contribute to the interruption of response preparation (Diesburg and Wessel, 2021; Rae et al., 2015; Sharp et al., 2010), neurophysiological evidence obtained in the medial frontal cortex of macaques demonstrate that neurons in this area modulate too late to directly contribute to stopping (Scangos and Stuphorn, 2010; Stuphorn et al., 2010). In our introduction, we proposed that this divergent evidence between macaques and humans may be resultant of differences intrinsic to the methodologies and analyses used.

We attempted to address this by mirroring recording and analysis techniques used in humans and combining them with intracortical recordings afforded by macaque neurophysiology. Through our previous work, we demonstrated that β-bursts in EEG had functionally similar properties to those in humans; their incidence decreased as a movement was being prepared and was slightly more common during response inhibition (Errington et al., 2020). Interestingly and unexpectedly, however, although we mirror the functional properties of EEG β-bursts in our intracortical signals, there was no association between the presence of a burst within the cortex and one within EEG. This is an important null result, the confidence in which is engendered by other clear, specific, and positive associations observed between neural signals in SEF with overlying EEG (Sajad et al., 2019; Herrera et al., 2020) and other examples of dissociated SEF LFP vs. EEG during performance monitoring (Westerberg et al., 2020). However, further investigation is needed to determine the extent to which EEG β-burst production sampled over other cortical areas is related to LFP β-bursts in the underlying dipole sources. We also, however, note one important confound: human studies classically employ a manual countermanding task, in which participants move a joystick or press a button, whilst macaque studies (including ours) employ a saccade countermanding task. Although pre-SMA and SEF have similar functions for their respective effectors (Fujii and Graybiel, 2003), and models of response inhibition are applicable across different modalities (Boucher et al., 2007; Logan and Irwin, 2000), there may be significant anatomical and functional differences in the stopping circuitry between these two which lead to variations in the contribution of different brain areas to the EEG signal. This warrants further investigation.

4.4.2 Origin of β-burst in agranular neocortex and in EEG

Bioelectric potentials have practical and clinical applications when their generators are known. For example, the electrocardiogram is useful in medicine because the physiological process associated with each phase of polarization is understood. Likewise, the electroretinogram is useful because the cell layers associated with each polarization are understood. In contrast, signals recorded over the scalp indexing cognitive operations will have limited utility until their neural generators are known. This work adds to a growing literature addressing this issue by simultaneously recording signals in the brain and on the cranial surface (Bimbi et al., 2018; Cosman et al., 2018; Heitz et al., 2010; Herrera et al., 2022; Purcell et al., 2013; Sandhaeger et al., 2019; Shin et al., 2017; Westerberg et al., 2020, 2022; Whittingstall and Logothetis, 2009).

By combining the laminar profile of β-bursts described in this study with knowledge of the anatomical and histological properties of SEF, we can gain further insight into how β-bursts can arise and how they relate to β-bursts recorded on the cranial surface. A recent model using 100 multi-compartment pyramidal cells in layer 3, another 100 in layer 5 with 35 inhibitory neurons in layer 3 and layer 5 showed how β-bursts can be produced via a transient dipole reversal caused by the interaction between a prolonged ($>$100 ms)

excitatory synaptic drive in the basal dendrites of pyramidal neurons that interacts with a transient pulse of excitatory drive in the apical dendrites persisting for at least 50 ms (Sherman et al., 2016; Shin et al., 2017). Consistent with this, we found β-bursts were more prevalent in upper cortical layers across all task epochs. The core mechanism embodied by the model of synchronous patterns of basal and apical dendritic events has been proposed in other contexts such as error detection (Cohen, 2014), and another recent biophysical model of layer 5 pyramidal neurons highlighted the contributions of apical dendrite calcium spikes (Herrera et al., 2020). The interpretation of the origin and role of β-bursts is enhanced by converging observations about the properties of neurons across the layers of SEF (Sajad et al., 2019), which demonstrate contributions to error detection, reinforcement expectation, response conflict, event timing, and working memory.

Interestingly, we found that LFP β-burst presence was not synchronized across the cortical column but instead coincided only on adjacent LFP channels. The spacing of the linear electrode contacts (150 μm) indicates that, assuming an isotropic medium, intracranial β-bursts occupy a volume no more than $\sim$0.014 mm3. These observations offer useful constraints on biophysical models of β-burst production (Sherman et al., 2016).

4.4.3 Cautionary considerations of β-burst activity

Sampling field potentials from within the cerebral cortex was supposed to improve the poor signal-to-noise ratio (SNR) of noninvasive EEG recordings. Nevertheless, we observed β-bursts as rarely in intracortical LFP as in cranial EEG recordings. Combining simultaneously recorded EEG and LFP signals allowed a direct comparison of β-bursts in two techniques thought to differ in SNR. Interestingly, at the individual channel level, we observed that β-bursts were as rare in intracranial LFPs as they were in EEG signals. Despite being similar in incidence, LFP β-bursts were not coincident with β-bursts observed in EEG. Furthermore, the incidence of β-bursts observed within MFC parallels that observed in human intracranial recordings (Diesburg et al., 2021; Mosher et al., 2021; Yu et al., 2021).

These results urge caution when interpreting β-bursts across cortical layers. By definition, β-bursts are detected when β-band power crosses a pre-defined threshold, based on the average overall power at a site. This approach may result in a lesser probability for a burst to be detected at sites with greater β-power. In this study we took two approaches to this potential problem, considering the threshold as a function of either the power at the given contact or the power across the entire site (all contacts within the cortex in the penetration). We believe taking the power across the entire site results in more balanced and fair opportunities for a burst to be captured and have presented data from that approach in this manuscript. Nevertheless, both approaches yielded qualitatively similar patterns that allow for the preposition of ideas on the functional architecture of β-burst signaling in the medial frontal cortex.

4.4.4 Conclusions

These results demonstrate that β-bursts are about as rare in the cerebral cortex as they are on the cranial surface and exhibit no more reliable association with response inhibition than that previously reported for EEG β-bursts. The findings demonstrate that β-bursts are incapable of mediating any mechanism of response inhibition. Nevertheless, β-bursts in MFC may index other processes underlying executive control.

CHAPTER 5

Functional Architecture of Executive Control and Associated Event-Related Potentials in Supplementary Eye Field

5.1 Introduction

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) task (Verbruggen et al., 2019b), during which macaque monkeys, like humans, exert response inhibition and adapt performance based on stimulus history, response outcomes, and the temporal structure of task events (Emeric et al., 2007).

Converging evidence from imaging, electrophysiology, and lesion studies indicates that MFC, including the supplementary motor complex, is essential for executive control (Kolling et al., 2016; Shenhav et al., 2016b; Cohen, 2014; Fu et al., 2022a). 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 function (Kok et al., 2004). Executive control of gaze is mediated by the SEF – an agranular area on the dorsomedial convexity of the medial frontal cortex. Here, neurons modulate too late to enable reactive control of eye movements and instead contribute to proactive control (Stuphorn et al., 2010). Electrical microstimulation of SEF improves performance in the countermanding task by slowing response time (RT) (Stuphorn and Schall, 2006) through postponing the accumulation of pre-saccadic activity in gaze control structures (Pouget et al., 2011). SEF also supports working memory (Bastos et al., 2018; Reinhart et al., 2012a) and signals surprising events (Kawaguchi et al., 2015), event timing (Ohmae et al., 2008; Wang et al., 2018), response conflict (Stuphorn et al., 2000), and consequences (Sajad et al., 2019).

Whilst these executive control signals are well documented in MFC, how they arise is uncertain (Cohen, 2014). 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 system (Hubel and Wiesel, 1968) the canonical cortical microcircuit derived from granular sensory areas (Bastos et al., 2012) does not explain agranular frontal areas like SEF (Beul and Hilgetag, 2014; Godlove et al., 2014; Ninomiya et al., 2015; Rapan et al., 2021; Shipp, 2005).

To understand mechanistically how executive control signals are generated in the medial frontal cortex, we sample 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 the error-related negativity (ERN) (Sajad et al., 2019). 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.

To understand mechanistically how executive control signals are generated in the medial frontal cortex, 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 the error-related negativity (ERN). 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 the first 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.

5.2 Methodology

5.2.1 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 (Godlove et al., 2014; Hanes and Schall, 1995). 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 previously (Godlove et al., 2011). 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 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.

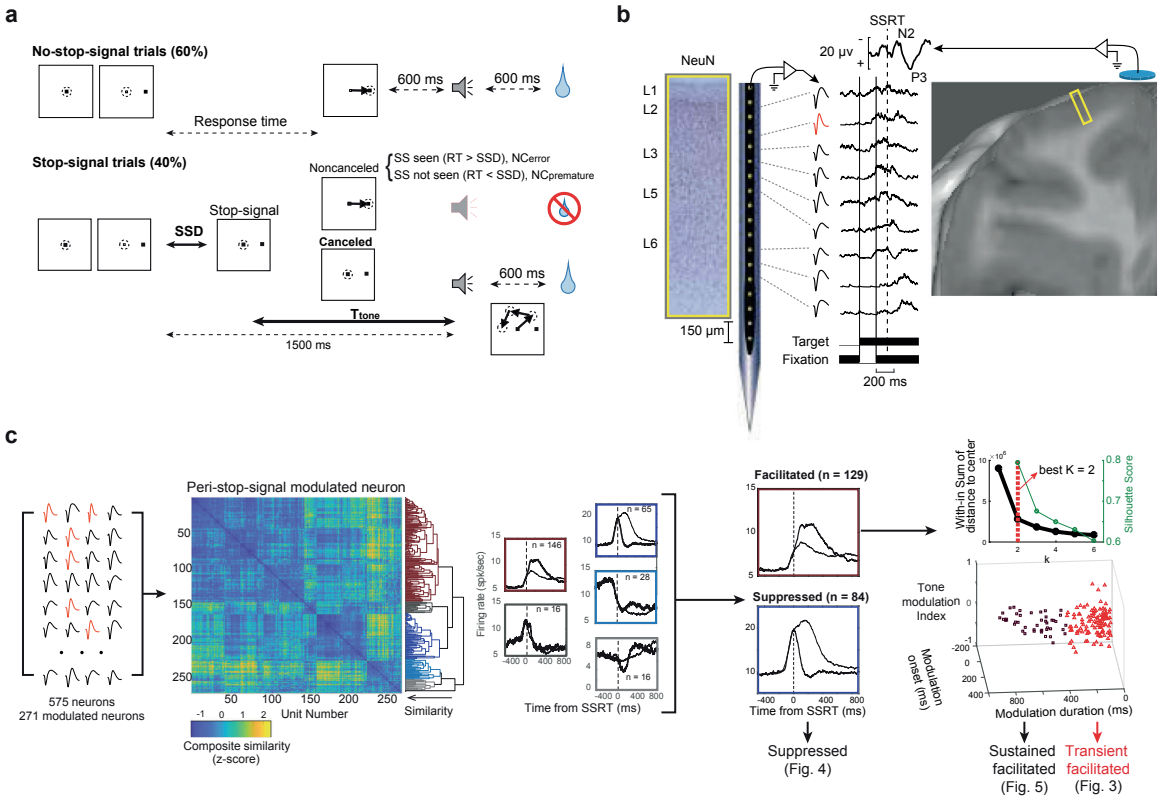

Figure 5.1: **Experimental Approach**. **A**. Saccade countermanding task. Monkeys earned fluid reward for shifting gaze (dashed circle) to a visual target unless a stop-signal appeared after a variable stop-signal delay (SSD) adjusted to achieve ∼50% canceled trials. Successful no-stop-signal or canceled 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 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 medial frontal cortex (10-20 location Fz). Yellow rectangle portrays 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 pipeline 27 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 was chosen based on the Elbow method and the Silhouette score. Therefore, three classes of neurons were analyzed in this study.

5.2.2 Countermanding task

The countermanding (stop-signal) task utilized in this study has been widely used previously (Verbruggen et al., 2019b). Briefly, trials were initiated when monkeys fixated at a central point. Following a variable time period, drawn from an aging function to avoid anticipation of the visual stimulus (Errington and Schall, 2020), 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 gaze to the target. Fixation on the target was required for 600 ms, until an auditory tone sounded, whereupon monkeys could shift gaze anywhere. 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 $\sim$50 ms for Monkey Eu (45 ± 15 ms) and by $\sim$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 stop-signal task. Different SSD values were used for the two subjects to account for between-subject differences in stopping performance (Verbruggen et al., 2019b). 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. Fluid reward was delivered 600 ms later. By design, the duration from target presentation until the tone was a fixed interval of 1500 ms. Thus, as SSD increased, the duration of fixation decreased (Figure 5.1a).

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 process (Logan and Cowan, 1984). The race model provides an estimate of the duration of the covert STOP process, the time taken to accomplish response inhibition, known as stop-signal reaction time (SSRT) (Boucher et al., 2007; Lo et al., 2009; Logan et al., 2015). 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; Matzke et al., 2017; Appendix C, Supplementary Figure 4a, 5a, 6a). Compared to weighted integration method, the Bayesian approach provides estimates of the variability in SSRT and the fraction of trigger failures for a given session (Matzke

et al., 2017b). Individual parameters were estimated for each session. Based on previous implementations of this approach, 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 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.

5.2.3 Data Acquisition

5.2.3.1 Cortical mapping and electrode placement

Chambers implanted over the medial frontal cortex were mapped using tungsten microelectrodes (2-4 M/*Omega*, FHC, Bowdoin, ME) to apply 200ms trains of biphasic micro-stimulation (333 Hz, 200 μs pulse width). The SEF was identified as the area from which saccades could be elicited using < 50 μA of current 84, 85. 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, 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 5 mm lateral to the midline and 31 mm anterior to the interaural line. In monkey X, the perpendicular penetrations sampled activity at site P2 and P3, located 5 mm lateral to the midline and 29 and 30 mm anterior to the interaural line, respectively. However, during the mapping of the bank of the cortical medial wall, we noted both monkeys had chambers place $\sim$1 mm to the right respective to the midline of the brain. This was confirmed through co-registered CT/MRI data. Subsequently, the stereotaxic estimate placed the electrodes at 4 mm lateral to the cortical midline opposed to the skull-based stereotaxic midline.

5.2.3.2 Acquiring EEG

EEG was recorded from the cranial surface with electrodes located over medial frontal cortex. 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 200ms before or after the analysis interval were removed.

5.2.3.3 Acquiring neural spiking

Spiking activity and local field potentials were recorded using a 24-channel Plexon U-probe with 150 μm between contacts, allowing sampling from all layers. The U-probes were 100 mm in length with 30 mm reinforced tubing, 210 μm probe diameter, 30° tip angle, with 500 μm 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 dura and entered 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 μm 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 μs 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.

5.2.3.4 Cortical depth and layer assignment

The retrospective depth of the electrode array relative to grey 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 pulsated 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 layers (Maier et al., 2010). Thirdly, an automated depth alignment procedure was employed which maximized the similarity of CSD profiles evoked by passive visual stimulation between sessions (Godlove et al., 2014).

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 pro-

teins parvalbumin (PV), calbindin (CB), and calretinin (CR) (Godlove et al., 2014). Additional information about laminar structure was assessed through the pattern of cross-frequency phase-amplitude coupling across SEF layers (Ninomiya et al., 2015). 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.

5.2.3.5 Acquiring eye position

Eye position data was 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.

5.2.4 Statistical Analysis

5.2.4.1 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 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 (ΔERP) was obtained by subtracting from the ERP recorded on canceled trials the ERP recorded on RT-matched no stop-signal trials.

5.2.4.2 Analysis of neural spiking

Spike widths exhibited a bimodal distribution (Godlove et al., 2014; Sajad et al., 2019), and neurons were distinguished as narrow- (peak-to-trough duration less than 250 μs) or broad-spikes ($\geq$ 250 μs). 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:

$$SDF = (1 - \exp(\frac{t}{\tau_g})) \times \exp(\frac{-t}{\tau_d})$$

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 post-synaptic 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. 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 algorithm 27 (Figure 5.1c, Appendix C, Supplementary Figure 1c-e). The input to this analysis pipeline was the spike-density function 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 (ΔSDF), 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 medial frontal cortex, 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 ΔSDF were used.

5.2.4.3 Mixed effects models

We fit variation in spike counts or EEG voltage to models of task events and performance outcomes (Figure 5.1, Appendix C, Supplementary Table 2, Supplementary Figure 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 ∼350ms were not included because too few canceled trials were obtained. The analysis of the facilitation after SSRT was based on ΔSDF (Figure 5.2, Figure 5.4), but the major conclusions held if the analysis used SDF. The analysis of the modulation before SSRT or the feedback tone (Figure 5.3) 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 3) 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 (Appendix C, Supplementary Figure 3b).

Mixed-effects models of ΔSDF, SDF, or ΔERP values in relation to the various performance measures were compared using Bayesian Information Criteria (BIC). We report the results of the most basic version of each model with a main effect term corresponding to the performance parameter and random intercepts grouped by neuron (for spiking) or session (for EEG). The values for each performance parameter were z-transform normalized for 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. ΔBIC (BICbest – BICcompeting) quantified the fit of the other models relative to the best with $\Delta BIC < 2$ offering weak support against the competing model, $2 < \Delta BIC < 6$ offering strong support against the competing model, and $\Delta BIC > 6$ conclusively ruling out the competing model 88, 89, More complex versions of these models resulted in similar conclusions. Mixed-effects models were performed using MATLAB's Statistical Toolbox.

5.2.4.4 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. A multiple linear regression model controlled for SSD to test whether spike rates varied with RT.

5.2.4.5 Relating N2/P3 and neural spiking

We used the method described previously to establish the relationship between spiking activity and the ERN 16. 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 Figure 5.5 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. Three factors were considered: (1) spiking activity in L2/3, (2) spiking activity in L5/6, plus (3) SSD to prevent its variation from confounding the relationship between ERP and neural spiking. However, as presented in the main text, the inclusion of this factor did not change the results.

5.3 Results

5.3.1 Countermanding performance and neural sampling

Neurophysiological and electrophysiological data was recorded from two macaque monkeys performing the saccade stop-signal task with explicit feedback tone cues preceding possible delivery of fluid reward (Figure 5.1) (Hanes and Schall, 1995). Data collection and analysis was informed by the consensus guide for the stop-signal task (Verbruggen et al., 2019b). Briefly, monkeys earned fluid reward for shifting 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 (Figure 5.1b): NC trials in which 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: 230 ± 83 ms) were systematically shorter than those on no-stop-signal-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 model (Logan and Cowan, 1984) 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, gaze will shift. Accordingly, neural modulation occurring before SSRT can contribute to reactive stopping but modulation occurring after SSRT cannot (Hanes et al., 1998; Stuphorn et al., 2010). 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 medial frontal cortex 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 (Figure 5.1a). SEF was localized by anatomical landmarks and intracortical electrical microstimulation (Godlove et al., 2014). 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 (Appendix C, Supplementary Figure 1b, Appendix C, 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.

5.3.2 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 model (Logan and Cowan, 1984) and validated neurophysiologically (Hanes et al., 1998) noncanceled trials result from faster responses than canceled trials (RT $<$SSD + SSRT, Appendix C, Supplementary Figure 2a) and therefore are associated with distinct neural mechanisms. Activity on noncanceled stop-signal trials were described previously (Sajad et al., 2019).

We classified neurons primarily based on the pattern of neural activation around SSRT (Figure 5.1c, Appendix C, Supplementary Figure 1c). A semi-supervised consensus cluster algorithm (Lowe and Schall, 2018) 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 (Figure 5.1c).

To identify 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

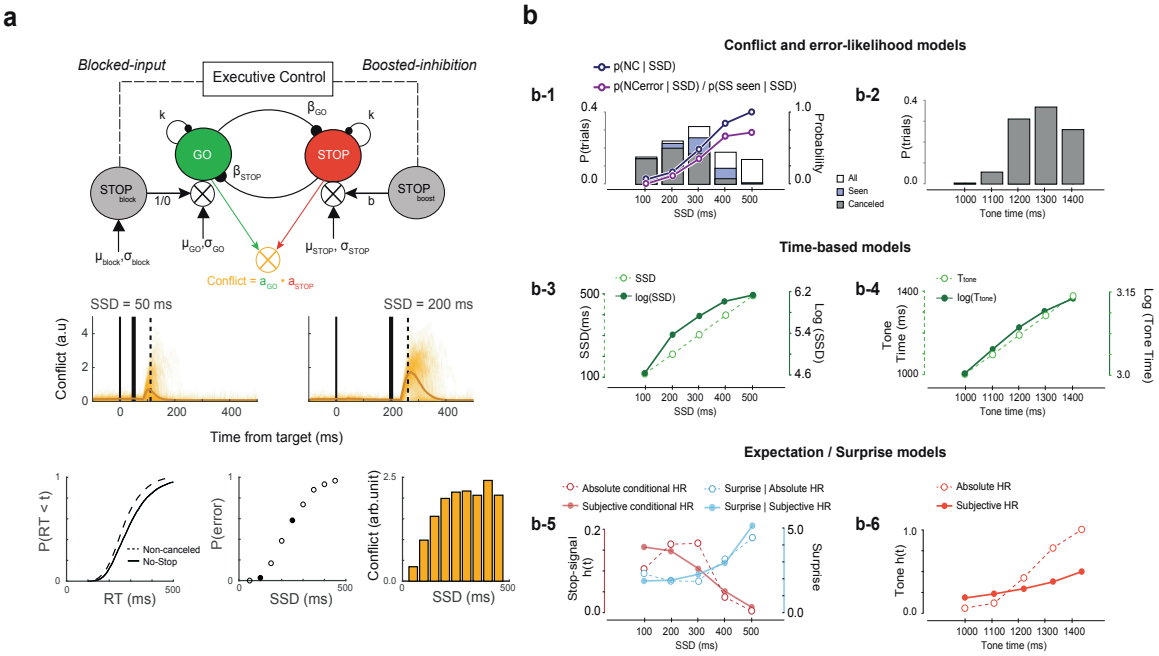

Figure 5.2: **Models of neural modulation.** **A**. Interactive race model architecture. In interactive race architectures 29, 30 GO and STOP units are racing, competing response channels, so the product of their co-activation measures response conflict in this task. In simulations with parameters simulating observed RT and p(NC | SSD) values (lower panels) the instantaneous product of aGO*aSTOP 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 countermanding performance is proportional to p(NC | SSD) (lower right panel). Consequently, p(NC | 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 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(NC | SSD) and error-likelihood with $p(NC_{error} \mid SSD) / p(SS_{seen} \mid SSD)$, 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 (which incorporates imprecision in duration in duration estimation) hazard rates (red) of SSD (b-5) and tone (b-6) and their associated surprise (blue) are shown. Other variants were also considered based on the underlying assumption about knowledge of task structure and imprecision in duration estimation (Appendix C, Supplementary Table 2, Appendix C, Supplementary Figure 3a). Expectation of SS was quantified by hazard rate conditional on stop-signal trial probability ($\sim$40%). Surprise as a violation of expectations was quantified by the Shannon information derived from the hazard rate.

enough to have been canceled had the stop-signal been presented (RT

$$\geq$$

SSD + SSRT; Appendix C, Supplementary Figure 2a). We distinguished canceled trials by SSD to account for variations in modulation dynamics arising from differences in the timing of task events (Figure 5.1a). We also examined the neural activity before the feedback tone which terminated operant control on behavior after successful stopping (Figure 5.1a). Variations in neural activity across SSDs were tested against a variety of parameters derived from different theories of MFC function described below (Figure 5.2b; Appendix C, Supplementary Table 2).

Performance of the stop-signal task is explained as the outcome of a race between stochastic GO and STOP processes (Logan and Cowan, 1984), instantiated by specific interactions enabling the interruption of the GO process by a STOP process (Boucher et al., 2007; Logan et al., 2015) (Figure 5.2a). A theory of medial frontal function posits that it encodes the conflict between mutually incompatible processes (Botvinick et al., 2001). 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(NC | 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 (Figure 5.2a). Therefore, the conflict model predicts variations in neural activity as a function of p(NC | 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 SSD (Brown and Braver, 2005). 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 $p(NC_{error} \mid SSD) / p(SS_{seen} \mid SSD)$. Neural modulation scaling positively (negatively) with this quantity encode the error (success) likelihood associated with each SSD. This quantity diverges from p(NC | SSD) at longer SSDs (Figure 5.2b, top-left panel).

Finally, due to the temporal regularities of SSD and other events, we determined whether neural modulation signaled interval timing (Kim et al., 2013; Ohmae et al., 2008; Tallot and Doyere, 2020; Wang et al., 2018), moment-by-moment expectation for events (Coull and Nobre, 2008; Janssen and Shadlen, 2005), or the surprise associated with the violation of the expectations (O'Reilly et al., 2013; Starkweather et al., 2017) (Figure 5.2b, Appendix C, Supplementary Table 2, Appendix C, Supplementary Figure 3). Guided by previous research on time perception (Gibbon, 1977; Janssen and Shadlen, 2005; Kim et al., 2013), we

tested whether neural modulation varied with elapsed time in linear or logarithmic scales. Expectation was operationalized by hazard rate,

$$h(t) = (f(t))/([1 - F(t)])$$

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,

$$s(t) = [-log]_2[h(t)]$$

Different hazard rate 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}; Figure 5.2b; Appendix C, Supplementary Figure 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).

5.3.2.1 Monitoring Conflict

The classification pipeline identified 75 neurons with transient facilitation on canceled trials compared to latency-matched no-stop-signal trials (Figure 5.1c). This facilitation was not a visual response to the stop-signal because it did not occur on noncanceled trials (Appendix C, Supplementary Figure 1e; Appendix C, Supplementary Figure 2b). It also cannot contribute to reactive response inhibition because for nearly all neurons (71/75) it arose after SSRT (Figure 5.3a, Appendix C, Supplementary Figure 4a). The facilitation started 90.4 ± 74 ms after SSRT with peak recruitment at ∼110 ms at which time ∼60% of neurons were active (Figure 5.3a).

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(NC | SSD) (Figure 5.3e; Appendix C, 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 ($\Delta BIC < 2$). The error-likelihood ($\Delta BIC = 3.52$) and other surprise models earned weak support ($\Delta BIC = 5.7$) or were rejected ($\Delta BIC > 6$) (Figure 5.3d; Appendix C, 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 activation (Stuphorn et al., 2000). 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 (Figure 5.3a, b; Appendix C, Supplementary Figure 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 previously (Sajad et al., 2019)(Appendix C, Supplementary Table 7). Some produced higher discharge rates on unrewarded trials (previously identified as Loss neurons); some, 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), χ^2 (3, N = 575) = 1.02, p = 0.791). The vast majority (65/75) were not modulated when stopping failed (Appendix C, Supplementary Figure 2b, Appendix C, Supplementary Table 7), reinforcing previous findings that conflict and error monitoring are distinct (Sajad et al., 2019; Stuphorn et al., 2000).

To investigate the microcircuit contribution of Conflict neurons, we examined their spike waveform duration and distribution across the layers (Appendix C, Supplementary Figure 1a). Neurons were distinguished as broad- and narrow-spiking, which may identify pyramidal neurons that can send extrinsic connections and intrinsic inhibitory interneurons, respectively (Lemon et al., 2021). 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), χ^2 (1, N = 575) = 0.2.05, p = 0.152).

Conflict neurons were observed in all recording sites but more prevalently at some (Chi-square test of homogeneity (one-tailed), χ^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 (Figure 5.3b). These neurons were sampled from all layers with an incidence corresponding to the sampling distribution (Chi-square test of homogeneity (one-tailed), χ^2 (4, N = 293) = 4.28, p = 0.369; Figure 5.3b; Appendix C, Supplementary Figure 1a; Appendix C, 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 (Figure 5.3b). The few neurons modulating with the tone were observed in all layers.

Thus, complementing previous observations (Stuphorn et al., 2000), 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 (Botvinick et al., 2001; Schall and Boucher, 2007). These neurons

are distributed across all SEF layers and are predominantly broad-spiking, as expected from the sampling distribution. This conflict signal has been observed in dopamine (DA) neurons in dorsolateral substantia nigra pars compacta (SNpc) during saccade countermanding (Ogasawara et al., 2018). The SNpc neurons modulate significantly earlier than those in SEF, but accounting for the long conduction delay, a dopaminergic signal cannot cause the SEF modulation (Appendix C, Supplementary Figure 9; Appendix C, Supplementary Table 9).

5.3.2.2 Event Timing

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 (Figure 5.4; Appendix C, Supplementary Figure 1c-e). This modulation cannot contribute to reactive response inhibition because for nearly all neurons (76/84) it happened after SSRT (Appendix C, Supplementary Figure 1d; Appendix C, Supplementary Figure 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 (Figure 5.4a). Noting the similarity of the ramping to earlier descriptions of time keeping neurons (Ohmae et al., 2008; Wang et al., 2018), we focused on the ramping activity preceding the suppression.

Model comparison revealed that the ramping magnitude varied best with log(SSD) (Mixed-effects linear regression (two-tailed) grouped by neuron, t(250) = 12.62, p = 0.001) with higher activity for longer SSD durations (Figure 5.4e; Appendix C, Supplementary Table 4). Absolute SSD earned weak support (ΔBIC = 2.7), but hazard rate models were rejected (ΔBIC $>$ 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 (Appendix C, Supplementary Figure 2b). Because the discharge rate decreased sharply after SSRT on canceled but not on noncanceled stop-

[1]**Figure 5.3. 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 difference between trials indicated by p(diff $>$ 0). Tone-aligned recruitment was difference in spiking on canceled trials (Acanceled) relative to baseline (BL) lowest spiking value ±500 ms from 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 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(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 difference in spiking across conditions. **D**. Model comparison table listing each tested model. The heatmap shows 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. Green circle indicates the best-fit model; white indicate candidate models (ΔBIC $<$2). Spike rate variation after SSRT was best predicted by the conflict model. Full statistics in Appendix C, Supplementary Table 3. **E**. Significant variation of spiking (residualized and adjusted for spiking across neurons) as a function of p(NC | SSD) proxy for conflict (normalized z-scale). Each point plots average spike-density and mean p(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 late-SSD bin was obtained due to too few trials. These data points were not included.

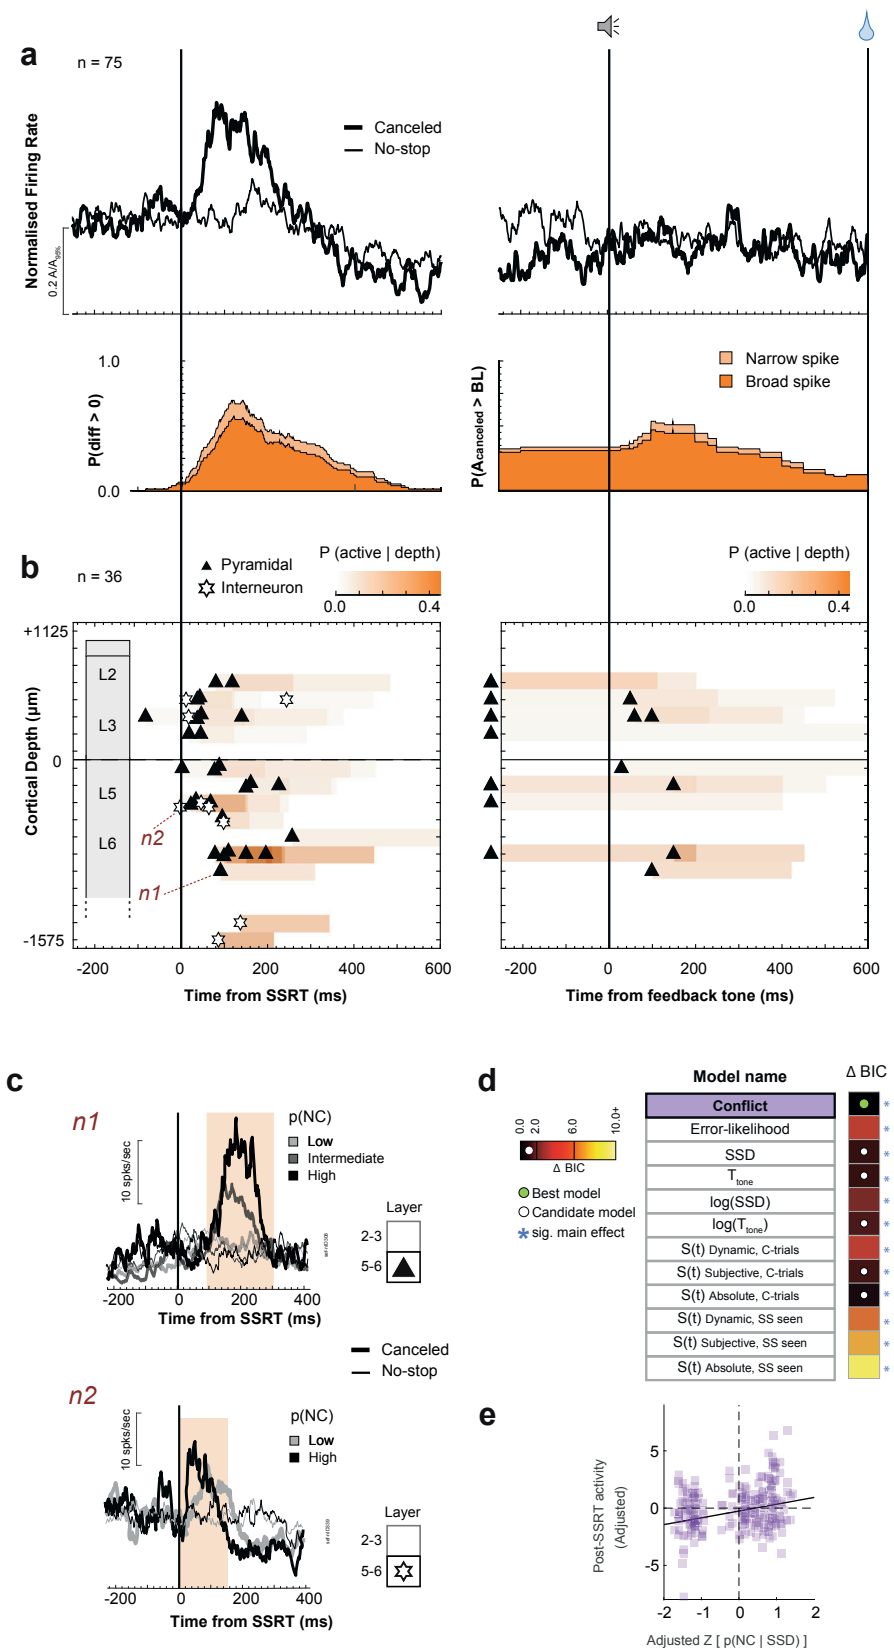

Figure 5.3: **Conflict neuron time-depth organization.** See footnote [1]

signal trials (Figure 5.4d, Appendix C, Supplementary Figure 1e; Appendix C, Supplementary Figure 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}, Figure 5.1a). 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 (Figure 5.4; Appendix C, Supplementary Figure 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 (Figure 5.4a).

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, $t(112) = 3.41$, $p < 0.001$; Figure 5.4h, Appendix C, Supplementary Table 4) with weak support for hazard rate models ($3.0 < \Delta BIC < 4.4$). The T_{tone} and $\log(T_{tone})$ models were indistinguishable ($\Delta BIC < 0.1$) (Figure 5.4g). The termination of this ramping activity was synchronized with the feedback tone and not the time at which fixation was interrupted following the feedback (Appendix C, Supplementary Figure 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 specificity in different neurons.

Next, we examined how the activity of these neurons relate to adjustments in behavior. Replicating previous findings, we found that RT was influenced by the SSD experienced on the previous trial, with slower RTs following longer SSD (Linear regression model (two-tailed), $t(83) = 2.64$, $p = 0.010$) Nelson et al. (2010). However, the peak of the peri-SSRT ramping activity of only 2 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 previously (Sajad et al., 2019). 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), χ^2 (3, N = 575) = 44.86, $p < 0.001$, Appendix C, Supplementary Table 7). The nature of the ramping before SSRT was invariant across multiplexing association. 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), χ^2 (1, N = 575) = 3.75, p = 0.053). These neurons were found in all penetrations but were more common in some sites (χ^2 (4, N = 575) $>$ 39.3, p $<$0.001; Appendix C, Supplementary Table 1). The time-depth organization of these neurons was revealed in the sample from perpendicular penetrations (49/84; Figure 5.4b). Those ramping before SSRT were found across all layers in proportion to the sampling distribution (Chi-square test of homogeneity (one-tailed), χ^2 (4, N = 293) = 7.33, p = 0.120; Figure 5.4b; Appendix C, Supplementary Figure 1a; Appendix C, Supplementary Table 1). However, those with ramping after SSRT until the tone were significantly more concentrated in lower L3 and L5 (χ^2 (2, N = 293) = 10.37, p = 0.006; Appendix C, 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 specificities (Zhang et al., 2019) and multiplex with error and reinforcement signals in different layers.

5.3.2.3 Goal Maintenance

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 (Figure 5.5). This neuron class was distinguished from Conflict neurons based on the duration of facilitation (K-means clustering, Figure 5.1c, Appendix C, Supplementary Figure 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 (Appendix C, Supplementary Figure 2b). It cannot contribute to reactive response inhibition because it arose after SSRT for effectively all neurons (53/54; Figure 5.5a; Appendix C, Supplementary Figure 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

[2]**Figure 5.4.** Conventions as in Figure 5.3. **A**. Average spike rate and suppression time of neurons with ramping spiking after the target (mean: vertical line; mix-max: gray-shaded rectangle) followed by suppression on successfully canceled relative to latency-matched no stop-signal trials. SSRT-aligned recruitment was 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 were used for analysis. **D**. Model comparison table for pre-SSRT activity. Best model is highlighted green. Variation of spiking was best predicted by log(SSD). Full statistics in Appendix C, Supplementary Table 4. **E**. Significant variation of spiking activity before SSRT as a function of log(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_{tone}). Full statistics in Appendix C, Supplementary Table 4. **H**. Significant variation of spiking activity before tone as a function of log(T_{tone}) with 38 neurons contributing 144 samples across early, intermediate, or late T_{tone}.

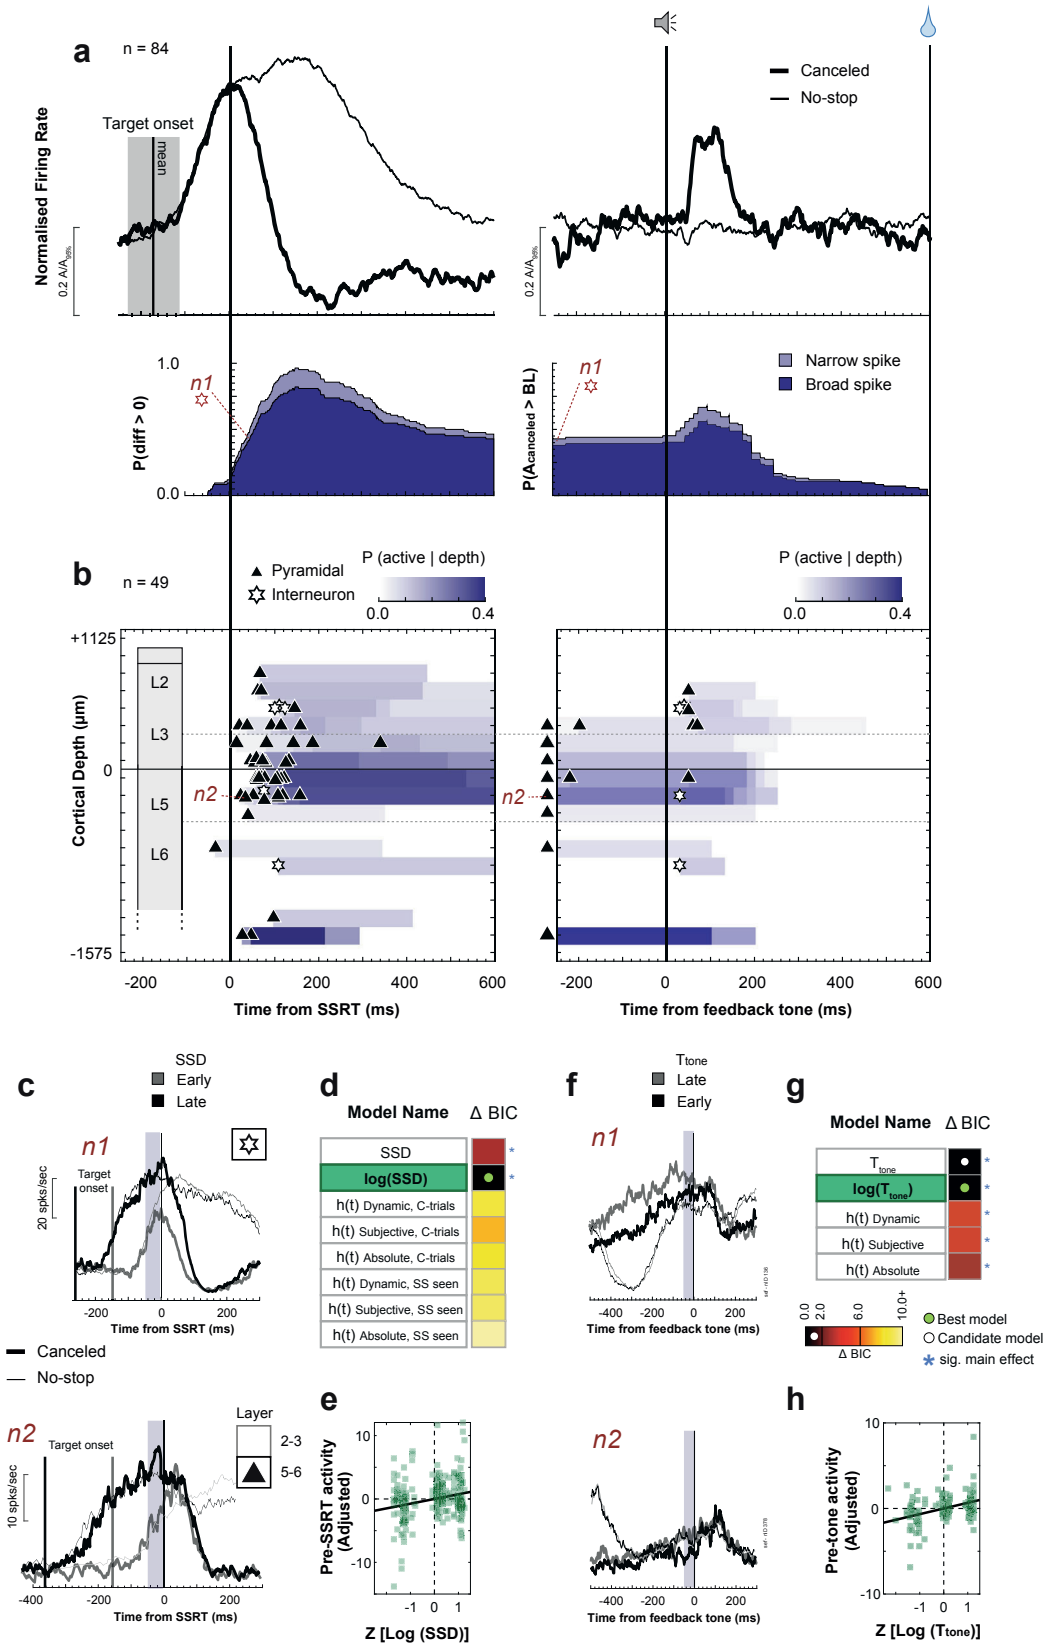

Figure 5.4: **Event timing neuron time-depth organization**. See footnote [2]

$<$0.001). The peak recruitment of these neurons reached $\sim$95% after $\sim$300 ms, later than that of the Conflict ($\sim$110 ms) and Event Timing neurons ($\sim$150 ms) (Figure 5.5a).

The variation in activity of these neurons was explained best by the surprise model (SSSubjective, SS seen; Mixed-effects linear regression (two-tailed) grouped by neuron, t(151) = -3.91, p $<$0.001 ; Figure 5.5d; Appendix C, Supplementary Table 5) with other surprise and time-based models also candidates (ΔBIC $<$2) and weak support for log(SSD) (ΔBIC = 3.0) and conflict (ΔBIC = 4.1) models. The error-likelihood model was rejected (ΔBIC $>$ 6.0). The variation of activity varied inversely with surprise and positively with T_{tone} and log(T_{tone}) (Figure 5.5c,e; Appendix C, Supplementary Table 5). Thus, besides having persistent activity, the modulation of these neurons was further distinguished from the conflict signal by the different (and opposite) relationship to performance and task parameters (Appendix C, Supplementary Figure 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), χ^2 (1, N = 129) = 27.3, p $<$0.001). The sustained activity attenuated $\sim$300 ms after the feedback tone that cued successful performance (Figure 5.5a). 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 (Appendix C, Supplementary Figure 6c).

Based on the evidence above and previous findings identifying SEF signals with working memory (Bastos et al., 2018; Reinhart et al., 2012a), we conjecture that these neurons contribute to maintaining a representation of task goals (e.g., sustain unblinking fixation) for 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; Appendix C, Supplementary Figure 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 signals (Sajad et al., 2019). 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 increased (Sajad et al., 2019) (Appendix C, Supplementary Figure 6d). The prevalence of this multiplexing pattern exceeded chance (Chi-square test of homogeneity (one-tailed), χ^2 (3, N = 575) = 19.43, p $<$0.001; Appendix C, 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), χ^2 (1, N = 575) = 9.27, p = 0.002). They were found in all penetrations but significantly more commonly at certain sites (χ^2 (4, N = 575) $>$ 39.3, p $<$0.001, Appendix C, Supplementary Table 1). Perpendicular penetrations revealed the time-depth organization of 34 Goal Maintenance neurons (Figure 5.5b). The distribution of these neurons across cortical layers was significantly different from the sampling distribution (χ^2 (4, N = 293) = 11.24, p = 0.024, Appendix C, Supplementary Figure 1a, Appendix C, Supplementary Table 1) with significantly more in L2/3 relative to L5/6 (Figure 5.5b, χ^2 (1, N = 293) = 10.37, p = 0.001). Their laminar distribution was also significantly different from those of Conflict (χ^2 (1, N = 70) = 11.54, p $<$0.001) and Event-Timing neurons (χ^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 (Bastos et al., 2018; Reinhart et al., 2012a), 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.

5.3.3 Functional classification of N2/P3 ERP related to successful stopping

To determine whether macaque monkeys produce ERP components associated with response inhibition homologous to humans (Kok et al., 2004), we sampled EEG from an electrode located over MFC (10-20 Fz) while recording neural spiking in SEF (Figure 5.6a). 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 (Figure 5.6b). 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 (Appendix C, Supplementary

[3]**Figure 5.5.** Conventions as in Figure 5.3. **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 (n2). Spiking of both neurons decreased after the tone. Shaded interval highlights significant difference in spiking across conditions. **D**.Model comparison table for post-SSRT activity. Variation of spiking was best predicted by the surprise model s(t) Subjective, SS seen. Full statistics in Appendix C, Supplementary Table 5. **E**. Significant variation of spiking activity after SSRT as a function of s(t) 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 late-SSD bin was obtained due to too few trials. These data points were not included.

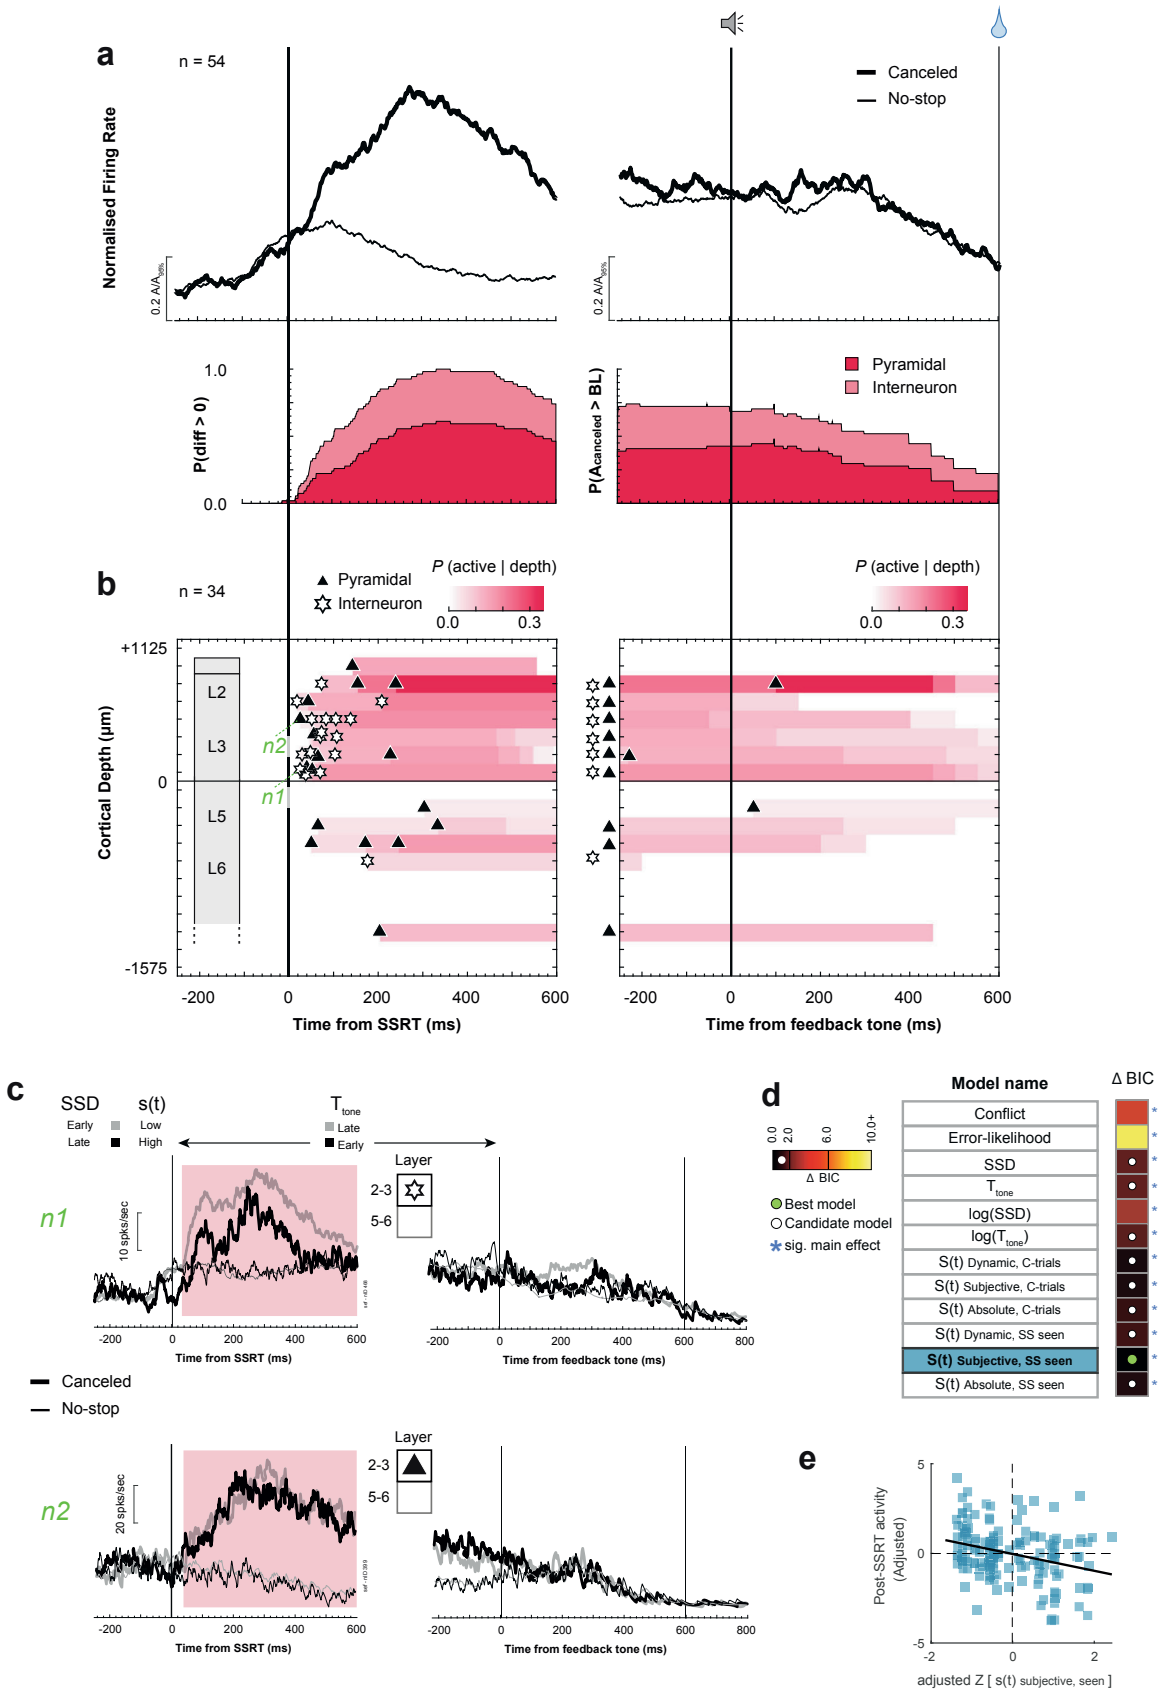

Figure 5.5: **Goal maintenance neuron time-depth organization**. See footnote [3]

Figure 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; Appendix C, Supplementary Figure 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; Figure 5.6d, Appendix C, Supplementary Table 6). Conflict, time-based, and surprise models were rejected (non-significant main effect, and ΔBIC > 3.0; Figure 5.6c). This result adds to the inconclusive evidence for the frontal N2 association with conflict monitoring and response inhibition (Kok et al., 2004).

The N2 was followed by a robust P3 beginning $\sim$300 ms and peaking 358 ± 17 ms after the stop-signal, homologous to the human P3 (Kok et al., 2004)(Figure 5.6a,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; Appendix C, Supplementary Figure 7c). Variation in the amplitude of the P3 was best described by the $\log(T_{tone})$ 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; Figure 5.6c,e; Appendix C, Supplementary Table 6). All time-based models were candidates (ΔBIC $<$1.30). The conflict (ΔBIC = 2.90) and most surprise models (4.5 $<\Delta$BIC $<$5.7) received weak support, and error-likelihood and one surprise model were rejected (ΔBIC > 6.0).

5.3.4 Linking neural spiking to ERP

We examined how neural spiking related to concomitant ERP (Sajad et al., 2019; Westerberg et al., 2022). Appreciating that EEG arises from $\sim$106 neurons and spikes are too brief to create scalp EEG, 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 (Figure 5.6b). The relationship between neural 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;

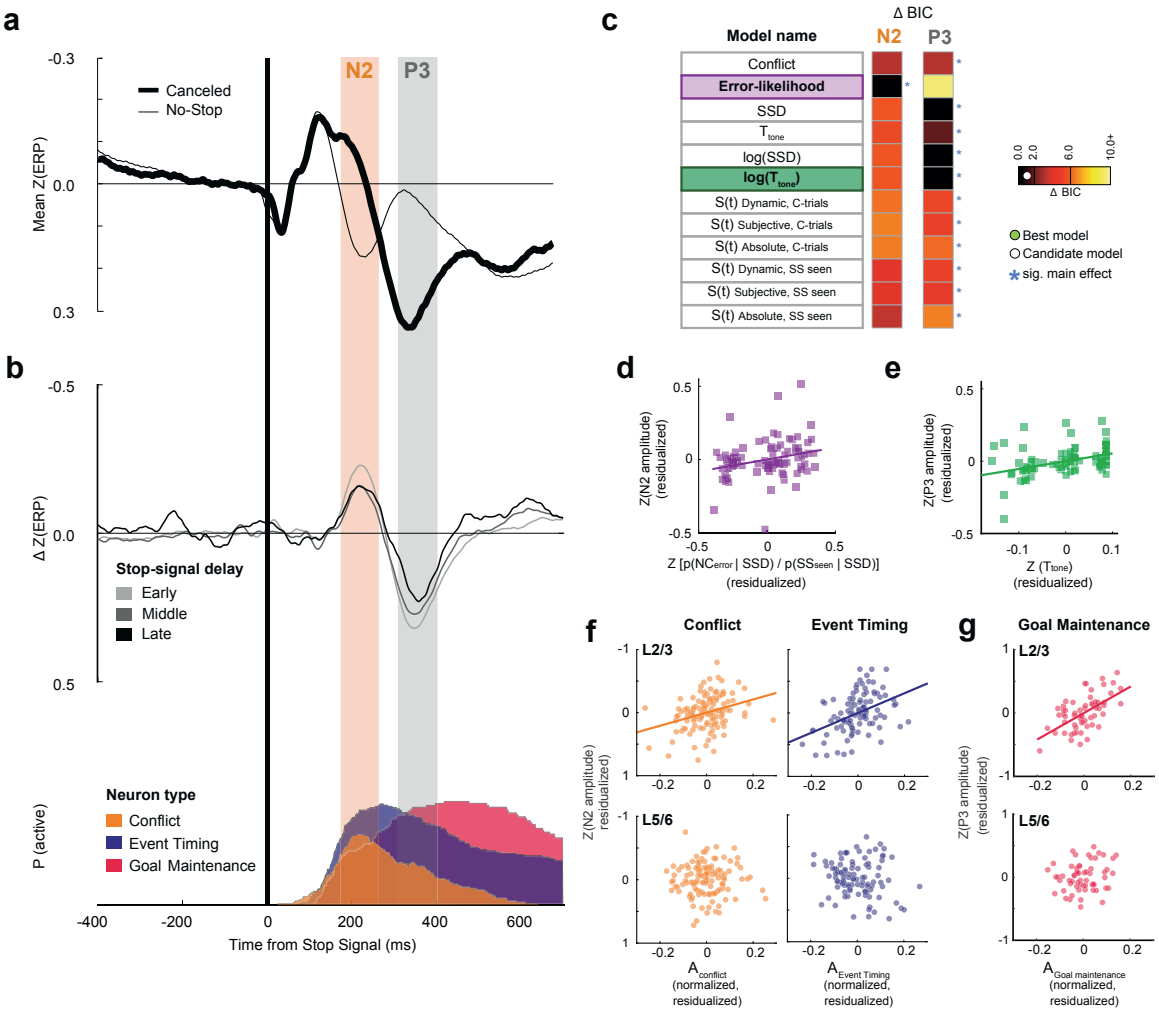

Figure 5.6: **Event-related potentials for successful response inhibition.** Conventions as in Figure 5.3. **A**. Grand average z-transformed EEG on canceled (thick) and latency-matched no stop-signal (thin) trials. **B**. Difference functions (top) remove stimulus-evoked ERP to highlight N2 and P3 components 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_{tone}). Full statistics in Appendix C, Supplementary Table 6. **D**. Significant variation of N2 amplitude as a function of p(NC_{error} | SSD) / p(SS_{seen} | 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.

Appendix C, Supplementary Figure 8a; Appendix C, 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 (Figure 5.6f). 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; Appendix C, Supplementary Figure 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; Figure 5.6g; Appendix C, Supplementary Figure 8c,d; Appendix C, 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; Appendix C, Supplementary Figure 8c; Appendix C, Supplementary Table 8).

5.4 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 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 SEF (Sajad et al., 2019). 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 connectivity of SEF (Huerta and Kaas, 1990; Parthasarathy et al., 1992; Shook et al., 1990), 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 (Figure 5.7). Also, complementing our earlier description of the source of the ERN (Sajad et al., 2019), 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.

5.4.1 Conflict, Surprise, Salience and Dopamine

In this study, we report a population of SEF neurons with a 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(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) (Hanes et al., 1998; Pare and Hanes, 2003). Neurocomputational models demonstrate how this network can instantiate the GO and STOP processes (Boucher et al., 2007; Lo et al., 2009; Logan et al., 2015). Noncanceled saccades happen when the gaze-holding STOP units do not interrupt the rise to 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 conflict (Botvinick et al., 2001). The multiplicative conflict between GO and STOP accumulator units scales with p(NC | SSD) and peaks following SSRT (Figure 5.2a) 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 countermanding (Ogasawara et al., 2018) offer an alternative interpretation for these neurons. The modulation of DA neurons in dorsolateral SNpc scales with p(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 (Bromberg-Martin et al., 2010; Ogasawara et al., 2018). 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 (O'Reilly et al., 2013; Starkweather et al., 2017). 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 cancelation after the infrequent stop-signal. To be most useful, such an eligibility trace must be salient and may be surprising.

5.4.2 Event Timing & 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 reward, monkeys must hold

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 (Errington and Schall, 2020; Nelson et al., 2010). A contribution of SEF and nearby areas in action timing and time production tasks has been demonstrated (Egger et al., 2020; Ohmae et al., 2008; Wang et al., 2018). 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 explanation (Stuphorn et al., 2010). 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 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 (Coull and Nobre, 2008; Egger et al., 2020; Kim et al., 2013; Ohmae et al., 2008; Tallot and Doyere, 2020; Wang et al., 2018; Zhang et al., 2019). 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 uncertainty (Zhang et al., 2019). 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 reinforced by the distinct laminar distribution of the two groups of neurons.

Another class of neurons produced persistent activity on canceled trials after SSRT lasting for most neurons 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 (Schultz et al., 2015; So and Stuphorn, 2010). However, the activity of these SEF neurons was definitely 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 rate 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 duration (Tallot and Doyere, 2020). 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 reward on canceled trials, monkeys needed to sustain unblinking fixation on the stop-signal until the feedback tone (Figure 5.1a). Hence, we conjecture that these neurons contribute to goal maintenance. This conjecture is consistent with an original theory of response inhibition (Logan and Cowan, 1984) and previous evidence linking SEF to working memory (Bastos et al., 2018; Reinhart et al., 2012a) and working memory to time estimation (Gibbon et al., 1984; Polti et al., 2018). Future work can discriminate time-based and surprise parameters and evaluate the link between surprise and goal maintenance.

5.4.3 Cortical Microcircuitry of Executive Control

By combining the laminar distribution of the neuron classes described in this study with 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 (Figure 5.7).

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 conjectures (Cohen, 2014) and recent biophysical modeling (Herrera et al., 2020) 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 errors (Bastos et al., 2012) 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 to 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 thalamus (Atsma et al., 2014). Consistent with its original conception, communicating conflict (or salience or surprise or just difficulty) to multiple areas simultaneously can coordinate adaptive changes in behavior (Botvinick, 2007).

The unexpected parallels between SEF and SNpc modulation patterns invites consideration of cause and effect. SEF is innervated by DA neurons in SNpc (Williams and Goldman-Rakic, 1998). Whilst SNpc DA neurons modulated significantly earlier than the SEF Conflict neurons (Appendix C, Supplementary Figure 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 axons (Grace and Bunney, 1980) ($\sim$100 ms conduction time from SNpc to SEF, Appendix C, Supplementary Figure 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 (Huerta and Kaas, 1990; Shook et al., 1990), 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 SC (Redgrave and Gurney, 2006). Through the Conflict neurons in L5, SEF can influence SC directly (Huerta and Kaas, 1990). Curiously, though, the modulation specifically after SSRT scaling with related performance parameters has not been observed in SC (Pare and Hanes, 2003).

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 (Buhusi and Meck, 2005; Egger et al., 2020; Emmons et al., 2020). Axon terminals from SEF are dense in the caudate nucleus (Parthasarathy et al., 1992), arising from pyramidal neurons in L3 and L5 (Griggs et al., 2017; Saint-Cyr et al., 1990). In fact, neurons with this pattern of modulation have been described in a recent investigation of the caudate nucleus of monkeys performing saccade countermanding (Ogasawara et al., 2018). 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 (Appendix C, Supplementary Figure 9, Ogasawara et al., 2018).

The rapid suppression of the ramping activity after SSRT merits consideration. One source can be intracortical inhibition from the narrow-spike, 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 dorsolateral prefrontal cortex (DLPFC) and selectively inhibit pyramidal neurons (Medalla and Barbas, 2009), although our methods are unlikely to sample their spikes. We note that although SEF is an agranular structure with weak interlaminar inhibitory connections (Beul and Hilgetag, 2014), CR neurons in L2/3 can potently inhibit L5 neurons through specialized projections on the apical dendrites (Gidon and Segev, 2012). 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 DLPFC and ACC, terminating in L2/3 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 signals (Fiorillo et al., 2008), can enhance these influences through higher density of D_1 relative to D_2 receptors in L2/3 (Lidow et al., 1991). The sustained discharge can be maintained through recurrent activation within SEF and between other structures (Bastos et al., 2018). 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 (Lim and Goldman, 2013; Wang et al., 2004). 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 D_2 neurons in the striatum (Griggs et al., 2017; Saint-Cyr et al., 1990). Intrinsic inhibitory Goal Maintenance neurons can suppress movement by inhibiting pyramidal neurons projecting the direct pathway D_1 and to the frontal eye field. As PV neurons in primates largely lack extrinsic connections (Tremblay et al., 2016), we propose that this can be mediated by the inhibition of other excitatory neurons (unidentified neurons and possibly Gain neurons identified in Sajad et al., 2019) that send projections to these motor structures (gray neurons in Figure 5.7, right panel). Therefore, Goal Maintenance neurons can achieve their role by altering the balance in the push-pull mechanism mediated by the direct (D_1) and indirect (D_2) 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 RT (Sajad et al., 2019). 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 task epoch.

We note that neurons with facilitated activity after SSRT were described in an investigation of the caudate nucleus of monkeys performing saccade countermanding (Ogasawara et al., 2018). The facilitation in the caudate nucleus coincided with that measured in SEF (Appendix C, Supplementary Figure 9). The parallel between SEF and the striatum in patterns of modulation associated with proactive but not reactive inhibition are surprisingly, but satisfyingly, clear.

5.4.4 Event Related Potentials

We showed that macaque monkeys exhibit a N2/P3 ERP complex homologous to humans (Kok et al., 2004). Disagreement persists about what the frontal N2 and P3 index (Huster et al., 2013, 2020). 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 (Kok et al., 2004; Smith et al., 2010).

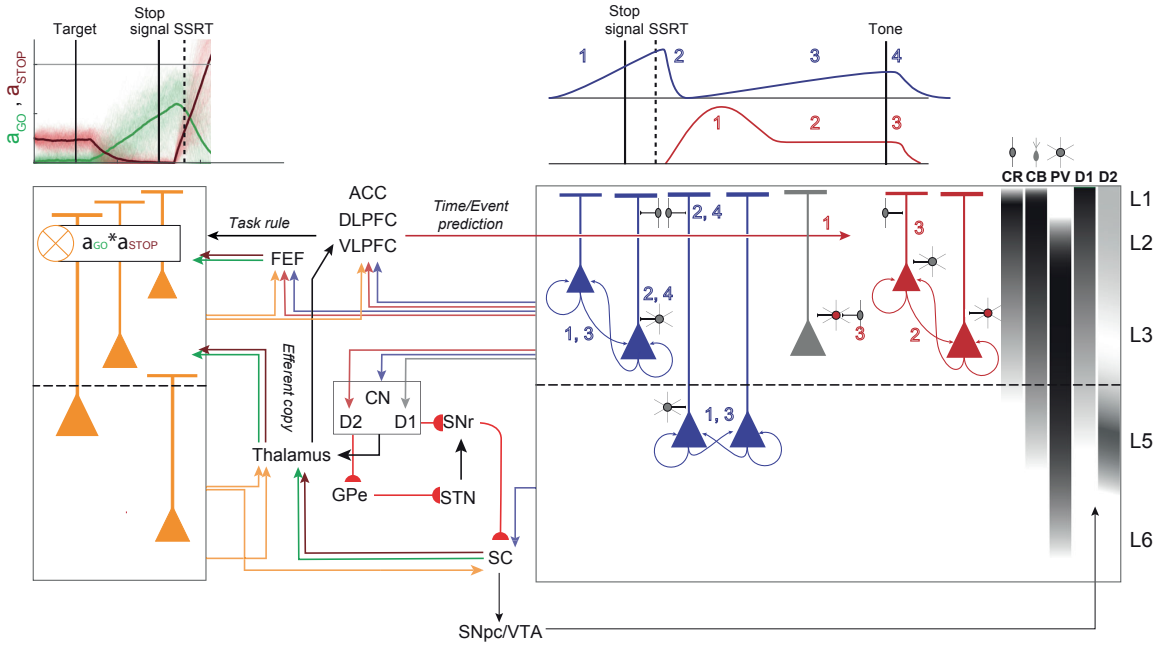

Figure 5.7: **Executive control circuitry.** Conflict (orange), Event Timing (dark blue), and Goal Maintenance (dark red) neurons are illustrated with selected anatomical connections and laminar densities of calretinin (CR), calbindin (CB), and parvalbumin (PV) neurons and of D_1 and D_2 receptors (right). 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 signal 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 about an upcoming event. Ramping results from recurrent connections (1, dark blue). SEF can receive information about 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 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 (D_2) pathway and inhibitory neurons can project to other pyramidal neurons, unobserved in this study (gray), that project to the direct (D_1) 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 though CR and CB inhibition of the sustained activity (3). Further details in text.

Consistent with previous reports of P3 indexing expectation and timing of behavior (Huster et al., 2020), 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 spiking of specific neuron classes in L2/3 and not L5/6. N2 magnitude was unrelated to spiking of Goal Maintenance neurons but co-varied with spiking of Conflict and Event Timing neurons, as well as 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 potentials (Buzsáki et al., 2012). 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 contributor. 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 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.

5.4.5 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 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 signals (Sajad et al., 2019). 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 (Ramakrishnan et al., 2017; Rigotti et al., 2013). 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 (Fu et al., 2022a; Rigotti et al., 2013), 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 towards formulating more specific hypotheses about how neural networks function (Dubreuil et al., 2021).

5.4.6 Conclusion

In conclusion, the present results add to the first catalogue of the laminar functional architecture of an agranular frontal lobe area. Pioneering insights into the microcircuitry and mechanisms of primary visual cortex began by describing the properties of neurons in different layers (Hubel and Wiesel, 1968). 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 formal (Logan and Cowan, 1984), algorithmic (Boucher et al., 2007; Logan et al., 2015), and spiking network (Lo et al., 2009) models of performance of a task with clear clinical relevance (Thakkar et al., 2011).

CHAPTER 6

Absence of conflict signaling in midcingulate cortex: Neural spiking in macaques during a stop-signal task

6.1 Introduction

We learn from mistakes. This ability is enabled by the medial frontal cortex—specifically the midcingulate cortex (MCC)—that detects conflicting plans, reward opportunities, and action outcomes (Ridderinkhof et al., 2004; Rushworth et al., 2004; Botvinick et al., 2001). These observations are explained by two theories of cognitive control that make explicit predictions about expected neural activity. Expected value of control theory predicts activity proportional to conflict or choice difficulty. The action-outcome theory predicts activity proportional to the evaluations of our actions (Shenhav et al., 2016a; Kolling et al., 2016).

Researchers agree that the MCC signals action outcomes and expected value based on single-neuron studies in macaques (Amiez et al., 2006; Blanchard and Hayden, 2014; Cai and Padoa-Schioppa, 2012; Hayden et al., 2011a,b; Ito et al., 2003; Kawai et al., 2015; Kennerley et al., 2011, 2009; Luk and Wallis, 2013; Nakamura et al., 2005; Oemisch et al., 2019). Researchers do not agree that MCC signals conflict because some report positive results (Ebitz and Platt, 2015) but others find no evidence for conflict in MCC (Ito et al., 2003; Nakamura et al., 2005).

No claim can be made to understand a cortical area without detailed knowledge of its functional architecture and intrinsic circuitry. This knowledge begins with information about the distribution of functional signals between distinct cortical areas. In macaques, the cytoarchitecture of the dorsal (dMCC) and ventral (vMCC) banks of the midcingulate cortex differ with the dorsal bank being an extension of the neighboring pre-supplementary motor area (Vogt et al., 2005). The ventral bank of the midcingulate cortex appears to have more subdivisions evidenced by different laminar densities of morphologically distinct neurons (Vogt, 2016; Rapan et al., 2021). Despite these differences, nothing known about how the functional architecture of the cortical areas in the cingulate sulcus contributes to performance monitoring. Neurophysiological investigations have mainly sampled spiking activity in the dorsal bank (Heilbronner and Hayden, 2015; Procyk et al., 2016) while only a few studies to date have compared the neural signals in the two banks of the midcingulate cortex in macaques (Cai and Padoa-Schioppa, 2012; Procyk et al., 2016; Shidara and Richmond, 2002; Monosov, 2017). Understanding the differences in functional signals between the dMCC and vMCC is as relevant to our understanding of how performance monitoring signals function as the differences between V1 and V2 or MT and V4 are to our understanding of visual processing. In this study, I report preliminary data on the activity across both dMCC and vMCC during successful inhibition.

6.2 Methodology

6.2.1 Animal care and usage

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 was collected from one male bonnet macaque (Monkey Da, *Macaca Radiata*, 9.5 kg) and one male rhesus macaque (Monkey Jo, *Macaca Mulatta*, 13.8 kg). Both animals were on a 12-hour 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.

6.2.2 Surgical Procedures

Prior to surgery, 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 anatomical sequence (TR = 8.729 ms; 130 slices, 0.70 mm thickness). These images were used to ensure Cilux recording chambers (Crist Instruments) were placed in the correct location and orientation. For monkey Da', the medial frontal chamber was implanted normal to the cortex centred on the midline, and 30 mm anterior to the interaural line. For monkey Jo', the medial frontal chamber was implanted normal to the cortex, -1.2 mm lateral to the midline, and 33.3 mm anterior to the interaural line.

All surgical procedures were carried out under aseptic conditions. Access to food was withdrawn 12 h before surgery. Animals were sedated with ketamine (10–30 mg/kg) and provided with an initial dose of buprenorphine (0.005–0.010 mg/kg) to alleviate postoperative discomfort. Ophthalmic ointment was applied to prevent corneal drying. Robinul (0.004–0.008 mg/kg) was administered to minimize mucosal secretions and help prevent vagal bradycardia. Animals were intubated, and catheters were inserted into the saphenous veins for the administration of support fluids throughout the procedure. Monkeys were anesthetized with an isoflurane-oxygen mixture (1–3% C3H2CIF5O), shaved, positioned in stereotax, and scrubbed. ECG, rectal temperature, respiration, and blood pressure were monitored. Expiratory CO_2 was maintained at $\sim$4%. After the subcutaneous administration of lidocaine ($\sim$1–2 ml of 2% solution), the subjects' skulls were exposed, and MRI-compatible headposts were firmly attached with orthopedic screws to immobilize the animals' heads during testing. Cilux recording chambers were also placed during this surgery. In consultation with attending veterinarians, analgesics (Bupronorphine, 0.005–0.010 mg/kg) and prophylactic antibiotics (Naxcel, 2.2 mg/kg) were administered for at least 3 days after surgery.

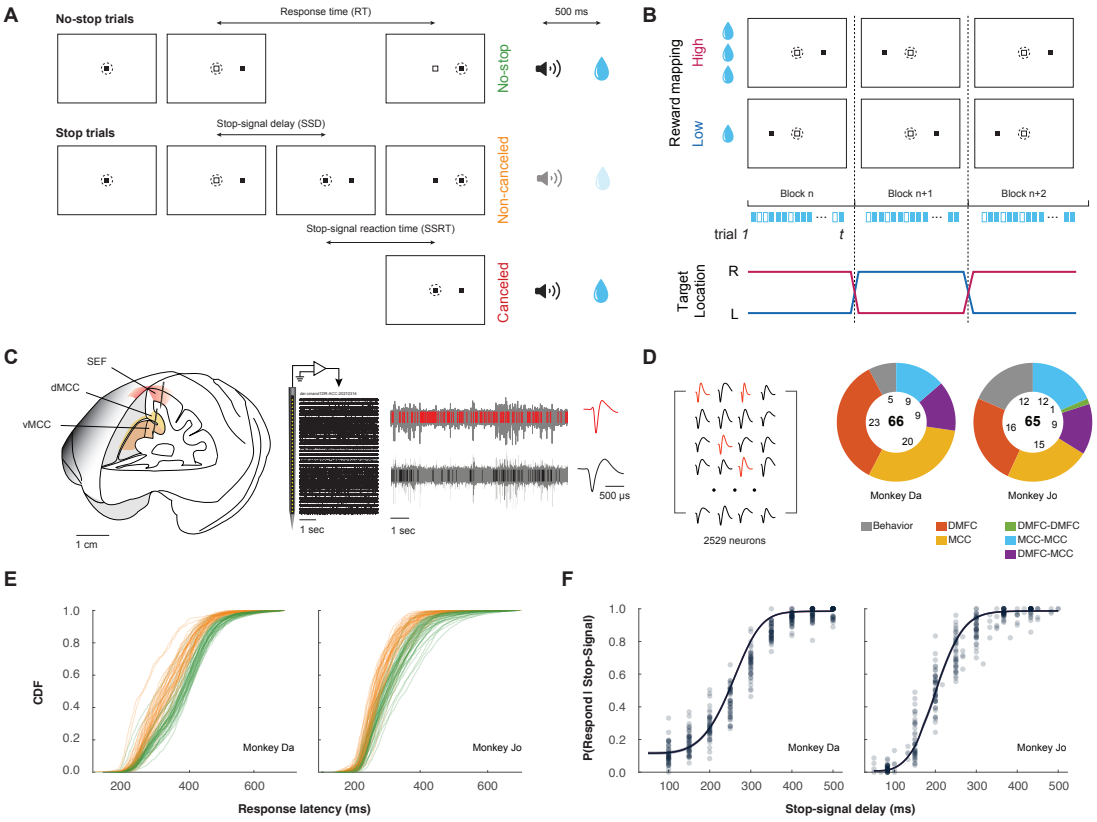

Figure 6.1: **Experimental approach to assessing single-neuron activity in MCC**. **A**. Saccade countermanding task. Monkeys initiated trials by fixating on a central point. After a variable time, the center of the fixation point was extinguished. A peripheral target was presented simultaneously at one of two possible locations. On no-stop-signal trials, monkeys were required to shift their gaze to the target, whereupon after 600 ± 0 ms a high-pitched auditory feedback tone was delivered, and 500 ms later fluid reward was provided. On stop-signal trials (40% of trials), after the target appeared the center of the fixation point was re-illuminated after a variable stop-signal delay, which instructed the monkey to cancel the saccade in which case the same high-pitched tone was presented after a 1,200 ± 0 ms hold time followed, after 500 ± 0 ms by fluid reward. The stop-signal delay was adjusted such that monkeys successfully canceled the saccade in 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**. Asymmetrical reward design. Monkeys received juice reward for generating a saccade to the target in no-stop trials, and for holding fixation at a central point in canceled trials. The amount of reward (0.1 ml, low; 0.3 ml, high) was cued by the location (left or right) of the target. This mapping switched on a block-by-block basis. **C**. Neurophysiology methods. Neural activity was recorded from supplementary eye field, dorsal midcingulate cortex, and ventral midcingulate cortex. Single units were isolated from the broadband signal. Units with both broad (black) and narrow (red) spikes were observed. **D**. Recording demographics. Across all penetrations, we preliminarily isolated 2529 units. These were a mix of broad (black) and narrow (red) spikes. In a given session, neural data was sampled from either a single penetration or multiple simultaneous penetrations; the sampling characteristics are represented in the donut charts. **E**. 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 for both monkeys. **F**. Inhibition function plotting the probability of responding across stop-signal delays. 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 both monkeys. The probability of responding increased with stop-signal delay.

6.2.3 Data Collection Protocol

An identical daily recording protocol across monkeys and sessions was carried out. In each session, the monkey sat in an enclosed primate chair with their head restrained 57 cm from a CRT monitor with a 60 Hz refresh rate (NEC MultiSync FE922-BK). Eye position data was collected at 1 kHz using an EyeLink 1000 infrared eye-tracking system (SR Research, Kanata, Ontario, Canada). Behavioral and neural data was streamed to a single data acquisition system (TDT System 3, Tucker-Davis Technologies, Alachua, FL). Time stamps of trial events, eye position, photodiode signals, and reward solenoid signals were recorded at 2035 Hz.

6.2.4 Neural signals

Electroencephalogram (EEG) was recorded from the cranial surface with an electrode placed over the medial frontal cortex, resting on top of the dura, next to the electrode penetration site. Intracortical spiking activity and local field potentials were recorded using NeuroNexus acute vector arrays and Plexon S-probes. Both electrode types had 50, 100, or 150 μm spacing between contacts, allowing unbiased sampling across cortical areas. NeuroNexus arrays (V1x32-Edge-10mm-150-177-V32-100-50) or Plexon S-probes (PLX-SP-32-15SE(150)-(420-100)(640-30)800-1-CON/32m-V) were housed within custom built guide tubes made from 21-gauge stainless steel hypodermic tubing (Small Parts Inc., Logansport, IN). This tubing had been cut to length, deburred, and polished so that they effectively supported the electrodes as they penetrated dura and entered the cortex. For both intracortical electrodes, contacts were referenced to the guide tube in which they were housed and grounded to the headpost. For EEG, the electrode was referenced and grounded to the earth.

Microdrive adapters were fit to recording chambers with < 400 μm 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 electrodes 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.

Electrodes were connected to the recording system through high-input impedance ZIF-clip head stages (ZD32, Tucker-David Technologies), and signals digitised through a preamplifier (PZ5, Tucker-David Technologies) and streamed to a single data acquisition system (TDT System 3, Tucker-Davis Technologies, Alachua, FL) at 24,414 Hz. Neural data was filtered online through the Synapse software (Tucker-Davis Technology). EEG signals were bandpass filtered between 0.7 and 170 Hz, with notch filters at 60 Hz and 120 Hz. Local field potentials were bandpass filtered between 3 and 300 Hz, with notch filters at 60 Hz, 120 Hz, and 180 Hz. Spiking data were derived from a broadband signal, bandpass filtered between 300 and 5000

Hz. Single units were extracted offline, identified using automated spike sorting via Kilosort, and manually curated thereafter.

6.2.5 Value-context stop-signal (countermanding) task

Data was recorded from monkeys performing a saccade countermanding task (Hanes and Schall, 1995). Recently, a set of guidelines has been proposed for designing and analysing the stop-signal task to allow for valid comparisons to be made across studies (Verbruggen et al., 2019a). 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, 30-40% of trials in our study were stop trials, compared to the recommended 25%. This higher value was used to achieve the necessary power to analyse neural data at the individual stop-signal level, but it did not introduce excessive slowing of responses (Emeric et al., 2007). Secondly, although we employed a staircase procedure, this stepped between one to three stop-signal delays. We do this to prevent monkeys from anticipating the staircase (Nelson et al., 2010).

Briefly, trials were initiated when monkeys fixated a central point. After a non-aging foreperiod, the centre of the fixation point was removed leaving an outline. At this point, a peripheral target was presented simultaneously on either the left (0^{o}) or right (180^{o}) hand of the screen. 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 centre of the fixation point was re-illuminated (stop-signal trials); this stop signal appeared at a variable time after the target had appeared (stop-signal delay; SSDs). An initial set of SSDs, separated by either 50 or 80 ms, were 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 (between 1 and 3 steps) 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 approximately 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 for 1000 ms. Following a correct response, an auditory tone was sounded 500 ms later, and followed another 500 ms later by a fluid reward.

Reward on correct trials could be low or high magnitude, depending on the stimulus-response mapping; the target location presented at the beginning of the trial cued the monkey to whether the trial would lead to a low or high magnitude of fluid reward. The stimulus-response mapping of location-to-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 between 20 ± 8 correct trials. For both monkeys, the higher magnitude reward was approximately three times the amount of the lower magnitude reward amount (i.e. 0.3 ml to 0.1 ml).

6.2.6 Analysis of neural activity distinguishing dorsal and ventral MCC.

To examine signatures that may differentiate between dorsal and ventral midcingulate cortex, we computed the power spectral density to determine the relative power of the oscillatory activity across cortical depth (Bastos et al., 2018; Maier et al., 2010; Westerberg et al., 2019). Local field potentials were filtered between 1 and 200 Hz with a 4th-order bidirectional Butterworth filter and full-wave rectified to estimate power at each frequency. Power was normalized relative to the mean of a given electrodes power at each frequency. In addition to comparing power over the electrode, we additionally calculated the signal correlations of the LFP activity between all channel combinations (Sotero et al., 2015; Westerberg et al., 2019). This allows for the distinction of contacts with similar activity.

6.2.7 Analysis of neural spiking.

Spike widths exhibited a bimodal distribution and neurons were distinguished as narrow- (peak-to-trough duration $\leq$ 250 μs) or broad-spikes ($>$ 250 μs). 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:

$$SDF = (1 - \exp(\frac{t}{\tau_g})) \times \exp(\frac{-t}{\tau_d})$$

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 post-synaptic potentials. 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. We first determine neurons that demonstrated significantly different activity between these trials in the -50 to 250 ms period around SSRT. These neurons were distinguished by patterns of modulation consisting of periods of facilitation or suppression using a consensus clustering algorithm (Fig. 6.3A, Lowe and Schall, 2018). The input to this analysis pipeline was the spike-density function on canceled trials and on latency-matched no-stop-signal trials during the 100 ms preceding SSRT and 600 ms following SSRT.

6.2.8 Analysis of conflict-related activity.

To examine the relationship between neural activity and the behavioral estimate of conflict, the performance and neural quantities were averaged within groups of early-, mid-, and late-SSD trials. The mid-SSD trial reflects the SSD with a p(respond | stop-signal) closest to 50%. Early and late SSD's are the SSD proceeding and following this middle SSD. To allow for fair comparisons between units, neural activity was normalized so that values were relative to the maximum firing rate in a 100 ms peri-SSRT window across all SSD's; i.e. the highest firing rate during this period was equal to one. This allowed for comparison of firing rates at each SSD to be collapsed across all units in each cluster.

To determine if monkeys demonstrate post-conflict slowing, we identified no-stop trials that directly proceed canceled trials, and no-stop trials that directly follow canceled trials. We then calculated the change in reaction time between these trials to get an estimate of post-conflict slowing.

6.3 Results

6.3.1 Countermanding performance & neural sampling

Neurophysiological and behavioral data was recorded from two macaque monkeys (Monkey Da & Monkey Jo) performing the saccade stop-signal task with explicit feedback tone cues preceding possible delivery of fluid reward (Figure 6.1A). Data collection and analysis was informed by the consensus guide for the stop-signal task (Verbruggen et al., 2019b). Briefly, monkeys earned fluid reward for shifting 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) stop-signal trials. Although unstudied in this paper, reward context was also manipulated within this task with the location of the target cueing whether correct performace would lead to a small or large reward (Figure 6.1B). Whilst monkeys performed this task, we sampled neural activity across the dorsomedial frontal cortex (including supplementary eye field) and the dorsal and ventral banks of the midcingulate cortex (MCC) using linear electrode arrays (NeuroNexus, 32/64 channels, 50/100/150 μm spacing, Figure 6.1C). In some sessions, we recorded from one cortical area only and in others we conducted multi-site recordings (Figure 6.1D). Here we focus on preliminary findings from sessions in which we recorded from the MCC.

Both monkeys exhibited typical sensitivity to the stop-signal. First, response latencies on non-canceled (error) trials were faster than those on no-stop trials (Figure 6.1E). Secondly, the probability of failing to inhibit a saccade was greater at longer stop-signal delays (Figure 6.1F). These two observations validated the assumptions of the independent race model (Logan and Cowan, 1984), allowing us to estimate the stop-signal reaction time (SSRT). Monkeys had indistinguishable mean SSRT (Da: 121.46 ± 1.61 ms; Jo: 114.82 ± 3.85

ms; Independent groups t-test, t (81) = 1.577, p = 0.119).

6.3.2 Distinguishing dorsal and ventral midcingulate cortex

We first aimed to distinguish activity between the dorsal and ventral midcingulate cortex. In macaques, the cytoarchitecture of the dorsal and ventral banks of the midcingulate cortex differ with the dorsal bank being an extension of the neighboring pre-supplementary motor area (Vogt et al., 2005). The ventral bank of the midcingulate cortex appears to have more subdivisions evidenced by different laminar densities of morphologically distinct neurons (Vogt et al., 2005, Rapan et al., 2021, Figure 6.2A). We used MRI guided penetrations into the cortex (Figure 6.2B) to compare and contrast signals across the cingulate suclus.

Considering the broad anatomy, the dorsal and ventral banks of the cingulate sulcus are folded such that vMCC is reflected horizontally in a coronal section relative to dMCC. Separately, recent findings have demonstrated that local-field potentials in the cortex have a ubiquitous spectrolaminar pattern of activity - gamma-frequency power is more pronounced in upper cortical layers, whilst alpha-frequency power is more pronounced in lower cortical layers (Mendoza-Halliday et al., 2022). Given the anatomical configuration of the cingulate sulcus and neurophysiological features of cortex, we should observe alpha power more prominently in the contacts distal to the sulcus and gamma power more prominently in the contacts proximal to the sulcus. As the sulcus has no neural activity, we shouldn't observe a clear dominance of any frequency. We observed this expected spectrolaminar pattern of actvity within our data (Figure 6.2C). We further observed clear clustering of contacts with highly correlated activity, differentiating between more distal and more proximal contacts (Figure 6.2D). These observations were reliable and robust across penetrations. Across all 104 penetrations into MCC, we confidently identified the cingulate sulcus in 94 penetrations (90%).

In the 94 penetrations into MCC in which we could identify the cingulate sulcus based on the pattern of power-spectral density and cross-contact correlation, we isolated 1599 units (Da: 938 units; Jo: 661 units). These units were a mix of narrow- (298/1599) and broad- (1301/1599) spiking neurons. We then assigned units to a depth relative to the cingulate sulcus, given their localisation to a particular electrode contact and the intercontact spacing of our electrodes. This approach yielded a clear bimodal distribution of neurons, with clear peaks approximately ±1000 um relative to the identified sulcus channel (Figure 6.2E).

6.3.3 Neural activity in midcingulate cortex during successful stopping

To identify 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 behavioral temporal dynamics, identified through a process of response 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).

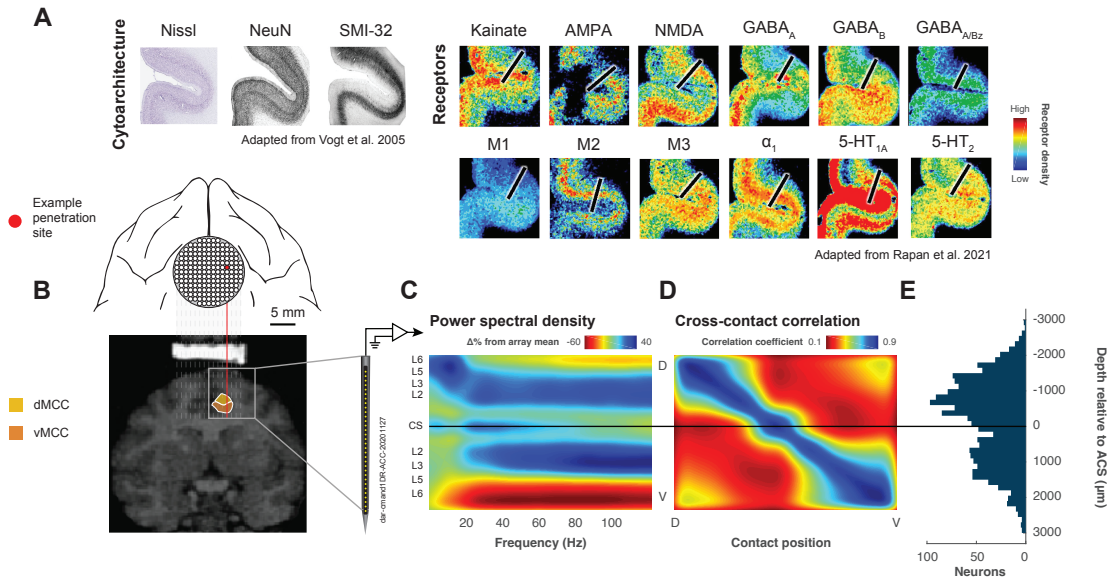

Figure 6.2: **Neural signatures of dorsal and ventral midcingulate cortex**. **A**. Cytoarchitecture (adapted from Vogt et al., 2005) and receptor density (adapted from Rapan et al., 2021) distinguishing the dorsal and ventral banks of the midcingulate cortex. Blue reflects low receptor density, red reflects high receptor density. **B**. MRI-guided recording. A chamber was placed over the medial frontal cortex of both monkeys. MRI images were captured following chamber placement and combined with the chamber grid projections to determine viable sites for MCC sampling. An example site at 30.5 mm anterior to the stereotaxic center and 2 mm lateral to the midline for monkey Da' is presented. **C**. Power spectral density. We computed the power spectral density to determine relative power of oscillatory activity across all contacts of our electrode. Blue and red represent higher and lower power relative to the mean power across the entire electrode. We observed a pattern of increased low-frequency power which then transitions into increased high-frequency pattern in the top half of our electrode. This pattern is mirrored in the bottom half of our electrode, divided by a band of increased low-frequency power. This reflects a hypothesized spectrolaminar pattern that allows for the distinction of dorsal and ventral banks. **D**. Cross-contact correlation. We computed the contact-to-contact correlation of activity within each of our penetrations. Blue and red represent higher and lower correlation coefficients. Activity appears to cluster into three components: dorsal, mid, and ventral. These correspond to contacts in which the dominant power moves from lower to higher frequency. E. Distribution of neurons relative to the putative cingulate sulcus. After determining the channel corresponding to the gap between dorsal and ventral banks, we assigned neurons to depth relative to the CS, given their localisation to a particular electrode contact and the intercontact spacing of our electrodes. This resulted in a bimodal distribution.

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Non-modulated	Total
dMCC	30	100	31	105	716	982
vMCC	10	80	21	46	460	617
Total	40	180	52	151	1176	1159
Chi-square	3.1958 p = 0.074	2.9371 p = 0.087	0.0733 p = 0.7865	4.463 p = 0.031	0.525 p = 0.469	-

Table 6.1: Number of neurons in dorsal and ventral MCC, divided by cluster

We then classified neurons based on the pattern of neural activation -50 ms to 600 ms relative to SSRT using a semi-supervised consensus cluster algorithm (Lowe and Schall, 2018, Figure 6.3A). This revealed four distinct populations of neurons that qualitatively differed by two features (Figure 6.3B).

First, clusters differed on whether they preferentially fired for inhibition or movement. Clusters 1 (n = 40) and 3 (n = 52) demonstrated significantly greater activity when a saccade was correctly generated compared to when one was inhibited. In contrast, clusters 2 (n = 180) and 4 (n = 151) demonstrated significantly greater activity when a saccade was inhibited (Figure 6.3B).

Second, clusters also differed in the temporal pattern of differential activity (Figure 6.3B). Clusters 2 and 3 distinguish between canceled and no-stop trials following the presentation of a stop-signal but prior to SSRT. Conversely, clusters 1 and 4 distinguish between canceled and no-stop trials well before the stop-signal, following target onset. This pattern of modulation warrants further interpretation. For neural activity to reactively contribute to response inhibition, then it must demonstrate a significant change in dynamics after the stop-signal appears but before SSRT. If we observe differential activity prior to the stop-signal, then this may instead reflect proactive control process that can bias the system towards going or stopping ahead of time. Given these criteria, the properties of our classes of neurons allow them to be considered in two separate conceptual categories: cluster two and three can contribute to reactive inhibition, whilst clusters one and four instead may reflect proactive inhibition signals.

To elucidate the contributions of these neurons to cortical circuitry, we examined spike duration and distribution across the midcingulate cortex. Narrow-spiking neurons were equally as common across all clusters (Figure 6.3C, Chi-square test of homogeneity (one-tailed) χ^2 (3, N = 423) = 0.039, p = 0.998), and reflected the sampling distribution (Chi-square test of homogeneity (one-tailed) χ^2 (1, N = 1599) = 3.441, p = 0.064). Spatially, we observed neurons in all clusters were found in equal proportions in dorsal and ventral MCC (Figure 6.3D , Table 6.1).

6.3.4 Conflict in the stop-signal task

Performance of the stop-signal task is explained as the outcome of a race between stochastic GO and STOP processes (Logan and Cowan,1984), instantiated by specific interactions enabling the interruption of the GO

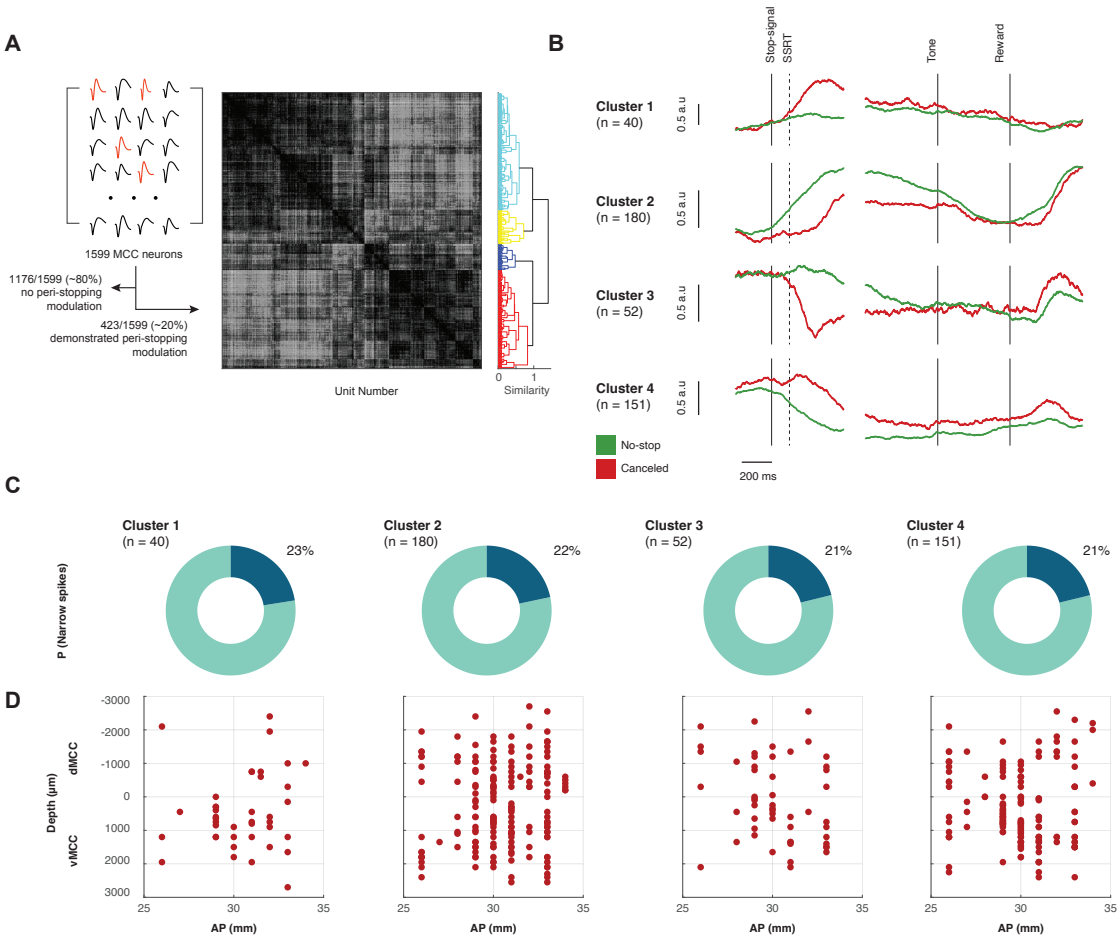

Figure 6.3: **Neural activity contributing to successful stopping in MCC**. **A**. Clustering approach. We observed 1599 units in penetrations that we could confidently identify the cingulate sulcus. 423 neurons demonstrated significant differential post-stopping modulation between trials in which a saccade was correctly generated and trials in which a saccade was correctly inhibited. Their respective spike density functions (SDF) were submitted to an semi-supervised consensus clustering pipeline (Lowe and Schall, 2018) yielding the composite similarity matrix (z-score color map) and associated dendrogram. Consensus clustering yielded 4 clusters. **B**. Cluster spike density functions. Clusters demonstrated clear and distinct patterns modulation on canceled (red) and no-stop trials. Activity is presented aligned on stop-signal (mean stop-signal reaction time is shown as a dashed line), and on tone. **C**. Cluster spike widths. Neurons in each cluster were a mix of broad (light) and narrow (dark) spikes, observed at an incidence expected given the sampling distribution. D. Spatial distribution of cluster neurons. Neurons were sampled across the anterior-posterior (AP, x-axis) and across a range of depths (y-axis).

process by a STOP process (Boucher et al., 2007, Logan et al., 2015; Figure 6.4A). A theory of medial frontal function posits that it encodes the conflict between mutually incompatible processes (Botvinick et al., 2001). Such conflict arises naturally as the mathematical product of the activation of the mutually incompatible GO and STOP units (Schall and Boucher, 2007; Sajad et al., 2022). We also can consider that the probability of responding at each SSD, p(respond | stop-signal), 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 (Figure 6.4B). Given this theoretical and computational foundation, we can consider variations in neural activity as a function of p(respond | stop-signal) to reflect conflict. Our simulations, in conjuction with prepositions in the interactive race model, suggest this conflict signal peaks around stop-signal reaction time.

In our preliminary analysis, we observe monkeys slow their responses following successful stopping (one-sample t-test, t (82) = -8.985, $p < 0.001$, Figure 6.4C). This mirrors previous observations of post-conflict slowing (Botvinick et al., 2001; Emeric et al., 2007; Nelson et al., 2010). We then considered how the dynamics of activity varies across different stop-signal delays (Figure 6.4D). The response inhibition function shows the fraction of non-canceled trials increases as a function of stop signal delay, as movements become less likely to be canceled as movement preparation progresses (Figure 6.4E, black points). In conjunction with the psychometric function, a neurometric function can also be derived to examine the activity of neurons as a function of stop-signal. We measured the averaging firing rate of midcingulate neurons in each of our clusters in the -50 to +50 ms peri-SSRT window on canceled trials for early, middle, and late stop-signal delays (Figure 6.4E, red points). We observed the activity of midcingulate neurons was unchanging across stop-signal delays and did not reflect changes in the probability of generating a saccade – a proxy of conflict (Figure 6.4E).

6.4 Discussion

In this study, we aimed to investigate activity at the single-neuron level in midcingulate cortex during successful stopping. Here, I discuss and interpret preliminary findings.

6.4.1 Summary of interim findings

In this study, we unbiasedly sampled neural activity from the medial frontal cortex, including from both banks of the anterior cingulate sulcus. We first developed a novel method in which we could distinguish between the dorsal and ventral midcingulate cortex based on the spectrolaminar profile and the similarity of activity across contacts in our recordings. This method allowed us to use a common landmark to align sessions and examine the differential contributions of these distinct cortical areas.

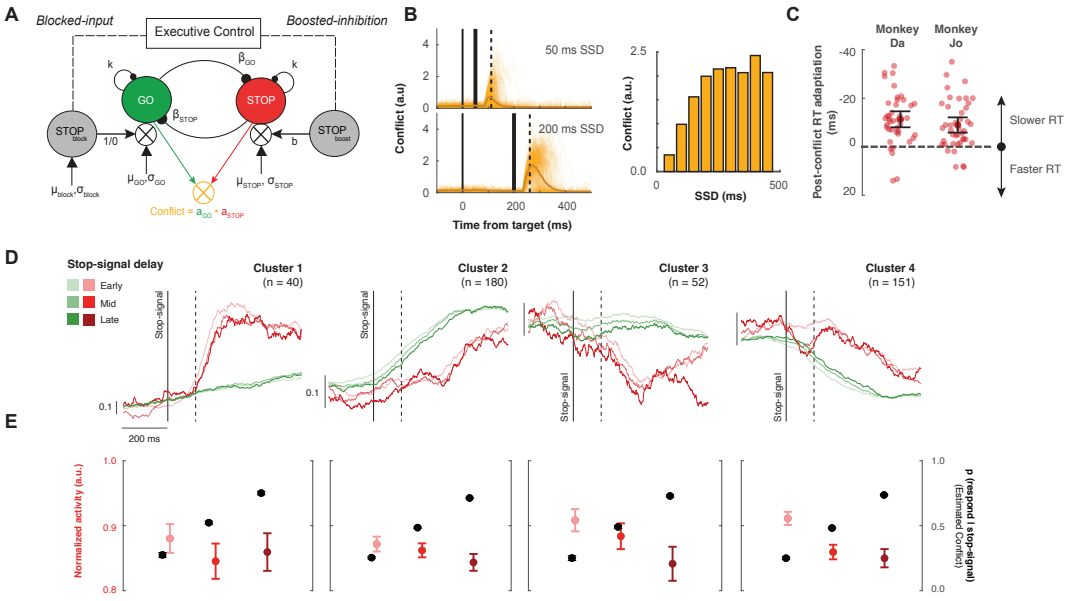

Figure 6.4: **Assessing conflict in midcingulate cortex using saccade countermanding**. **A**. Interactive race model architecture. In interactive race architectures, GO and STOP units are racing, competing response channels, so the product of their co-activation measures response conflict in this task (Boucher et al., 2007; Logan et al., 2015). **B**. Estimating conflict as product of GO and STOP unit. In simulations with parameters simulating observed RT and p(respond | SSD), the instantaneous product of aGO*aSTOP of individual (thin) and average (thick) increases around and peaks after SSRT. The product of the activation of GO and STOP units producing countermanding performance increases as a function of stop-signal delay, similar to p(respond | stop-signal). **C**. Post-conflict adaptation. Both monkeys demonstrate significantly slower response latencies following canceled trials. Individual light-red points are presented for each session. The solid red point and attached black bars represent mean ± SEM. **D**. Cluster activity on canceled (red) and latency-matched no-stop (green) on early (light), middle (regular), and late (dark) stop-signal delays. Activity is presented aligned on stop-signal (mean stop-signal reaction time is shown as a dashed line), and on tone. Activity does not appear to systematically vary across SSD for any cluster. **E**. Mismatch between behavioral estimates of conflict (black, right ordinate scale) and neural activity (red, left ordinate scale) during the expected conflict period (-50 to 50 ms, peri-SSRT). Neural activity did not reflect the estimated conflict.

We then used a semi-supervised clustering approach to find systematic modulations in our sample that may provide insight into the function of MCC. This approach revealed clear patterns of modulation that differentiated between trials in which a saccade was successfully inhibited, and trials in which a movement was correct generated. These neurons were similar in spike width and in their incidence. However, we observed interesting temporal properties of these neurons: some demonstrated prolonged firing rates, beginning even before the onset of the stop signal, whilst others modulated after the presentation of the stop-signal. This distinction is discussed further below.

6.4.2 Reactive and proactive control

Adaptive behavior relies on the ability to meet and adjust to the demands of a changing, complex environment. Response inhibition allows us to stop actions that are inappropriate or may no longer be warranted, given new information about our environment. Whilst this can be exerted reactively and "in the moment," behavior can also be proactively adapted in advance of potential situations where abrupt changes are anticipated ahead of time (Braver, 2012; Aron, 2011). Successful performance in the stop-signal paradigm requires both: control processes need to be engaged in reaction to the perception of a stop-signal but also need to monitor the race between going and stopping and adjust response strategies to balance the competing demands of the two processes (Verbruggen and Logan, 2009b). Stop-signal reaction time provides us with a measure of how long it takes for a planned action to be stopped (Logan and Cowan, 1984). As such, it also allows for a measure in which can be used to distinguish the specific contributions of neural activity to successful inhibition. First, if the neural activity is reactive to the stop-signal and directly inhibits movements, then it must demonstrate a significant change in dynamics after the stop-signal appears and before the inhibition is completed (as indexed by SSRT). Alternatively, if neural activity occurs after SSRT then it is too late to contribute to response inhibition; instead, this activity may reflect aspects of the outcome and performance that need to be monitored or updated for future adjustments in behavior. Finally, if activity differs prior to the onset of the stop-signal between trials in which movements are inhibited and trials in which movements are generated, then this activity may instead reflect proactive control processes moderating or biasing the system to facilitate or suppress movement.

Observations in the motor system describe neurons that suppress their firing rate around the time of the saccade and elevate their firing rates during periods of fixation (Izawa et al., 2009). Other studies in our lab have further described these neurons in the stop-signal task, where they demonstrate a rapid increase in firing rate following the stop-signal and before SSRT when a planned saccade was inhibited (Hanes et al., 1998). Given the criteria outlined above, we can consider fixation neurons in frontal eye field to contribute to reactive inhibition. In our sample, we observed classes of neurons with a similar profile of activity that

could reflect reactive control processes. These neurons did not respond in a similar pattern on non-canceled trials, so these neurons did not simply reflect a visual response to the stop-signal. Furthermore, there is little previous evidence of fixation neurons in midcingulate cortex. Although this doesn't rule out the possibility that these neurons simply impose gaze-holding, the exact function of these signals need to be investigated further.

We also present the first evidence of neurons that may reflect proactive control processes. The probability of successful stopping can also be increased by proactively recruiting control processes in anticipation of a stop-signal. Several studies have demonstrated that participants change their response strategies dependent on the likelihood of encountering a stop-signal (Ramautar et al., 2004; Dimoska and Johnstone, 2008; Emeric et al., 2007) and previous experiences in the task (Emeric et al., 2007). From a cognitive perspective, we can consider that the countermanding task involves balancing two separate task demands: one task is to respond as quickly as possible to a signal, and another is to stop any movements from occuring. As such, success in the going process can imply failure in the stop task and vice versa: faster GO processes yield a higher probability of responding given a stop-signal, whilst slower GO processes result in a lower likelihood of incorrectly responding to the stop-signal. In our sample, we observed two sets of neurons that appear to modulate just after target onset, well in advance of the stop-signal. These neurons appear to be complementary – one group facilitated on canceled trials, and the other facilitated on no-stop trials. As these neurons are recruited well before the stop-signal, we can hypothesize that this activity may reflect a proactive control signal that biases the primary task goal towards going or stopping. This signal may then modulate downstream attentional, perceptual, or motor systems. This requires further validation.

6.4.3 Conflict in the stop-signal task

Early work built on advances in neuroimaging demonstrated the midcingulate cortex showed pronounced activation in several tasks that required that overriding of automatic or learnt behaviors (Carter et al., 1995; Pardo et al., 1990; Corbetta et al., 1991). These observations were later formalized through the conflict monitoring theory – a theory that proposes cognitive control is essential when are multiple competing and incompatible demands, referred to as response conflict (Botvinick et al., 2001).

Conflict can also arise in the stop-signal task when a movement is inhibited. Performance in this task can be considered as a representation of the outcome of a race between stochastic GO and STOP processes (Logan and Cowan, 1984). Successful inhibition occurs when the STOP process is faster than the process. Although behaviorally, these two processes are considered to be independent runners, at the neuronal level we observe that movement neurons reduce their activity when fixation neurons become active, as though they are directly inhibited. Development of neurocognitive models account for this, stipulating that the GO and

STOP processes do not influence each other for most of the race and are independent until very late in the STOP process (Boucher et al., 2007, Logan et al., 2015; Figure 6.4B). At this point, the two units interact, and the STOP unit exerts robust and potent inhibition over the GO unit. Here, there are two competing and incompatible action plans – one to shift-gaze, and one to hold-gaze. This is in line with earlier definitions of conflict (Botvinick et al., 2001). The framework proposed by Botvinick et al. allows us to quantify conflict as the mathematical product of the activation of the mutually incompatible GO and STOP units (Schall and Boucher, 2007; Sajad et al., 2022). The probability of noncanceled trials at each SSD, p(respond | stop-signal), served as a proxy for conflict because it is the outcome of the race between the conflicting processes and this was validated by simulations of the GO and STOP units in the interactive race model.

Previous work from our lab has observed neurons with these properties in the supplementary eye field but not the midcingulate cortex. Neurons in the supplementary eye field demonstrated pronounced, transient facilitation after successful response inhibition and were predominantly broad-spiking neurons found across all layers (Stuphorn et al., 2000; Sajad et al., 2022). Neurons in the midcingulate however show no apparent modulation of their firing rate after successful inhibition (Ito et al., 2003). We replicate these findings, but importantly do so using less biased methods and sampling from a larger area within the midcingulate cortex. As such, we conclude that single-unit activity in midcingulate cortex does not reflect response conflict in the stop-signal task.

6.4.4 Conclusion

Conflict has been well described in tasks that require over-riding an automatic or learnt response. This function has been localized to the midcingulate cortex in humans, but evidence in macaque neurophysiology fails to support this. In this study, we observe clear patterns of modulation across both dMCC and vMCC that distinguish between contexts in which a movement is generated or inhibited. Two clusters in our sample that demonstrate properties that allow them to reactively contribute to stopping, whilst two clusters distinguish between going and stopping even prior to the onset of the stop-signal. None of our identified clusters varied their activity proportional to the amount of conflict experienced. This work adds further evidence against the role of midcingulate cortex in conflict detection and monitoring, but describe several new signals that may contribute more broadly to other aspects of cognitive control.

CHAPTER 7

Discussion

7.1 Summary of Findings

In the introduction of this dissertation, I described how cognitive control is essential in ensuring we make the most appropriate action or decision in a situation. Part of this control entails stopping actions or decisions that may no longer be appropriate – this is referred to as response inhibition. Neural signals reflecting inhibition over motor plans have been observed in motor networks and are well explained through neurocognitive models. Whilst understanding these direct inhibition processes are essential, adaptive and flexible behavior also requires the continuous assessment and evaluation of our actions and environment. This monitoring allows us to learn the statistics of our environment, allowing for control to be employed proactively and ahead of anticipated periods of high cognitive demands. The medial frontal cortex has been proposed as an important region for monitoring action, supported by a range of evidence across modalities and species (Ridderinkhof et al., 2004; Rushworth et al., 2004; Schall et al., 2002; Passingham et al., 2010). However, activity within the medial frontal cortex is complex, and signals vary markedly with no apparent pattern or adherence to a single cognitive theory or computational model. Furthermore, it is unclear how different areas within the medial frontal cortex interact with one another to promote cognitive control. In this dissertation, I aimed to understand the contributions of the medial frontal cortex to cognitive control by examining several perspectives: from behavior to variation in EEG and LFP signals to differences in patterns of activity at the laminar and single-neuron level. Here, I discuss the key findings and consider their broader implications.

7.1.1 Establishing cognitive control over seemingly automatic behaviors

In chapter two, I showed that monkeys produced multi-modal distributions of saccade latencies during a saccade countermanding task. Whilst later modes peak around 300 ms, the earliest mode peaks before 100 ms and demonstrates little variability. This is consistent with earlier reports of express saccades that approach the lower limits imposed by conduction delays in the oculomotor system (Boch and Fischer, 1986; Boch et al., 1984; Fischer and Boch, 1983; Schiller et al., 2004). These observations are unusual for the countermanding task, where participants typically demonstrate slower response latencies than simple reaction time tasks (Logan and Cowan, 1984).

I demonstrated that this might result from two factors: temporal expectation and reward maximization. First, I considered whether monkeys might be predicting the time of a target and priming their oculomotor system accordingly. The time of the target cueing an eye movement occurred relative to when the monkey

initiated the trial by fixating on a central cue. This period is referred to as the foreperiod. Although we observed idiosyncrasies between monkeys, there was some evidence that monkeys responded faster when the foreperiod was predictable. However, this only influenced the latency of regular saccades; express saccades were unaffected. Second, I considered whether monkeys varied express saccade production to maximize reward opportunities. If a trial length is not fixed within the experiment, faster responses result in shorter trials. In turn this allows for more trials and thus more opportunities for reward in each session. Using simulations of response latencies, we demonstrated that producing a small proportion of express saccades could elevate the average reward rate under particular task conditions. We found that monkeys made a proportion of express saccades within this range, and not so many as to become detrimental to the reward rate. This behavior appeared to be learnt, developing over training, and extinguishing when no longer rewarded.

These observations suggest that higher-order representations of the task influence the mechanisms of express saccade generation. This includes temporal regularities and reward rates, which require integrating information over time. This may occur by higher cortical areas involved in performance monitoring (such as medial frontal cortex) modulating regions thought to generate express saccades (such as the superior colliculus). However, this requires further investigation.

7.1.2 Distinguishing the contribution of β-bursts to response inhibition and cognitive control

In chapters three and four, I build on recent findings in the human stopping literature that propose beta-bursts in EEG signals recorded over the frontal cortex directly contribute to response inhibition (Jana et al., 2020; Wessel, 2020). I first demonstrate that qualitatively similar activity patterns can be observed in EEG signals recorded over the medial frontal cortex of macaques performing a saccade countermanding task. Mirroring observations in humans, beta-bursts over the medial frontal cortex were more prevalent during the stopping period when a movement was inhibited compared to when a movement was executed. However, investigating the functional properties of these signals further, we reported bursts were infrequent, did not strongly differentiate between trials in which a movement was made or inhibited, and didn't reflect stopping behavior. These observations unfortunately raise doubts about the proposal that β-bursts are a viable mechanism of reactive response inhibition.

However, although this may accurately represent the association between beta-bursts and stopping, several alternative reasons have been proposed for this lack of causal efficacy. Most notably, the rarity of β-bursts has been argued to result from EEG recordings' low signal-to-noise ratio (SNR). I address these issues in chapter four, where I extend the work I present in chapter three to look at activity directly within the cortex. Recording directly from the cortex reduces any filtering of neural activity by the skull, muscles, and other tissue, which contribute to the low SNR in EEG. However, despite recording directly from the cortex, adjacent

to where the EEG signals are recorded, I found β-bursts were equally as infrequent as they were in EEG, and their properties did not account for the likelihood of canceling a planned response. As such, just like in EEG, β-bursts within the supplementary eye field cannot act as a causal mechanism for reactive response inhibition.

Interestingly, aside from response inhibition, we observed β-bursts in both EEG and LFP signals were most prevalent during other periods that contribute to cognitive control. Specifically, we observed an increase in bursts following the execution of an action – be it correctly stopping a movement or incorrectly generating a movement. These findings align with a wide range of literature in both humans and macaques, supporting the claim that the medial frontal cortex contributes to action and outcome monitoring (Bonini et al., 2014; Passingham et al., 2010; Rushworth et al., 2004; Schall et al., 2002; Ridderinkhof et al., 2004).

7.1.3 Cognitive control in the medial frontal cortex at the neuron level

In this dissertation, I presented evidence of cognitive control signals in single-units within two areas of the medial frontal cortex: supplementary eye field and midcingulate cortex. Whilst cognitive control signals have been well documented in both these areas, how they are generated is uncertain. This can be addressed by examining signals at laminar resolution: neurons in different cortical layers have different biophysical properties and anatomical connections. By describing patterns of activity across cortical layers and combining them with knowledge of the connectivity and cellular physiology of the region, these findings help constrain hypotheses about the microcircuitry of processes contributing to successful response inhibition.

In chapter five, I looked at activity within the supplementary eye field. Using semi-supervised clustering techniques, we observed three distinct classes of neurons that were modulated around the time of successful inhibition. We proposed alternative hypotheses about the functions of these neurons based on previous literature and theories of medial cortex function and combined this with a model comparison method to determine the function that best supported the variation in neural activity. This method labeled these three classes as conflict, goal maintenance, and event timing neurons. Conflict neurons were transiently facilitated after successful stopping and scaled their activity with the probability of responding at a given SSD – a measure previously described to reflect the amount of conflict between GO and STOP processes (Schall and Boucher, 2007; Botvinick et al., 2001). Goal Maintenance neurons demonstrated sustained facilitation, ending after a feedback tone signaling the termination of operant control on behavior. Finally, Event Timing neurons exhibited clear facilitation in the lead-up to inhibition but immediately suppressed their activity following successful inhibition. To understand how these signals may arise, we considered their distribution across cortical layers and their waveform properties. In contrast to Conflict and Event Timing neurons which were found in all layers and were a mix of narrow and broad spikes, Goal Maintenance neurons were mainly found in L2/3 and were primarily narrow spikes.

We then considered the relationship of this spiking activity to EEG signals recorded outside of the skull. 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 function. In chapter five, we also aim to link this single neuron activity to event-related potentials recorded in the EEG. Functionally, we found variation in the N2 component was best explained by the likelihood of an error at a given SSD, whilst variation in the P3 was best explained by variation in expectations about behavior. The N2 generally coincided with the peak recruitment of conflict and Event Timing neurons. When the discharge rate of the L2/3 neurons in these classes was higher, the N2 exhibited more negativity. The P3 coincided with the peak recruitment of Goal Maintenance neurons. When the discharge rate of the L2/3 neurons in this class was higher, the P3 exhibited more negativity.

In chapter six, I develop this work to examine cognitive control signals in the midcingulate cortex. Using the same semi-supervised clustering approach I used in chapter 5, we observed four distinct classes of neurons that differentiated their activity on trials in which a movement was generated or inhibited. In two populations of neurons, we observed patterns of modulation that could directly contribute to the reactive inhibition of a movement. These were similar to previous reports of fixation neurons in motor structures such as frontal eye field (Hanes et al., 1998), but were unexpected in midcingulate cortex. We also observed two classes of neurons that demonstrated clear differentiation between canceled and no-stop trials well before the presentation of a stop-signal. These signals were prolonged, lasting from target onset until after reward delivery. Preliminary interpretations label this activity as reflecting features of proactive control - signals that may bias attentional, perceptual, and motor systems towards going or stopping. As far as we are aware, this is the first neurophysiological evidence of such activity in midcingulate cortex. In contrast to observations in the supplementary eye field (described in chapter five), activity of neurons in the midcingulate cortex did not vary with p (respond | stop-signal) - a proxy for response conflict - challenging some theories of MCC function.

7.2 Interpretations & implications

7.2.1 Medial frontal cortex contributions to cognitive control

Cognitive control is especially essential when automatic or learned actions are detrimental to achieving a goal. In such situations, these actions must be stopped before execution; if not inhibited, these processes may be incorrectly executed outright or interfere with task-relevant behaviors. The contribution of the medial frontal cortex to response inhibition has been the subject of much debate. Whilst human studies emphasize that the pre-SMA is active during response inhibition and can directly contribute to the interruption of response preparation (Diesburg and Wessel, 2021; Rae et al., 2015; Sharp et al., 2010), neurophysiological evidence

obtained in the medial frontal cortex of macaques demonstrate that this area modulates too late to directly contribute to stopping (Scangos and Stuphorn, 2010; Stuphorn et al., 2010). Through this dissertation, I attempted to address this discrepancy more directly by applying the same methods and analyses in macaques that have been used to argue for reactive control in humans. Using this approach, I demonstrated that although we can observe similar activity in the macaque medial frontal cortex, this activity could not causally contribute to the inhibition of a planned movement. This discrepancy between humans and macaque can occur for a number of reasons, and are discussed later in this chapter.

Aside from response inhibition, effective cognitive control also requires continuous assessment and evaluation of our actions and environment. Evidence across various modalities and species has demonstrated activity within the medial frontal cortex in response to action outcomes, rewards, and value (Ridderinkhof et al., 2004; Rushworth et al., 2004; Schall et al., 2002; Passingham et al., 2010). Through this dissertation, I have shown that signals within the medial frontal cortex demonstrate clear patterns of modulation that occur after an action, i.e. after successful inhibition or the execution of an erroneous saccade. Most notably, in chapters four, five, and six, I described several signals within the medial frontal cortex that can contribute to effective and correct stopping, albeit not to the response inhibition process. This distinction is important: effective stopping requires the active and direct inhibition of a planned movement (response inhibition) and requires this inhibition to be held until we receive feedback that our action is correct. Linking back to my earlier discussion of inhibition in the introduction of this dissertation, we may consider that this period following action inhibition requires effective cognitive inhibition. Cognitive inhibition is proposed to prevent self-generated and irrelevant task goals or ideations into working memory (Nigg, 2000). Our studies found signals in SEF that were facilitated for a prolonged period after successful stopping. In our spiking data, we observed mainly narrow-spiking neurons that only became active once action inhibition was achieved and were sustained until the tone cueing that fixation at the central point was no longer required to obtain a juice reward. Similar observations were made in our LFP data, where beta-bursts were significantly more common during the stop period. These observations may also be considered in the context of the interactive race model 2.0 (Logan et al., 2015), in which inhibition is achieved at the neural level by either removing the drive to a movement unit (blocked-input) or by boosting input to a fixation unit (boosted-fixation). Here, we can imagine our observed neural activity contributing to either of these processes; narrow spikes are putative inhibitory neurons, and increased beta-band activity is thought to reflect inhibitory processes. With this, increased activity can reflect processes inhibiting a drive to the GO process or providing supplemental inhibition of processes that facilitate fixation. Future research must further distinguish these models and their underlying mechanisms within the brain, elucidating the role of prolonged and persistent activity and the relationship between beta-bursts and spiking activity.

7.2.2 Linking intracortical and EEG signals

Electroencephalogram (EEG) and magnetoencephalogram (MEG) are the most powerful methods for allowing us to non-invasively measure the electrophysiological dynamics of the brain. Whilst such dynamics offer a clear window into how the brain performs cognitive operations, these signals have limited utility until their neural generators are known (Cohen, 2014, 2017). Indeed, for EEG to have practical and clinical applications, we need to understand the function and anatomical properties of the circuits within the cortex that can generate the larger signals observed on the scalp. Through this dissertation, I attempted to uncover how these broad signals recorded over the scalp reflected features of intracortical signals.

Two prominent features of the EEG signal have been classically considered: oscillations and event-related potentials. Neural oscillations reflect the synchronous activity of a large population of neurons. They are well described in EEG and LFP signals and are observed across a range of nervous systems (Varela et al., 2001; Buzsáki et al., 2013). Over the past several decades, a range of cognitive, perceptual, and motor behaviors have been associated with specific changes in oscillatory activity (Buzsáki and Draguhn, 2004). This dissertation focused on the role of beta-band oscillations ($\sim$13–30 Hz) within the cortex and EEG which have long been considered in the context of movements (Jasper and Penfield, 1949). Classically, beta power decreases before a movement and rebounds following its execution. Whilst earlier work focused on trial-averaged changes in the power of these oscillations between conditions of interest, our understanding of beta oscillations has developed over the past few years. What appear to be gradual changes in beta-activity when averaged across trials are instead transient burst-like events at the individual trial level (Sherman et al., 2016; Shin et al., 2017)

In chapter four, I demonstrated no association between the observance of a β-burst in LFP and EEG. Interestingly, these signals showed similar functional patterns of activity and modulated similarly in both signals. We also observed that β-burst activity was not synchronized across a cortical column: a burst on one column showed no temporal or spatial relationship with bursts on channels other than those directly adjacent to it. This is unlikely to reflect multiple distinct bursts but instead reflect the same burst across multiple channels. With this assumption our observations can provide constraints on the spatial volume of a β-burst which can be implemented in biophysical models. Despite this, the lack of association between two similar and nearby signals raises questions about the biophysical link between them, which can be addressed through future research.

In contrast to oscillatory activity in these signals, we can also observe apparent variations in the voltages of the EEG signal in response to a particular event or stimuli (Luck, 2014). These are referred to as event-related potentials. Like oscillatory activity, these potentials reflect the summed activity of postsynaptic

potentials from a large population of similarly oriented neurons (Peterson et al., 1995). For consideration in this dissertation, previous studies have described an ERP complex referred to as the N2/P3 complex that was elicited on stop trials. The nature of this ERP is debated but is thought to reflect response inhibition processes (Kok et al., 2004; Ramautar et al., 2004; Waller et al., 2021). In this dissertation, I report the first observation of this component in the macaque model.

In chapter five, I investigated how signals in the supplementary eye field and the EEG recorded over the cortex reflected various performance monitoring processes contributing to successful inhibition. In contrast to my reported observations between beta activity in the cortex and EEG, I observed that the activity of neurons in L2/3, but not those in L5/6, which modulated during the peak of the N2 and P3 event-related potentials, significantly predicted the variation in the amplitude of the respective ERPs. However, it's important to note that whilst the populations of neurons that covaried with the ERPs consisted of less than 100 neurons, EEG reflects the activity of $\sim 10^6$ neurons – 10,000 times this amount. In conjunction with this proposition, the generation of a single spike is too brief and weak to create a notable contribution to the EEG signal. As such, it is improbable that these L2/3 neurons directly cause the variation in the ERPs; instead, it is more likely that the variation in spiking rate and ERP may reflect a common underlying mechanism or drive. This requires further investigation.

7.2.3 Translating findings between macaques and humans

Much of our knowledge about how the brain functions has come from noninvasive methods such as EEG, fMRI, and transcranial magnetic stimulation (TMS) in humans. These approaches have yielded useful insight into how neural systems respond to various cognitive processes but have fundamental constraints on the spatial and temporal resolution at which they can measure activity. These methods provide a window into how the brain performs particular functions but fall short of providing more mechanistic descriptions.

This can be addressed by using higher resolution techniques (such as single-unit recordings and microstimulation in neurophysiology) which allow for finer probing of the circuit. However, applying these techniques in humans carries several ethical issues and considerations, although they are possible in specific patient populations. Conversely, whilst these techniques can be used to rats and other animals, primates provide a useful translatable model of the brain given their similarities with humans (Passingham, 2009; Lear et al., 2022). Indeed, monkeys have been used as an effective model for studying the neural mechanisms of cognition for nearly 90 years (Jacobsen, 1936), with most of this work using rhesus macaques. Here, I consider the feasibility and extent to which factors of the work presented in this dissertation may be applied to human models.

We can assess the translatability of findings from our macaque model to humans from several perspec-

tives. First, we can consider the functional similarities. Behaviorally, stopping behavior is quantitatively similar between species – both demonstrate an increased probability of responding at later stop-signals, faster response latencies for errors, and performance based adjustments (Emeric et al., 2007). These observations allow for the application of the independent race model to provide an estimate of stop-signal reaction time (Logan and Cowan, 1984). As described in previous section, this model can appropriately capture observations across different effector types and different species (for example, see Logan and Irwin, 2000, and Emeric et al., 2007). We also see functional similarities between species in neural signals. As described above, the pattern of beta bursts we observed in our EEG data clearly mirrors those observed in humans (Wessel, 2020; Jana et al., 2020). We also observed similar event-related potentials to those observed in humans, providing the first reports of the N2/P3 component during stopping in macaques in chapter five; this component had similar properties to those described in human studies (Kok et al., 2004; Ramautar et al., 2004). Outside of the work presented within this dissertation, neurons in both regions in both species signal when errors are committed, whether positive or negative reinforcement is anticipated or received, or whether other forms of an outcome have been achieved (Ullsperger et al., 2014; Fu et al., 2019, 2022a).

We can also consider the anatomical similarities between species. Humans and macaques share many features of the medial frontal cortex that allow for greater cross-species translatability; this contrasts with other species such as rodents (Balsters et al., 2020, Schaeffer et al., 2020, but see Heukelum et al., 2020). First, considering the gross anatomy, we know these agranular areas in the frontal cortex seem largely homologous in organization and function between macaques and humans (Sallet et al., 2013). Second, the cytoarchitecture of these areas is also mostly similar between species (Heukelum et al., 2020; Vogt et al., 2005, 1995). One notable difference is that humans, but not macaques, have large spindle neurons in layer 5 of the cingulate cortex (Nimchinsky et al., 1995); it is unclear what functional differences this may elicit. Nevertheless, the mostly similar structure of the medial frontal cortex at the cellular level allows for translatability in our microcircuit and biophysical models of medial frontal cortex activity. Finally, in the context of cognitive control, monitoring and control signals in the medial frontal cortex must be broadcast to other areas. Using resting-state-fMRI, the medial frontal cortex in both species share common patterns of functional connectivity to other brain areas (Neubert et al., 2015). As such, the anatomy of the medial frontal cortex is highly similar between macaques and humans and is conducive for translating findings between species.

Despite these similarities between humans and macaques, some functional findings in the medial frontal cortex have highlighted clear distinctions between the species which must also be considered. These divergent observations may result from notable differences between these species. These are discussed in Fu et al. (2022b), Cole et al. (2009), and Schall and Emeric (2010), but briefly described here. First, whilst the gross anatomy of the medial frontal cortex is similar between the species, there are more nuanced distinctions in the

location of particular subregions and sulcal morphology (Sallet et al., 2013). In both humans and macaques, SMA is located immediately rostral to the primary motor cortex, and pre-SMA is located rostrally to SMA. In both species, pre-SMA is found on the medial surface. However, whilst much of SMA and all of SEF is situated on the dorsal convexity in macaques, in humans, they are located on the medial wall, with SEF centered on the paracentral sulcus. These differences in gross anatomy are not vital in understanding the brain-function relationship but may have important implications when considering the relationship between different signals. Such differences may be conducive to variation in recorded neural signals and must be considered when developing biophysical models. Second, other technical and conceptual differences exist in studying cognitive signals between humans and macaques. This includes variations in recording methodologies (i.e., microelectrodes in macaques, fMRI and EEG in humans) and training on the task. Unlike experiments in humans where participants can verbalize instructions and choices, experiments with monkeys must rely on sensory-motor paradigms in which performance is shaped through operant conditioning. In these tasks, participants express choices through an overt action, typically an eye, forelimb, or digit movement. Furthermore, perhaps most importantly, there are apparent differences in task design for studying stopping between the two species. Notably, whilst macaque models have established stopping in an oculomotor version of the stop-signal task, human stop-signal research has instead focused on inhibiting manual movements such as button presses. Although motor control systems may share some common systems (such as the basal ganglia pathways), the nature of the tested movements are drastically different in their complexity and the degree to which they are exposed to external influences. Whilst the eyes are set within the skull and have only a few degrees of freedom, arm movements occur away from the body (and thus can be acted upon by objects and forces in our environment) and have significantly more degrees of freedom. These features will undoubtedly result in significant variation in action monitoring and control networks which need to be considered when extrapolating findings between species.

7.2.4 Practical applications

Given the functional and anatomical similarity of the medial frontal cortex across species, we can then conjecture about potential practical applications of our findings. Here, I consider two applications: clinical and brain-machine interfaces.

7.2.4.1 Clinical applications

Response inhibition can be altered in various perturbed physiological states and pathological conditions. Studying how these disorders impact response inhibition can provide valuable insight into how these deficits can be addressed within a patient population and can also help us understand the neuroanatomical, neuro-

physiological, and neuropharmacological basis of response inhibition. Impairment in response inhibition has typically been assessed by comparing stop-signal reaction time across a range of disorders, including schizophrenia (Thakkar et al., 2011), obsessive-compulsive disorder (Mar et al., 2022), substance abuse disorders (Smith et al., 2014), and attention-deficit/hyperactivity disorder (Alderson et al., 2007; Lijffijt et al., 2005). Across these populations, patients show similar levels of slowing in SSRT. This similarity in impairments between these different psychiatric disorders warrants the consideration that they may share a common underlying neural mechanism.

Given the proposed role of the basal ganglia in several neural models of response inhibition, further insight can be gained from studying disorders of this system, such as Parkinson's disease (PD). Parkinson's disease is a neurodegenerative disorder characterized by dopaminergic and monoaminergic neuronal loss. This leads to motor impairments as the basal ganglia network becomes unbalanced, resulting in the STN becoming overactive, acting as a brake over movement. Relevant to some of the work presented in this dissertation, a recent study has shown that motor impairments in Parkinson's disease are associated with abnormal beta rhythms in the STN (Tinkhauser et al., 2017). In addition to motor impairments, Parkinson's disease patients have impaired response inhibition reflected by slower SSRTs. This impairment is independent of the severity of motor slowing (Gauggel et al., 2004). Dopaminergic therapies (such as Levodopa) and STN deep brain stimulation (STN-DBS) both improve response inhibition (Ray et al., 2009; Choudhury et al., 2019; van den Wildenberg et al., 2006). Furthermore, recent work has demonstrated estimates of SSRT can be used to optimize the parameters of STN-DBS, leading to significant clinical improvement of symptoms (Roy et al., 2020).

Although considered as one disorder, the clinical presentation of Parkinson's disease is heterogeneous and invites consideration of different subtypes (Sauerbier et al., 2016). Notably, patients can vary as to whether their primary presentation is described by more significant deficits in motor function or cognitive function. Interestingly, SSRT is similar among patients regardless of motor severity or whether they have a co-morbid impulse control disorder (Gauggel et al., 2004; Claassen et al., 2015). Furthermore, a recent study has detailed that reactive inhibition in early Parkinson's disease is impaired but proactive control processes are spared (Di Caprio et al., 2020), which may provide further evidence of a distinction between motor and cognitive processes in response inhibition. These findings present a clear difference between those patients with more significant motor impairment and those with greater cognitive impairment, which deserves additional consideration in future research. This may be achieved by considering differences between patients with different clinical presentations of the disorder; this may also be achieved by extending some of the work described about to also consider patients with early-stage Lewy-body dementia - which has similar pathology to Parkinson's disease but is characterized by greater cognitive impairment than motor (Gnanalingham et al.,

1997; O'Dowd et al., 2019).

7.2.4.2 Brain-computer interfaces

The neural control of movement has been of great interest to those developing brain-computer interfaces (BCI). Brain-computer interfaces acquire brain signals, process and analyze them, and translate them into commands relayed to output devices carrying out desired actions.

Our findings on the role of beta-band activity during stopping can be considered in the context of brain-computer interfaces. Whilst behavioral evidence has demonstrated that beta-bursts lack the functional properties to directly inhibit a response, a causal role for neural oscillations in behavior can only be attained by experimental manipulation (Herrmann et al., 2016; Vosskuhl et al., 2018). Although manipulations of neural activity can be achieved non-invasively by transcranial direct-current stimulation (tDCS) or transcranial magnetic stimulation (TMS), or invasively through microstimulation, these approaches cannot target band-specific activity. Brain-computer interfaces can be practical tools for making causal manipulations of oscillatory activity. This can be achieved by training participants to use BCIs to self-modulate brain signals through a process referred to as neurofeedback. Neurofeedback is a form of operant conditioning implemented through BCI in which oscillations are measured, processed, and translated in real-time into visual or auditory representations (Sitaram et al., 2017). Through learning, participants can then learn to self-manipulate aspects of these signals (Ros et al., 2014). These approaches have demonstrated potential as an alternative therapy for several disorders (Hammond, 2005; Begemann et al., 2016).

Linking back to the aforementioned clinical applications, neurofeedback approaches have been used to train patients with Parkinson's disease to regulate their STN beta-rhythm (Bichsel et al., 2021). After a single training session, patients learned to downregulate beta-activity, resulting in improved motor behaviors. Interestingly, patients could employ these strategies after the training session without real-time neurofeedback. Enz et al. (2022) used a similar approach in a cohort of healthy young adults to investigate the role of beta observed over the right front cortex in response inhibition. They found that training-related modulation of beta-activity was only achieved during the neurofeedback task but not during the stop-signal task. Furthermore, neurofeedback training of the right frontal beta activity did not change stopping behaviors. As such, BCIs may be effective as an alternative therapeutic technique in disorders characterized by abnormal beta-activity, but their lack of causal efficacy in healthy populations invites further consideration.

Methods such as BCI-neurofeedback are only effective when signals can be appropriately captured and are actually associated with observable changes in behavior. Current results indicate that beta-bursts may not be an effective target for neurofeedback interventions.

7.3 Open questions and future directions

Although addressing several important questions on the role of the medial frontal cortex in cognitive control, the research presented through this dissertation has produced many more to be answered.

Throughout this dissertation, I have demonstrated several neural signals that reflect the outcomes of our behavior (i.e., an error) and the state of the environment (i.e., the timing of key events, whether control is required, etc.). However, the neural mechanisms through which these signals elicit behavioral adaptations are untested, although we can form some speculations based on previously reported findings. Through other work in our lab, the baseline activity of movement neurons in the frontal eye field and superior colliculus have been shown to modulate in response to changes in speed-accuracy trade-off (Heitz and Schall, 2012; Reppert et al., 2018) - a form of proactive control. These observations were formalized through a neurocognitive model of speed-accuracy trade-off (Servant et al., 2019). Interestingly, other work from the countermanding task showed that the onset of accumulation of movement neurons in the superior colliculus is delayed following successful stopping, compared to when a saccade was correctly generated (Pouget et al., 2011). This finding suggests that proactive control influences the onset of accumulation or baseline activity of motor neurons, resulting in slower response latencies.

Current models of response inhibition primarily focus on how reactive stopping occurs but fail to account for such performance-based or contextual adjustments in behavior. Considering the findings above that outline how RT adjustments may occur, we can speculate how these can be applied to current cognitive models of the saccade countermanding task - specifically, the interactive race model and the spiking neural network model (Lo et al., 2009; Boucher et al., 2007; Logan et al., 2015). Building from the observation that proactive control is reflected through baseline changes in the speed-accuracy trade-off task, adaptation may result from fixation neurons exerting inhibition over movement neurons. Cognitively, this can be considered as a shift in the task goal as proposed by Verbruggen and Logan (2009b) – greater inhibition reflects a shift towards the STOP task, whilst less inhibition demonstrates a shift towards the GO task. However, if we build on findings from Pouget et al. (2011), RT adjustments result from delayed accumulation onset. This can be achieved through a gating mechanism (i.e., see Purcell et al., 2012, and Servant et al., 2019) in which the onset of motor plans is delayed until more evidence is gained for the correct response. In the context of ideas presented in the introduction, we can consider this as requiring more information to "decide that" our action is warranted before we "decide to" act (Schall, 2001).

Although we can combine these observations with findings from previous research to generate ideas of how behaviors are adapted, these associations need to be more directly tested. Future work can achieve this through a multi-electrode approach in which neural activity is simultaneously recorded across several key

regions in the system rather than from just one. Such techniques have been useful in understanding how learning signals are conveyed across a fronto-striatal network (Voloh et al., 2020). In the context of this work, a future study may opt to simultaneously sample activity in the supplementary eye field and frontal eye field. Once we establish evidence across systems as described above, this needs to be formalized through neurocomputational models of executive control. These models will have to also be reconciled with current theories of medial frontal cortex function (Shenhav et al., 2016a; Kolling et al., 2016).

Finally, perhaps most curiously, several findings from this study draw comparisons to complementary findings from ascending neuromodulatory systems. In chapter two, I described that monkeys elicited express saccades at an incidence proportional to the average reward rate of the session. Although activity in the medial frontal cortex is thought to reflect outcomes and rewards on a trial-by-trial basis, recent work has described activity in the DRN of macaques that represents the state of the environment over longer time scales, including the average value (Wittmann et al., 2020; Khalighinejad et al., 2021) or average uncertainty (Grossman and Cohen, 2022) of reward. In chapter five, I described a class of neurons in the supplementary eye field that demonstrates transient activity scaling with the amount of experienced conflict. These observations mirrored, with surprising similarity, recent findings of dopaminergic neurons in the dorsolateral aspect of the substantia nigra pars compacta (Ogasawara et al., 2018) which were interpreted as reflecting salience or surprise; this is discussed further in chapter five. This observation can fit into a broader range of work that highlights the role of dopamine in stopping. Previous work demonstrated that reduced blink rate, an index of dopaminergic function, is associated with poorer inhibition - indexed by slower SSRTs (Colzato et al., 2009). This is also conveyed through observations that clinical populations characterized by dopaminergic dysfunction also have slower SSRTs (Colzato et al., 2009; Enticott et al., 2008). Collectively, these findings invite future work to consider how contributions from ascending neuromodulatory systems, such as the dopaminergic (DA) midbrain and the serotonergic (5HT) brainstem, may influence or are influenced, by higher-cortical areas such as the medial frontal cortex. This may be achieved by recording and stimulating these structures and observing responses during the task or by studying changes in stopping behaviors and neural activity following pharmacological manipulations.

7.4 Conclusion

In our interactions with the world, we are constantly monitoring the state of our environment and using this information to guide decisions about our actions. This control grants us flexibility, allowing us to take the most appropriate action in a given context, hopefully leading to the outcome we want. Through this dissertation, I combined neurophysiological methods across temporal and spatial scales with established cognitive models to understand how the medial frontal cortex contributes to cognitive control. In part, this work has added to

a growing body of literature supporting that the medial frontal cortex represents the outcomes of our actions (such as conflict and errors). However, this work has also provided significant developments to our existing knowledge. Importantly, this work has developed our understanding of how cognitive control signals in the cortex are related to those observed outside the cortex. Such knowledge is vital for advancing the practical and clinical application of EEG and other neural signals. Furthermore, this dissertation has described several new signals extending our understanding of how response inhibition occurs; these findings invite us to think more broadly and consider response inhibition as more than just a process that stops a physical action but, instead, as a series of processes that support the need to stop. In closing, the findings presented in this dissertation have advanced our understanding of the mechanisms through which cognitive control is represented in the medial frontal cortex and has provided new ideas and directions for consideration in future research.

References

Afsharpour, S. (1985). Topographical projections of the cerebral cortex to the subthalamic nucleus. *Journal of Comparative Neurology*, 236(1):14–28.

Alderson, R. M., Rapport, M. D., and Kofler, M. J. (2007). Attention-deficit/hyperactivity disorder and behavioral inhibition: a meta-analytic review of the stop-signal paradigm. *Journal of Abnormal Child Psychology*, 35(5):745–758.

Alexander, W. H. and Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, 14(10):1338–1344.

Amarante, L. M., Caetano, M. S., and Laubach, M. (2017). Medial frontal theta is entrained to rewarded actions. *Journal of Neuroscience*, 37(44):10757–10769.

Amatya, N., Gong, Q., and Knox, P. C. (2011). Differing proportions of 'express saccade makers' in different human populations. *Experimental Brain Research*, 210(1):117–129.

Ameqrane, I., Pouget, P., Wattiez, N., Carpenter, R., and Missal, M. (2014). Implicit and explicit timing in oculomotor control. *PLoS ONE*, 9(4):e93958.

Amiez, C., Joseph, J.-P., and Procyk, E. (2006). Reward encoding in the monkey anterior cingulate cortex. *Cerebral cortex*, 16(7):1040–1055.

Amiez, C. and Petrides, M. (2009). Anatomical organization of the eye fields in the human and non-human primate frontal cortex. *Progress in Neurobiology*, 89(2):220–230.

Andersen, R. A. and Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, 25:189–220.

Anderson, M. C. and Levy, B. J. (2009). Suppressing unwanted memories. *Current Directions in Psychological Science*, 18(4):189–194.

Anderson, M. E. and Horak, F. B. (1985). Influence of the globus pallidus on arm movements in monkeys. III. Timing of movement-related information. *Journal of Neurophysiology*, 54(2):433–48.

Aron, A. R. (2007). The neural basis of inhibition in cognitive control. *Neuroscientist*, 13(3):214–28.

Aron, A. R. (2011). From Reactive to Proactive and Selective Control: Developing a Richer Model for Stopping Inappropriate Responses. *Biological Psychiatry*, 69(12):e55–e68.

Aron, A. R., Behrens, T. E., Smith, S., Frank, M. J., and Poldrack, R. A. (2007). Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *Journal of Neuroscience*, 27(14):3743–3752.

Aron, A. R., Herz, D. M., Brown, P., Forstmann, B. U., and Zaghloul, K. (2016). Frontosubthalamic circuits for control of action and cognition. *Journal of Neuroscience*, 36(45):11489–11495.

Aron, A. R. and Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience*, 26(9):2424–2433.

Aron, A. R., Robbins, T. W., and Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: one decade on. *Trends in Cognitive Science*, 18(4):177–85.

Arrington, C. M. and Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological Science*, 15(9):610–5.

Asrress, K. N. and Carpenter, R. H. (2001). Saccadic countermanding: a comparison of central and peripheral stop signals. *Vision research*, 41(20):2645–2651.

Atsma, J., Maij, F., Corneil, B. D., and Medendorp, W. P. (2014). No perisaccadic mislocalization with abruptly cancelled saccades. *Journal of Neuroscience*, 34(16):5497–504.

Baddeley, A. (1996). Exploring the central executive. *The Quarterly Journal of Experimental Psychology Section A*, 49(1):5–28.

Badry, R., Mima, T., Aso, T., Nakatsuka, M., Abe, M., Fathi, D., Foly, N., Nagiub, H., Nagamine, T., and Fukuyama, H. (2009). Suppression of human cortico-motoneuronal excitability during the Stop-signal task. *Clinical Neurophysiology*, 120(9):1717–23.

Bahill, A. T., Clark, M. R., and Stark, L. (1975). The main sequence, a tool for studying human eye movements. *Mathematical Biosciences*, 24(3-4):191–204.

Balsters, J. H., Zerbi, V., Sallet, J., Wenderoth, N., and Mars, R. B. (2020). Primate homologs of mouse cortico-striatal circuits. *eLife*, 9:e53680.

Band, G. P., Molen, M. W. v. d., and Logan, G. D. (2003). Horse-race model simulations of the stop-signal procedure. *Acta Psychologica*, 112(2):105–142.

Barbas, H. and Pandya, D. (1987). Architecture and frontal cortical connections of the premotor cortex (area 6) in the rhesus monkey. *Journal of Comparative Neurology*, 256(2):211–228.

Bari, A. and Robbins, T. W. (2013). Inhibition and impulsivity: Behavioral and neural basis of response control. *Progress in Neurobiology*, 108:44–79.

Basso, M. A. and Sommer, M. A. (2011). Exploring the role of the substantia nigra pars reticulata in eye movements. *Neuroscience*, 198:205–12.

Basso, M. A. and Wurtz, R. H. (2002). Neuronal activity in substantia nigra pars reticulata during target selection. *Journal of Neuroscience*, 22(5):1883–94.

Bastin, J., Polosan, M., Benis, D., Goetz, L., Bhattacharjee, M., Piallat, B., Krainik, A., Bougerol, T., Chabardes, S., and David, O. (2014). Inhibitory control and error monitoring by human subthalamic neurons. *Translational Psychiatry*, 4(9):e439.

Bastos, A. M., Loonis, R., Kornblith, S., Lundqvist, M., and Miller, E. K. (2018). Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. *Proceedings of the National Academy of Sciences*, 115(5):1117–1122.

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711.

Battaglia-Mayer, A., Buiatti, T., Caminiti, R., Ferraina, S., Lacquaniti, F., and Shallice, T. (2014). Correction and suppression of reaching movements in the cerebral cortex: physiological and neuropsychological aspects. *Neuroscience Biobehavioral Reviews*, 42:232–51.

Begemann, M. J., Florisse, E. J., van Lutterveld, R., Kooyman, M., and Sommer, I. E. (2016). Efficacy of EEG neurofeedback in psychiatry: A comprehensive overview and meta-analysis. *Translational Brain Rhythmicity*, 1(1):19–29.

Belin, T. R. and Rubin, D. B. (1995). The analysis of repeated-measures data on schizophrenic reaction times using mixture models. *Statistics in Medicine*, 14(8):747–768.

Bell, A., Meredith, M., Van Opstal, A., and Munoz, D. (2006). Stimulus intensity modifies saccadic reaction time and visual response latency in the superior colliculus. *Experimental Brain Research*, 174(1):53–59.

Bender, A. D., Filmer, H. L., Garner, K., Naughtin, C. K., and Dux, P. E. (2016). On the relationship between response selection and response inhibition: An individual differences approach. *Attention, Perception, & Psychophysics*, 78(8):2420–2432.

Benevento, L. A. and Fallon, J. H. (1975). The ascending projections of the superior colliculus in the rhesus monkey (Macaca mulatta). *Journal of Comparative Neurology*, 160(3):339–61.

Benis, D., David, O., Piallat, B., Kibleur, A., Goetz, L., Bhattacharjee, M., Fraix, V., Seigneuret, E., Krack, P., Chabardes, S., and Bastin, J. (2016). Response inhibition rapidly increases single-neuron responses in the subthalamic nucleus of patients with Parkinson's disease. *Cortex*, 84:111–123.

Beul, S. F. and Hilgetag, C. C. (2014). Towards a "canonical" agranular cortical microcircuit. *Frontiers in Neuroanatomy*, 8:165.

Beyer, F., Münte, T. F., Fischer, J., and Krämer, U. M. (2012). Neural aftereffects of errors in a stop-signal task. *Neuropsychologia*, 50(14):3304–3312.

Bibi, R. and Edelman, J. A. (2009). The influence of motor training on human express saccade production. *Journal of Neurophysiology*, 102(6):3101–3110.

Bichot, N. P., Schall, J. D., and Thompson, K. G. (1996). Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature*, 381(6584):697–699.

Bichsel, O., Stieglitz, L. H., Oertel, M. F., Baumann, C. R., Gassert, R., and Imbach, L. L. (2021). Deep brain electrical neurofeedback allows parkinson patients to control pathological oscillations and quicken movements. *Scientific Reports*, 11(1):1–10.

Bimbi, M., Festante, F., Coudé, G., Vanderwert, R. E., Fox, N. A., and Ferrari, P. F. (2018). Simultaneous scalp recorded EEG and local field potentials from monkey ventral premotor cortex during action observation and execution reveals the contribution of mirror and motor neurons to the mu-rhythm. *NeuroImage*, 175:22–31.

Blanchard, T. C. and Hayden, B. Y. (2014). Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *Journal of Neuroscience*, 34(2):646–655.

Boch, R. and Fischer, B. (1986). Further observations on the occurrence of express-saccades in the monkey. *Experimental Brain Research*, 63(3):487–494.

Boch, R., Fischer, B., and Ramsperger, E. (1984). Express-saccades of the monkey: Reaction times versus intensity, size, duration, and eccentricity of their targets. *Experimental Brain Research*, 55(2):223–231.

Bohlen, M. O., Warren, S., and May, P. J. (2016). A central mesencephalic reticular formation projection to the supraoculomotor area in macaque monkeys. *Brain Structure and Function*, 221(4):2209–29.

Bompas, A., Campbell, A. E., and Sumner, P. (2020). Cognitive control and automatic interference in mind and brain: A unified model of saccadic inhibition and countermanding. *Psychological Review*, 127(4):524–561.

Bonini, F., Burle, B., Liégeois-Chauvel, C., Régis, J., Chauvel, P., and Vidal, F. (2014). Action monitoring and medial frontal cortex: Leading role of supplementary motor area. *Science*, 343(6173):888–891.

Bonnevie, T. and Zaghloul, K. A. (2019). The subthalamic nucleus: unravelling new roles and mechanisms in the control of action. *The Neuroscientist*, 25(1):48–64.

Born, S., Mottet, I., and Kerzel, D. (2014). Presaccadic perceptual facilitation effects depend on saccade execution: Evidence from the stop-signal paradigm. *Journal of Vision*, 14(3):7–7.

Botvinick, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cognitive Affective Behavioral Neuroscience*, 7(4):356–66.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., and Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3):624–52.

Boucher, L., Palmeri, T. J., Logan, G. D., and Schall, J. D. (2007). Inhibitory Control in Mind and Brain: An Interactive Race Model of Countermanding Saccades. *Psychological Review*, 114(2):376–397.

Brass, M. and Haggard, P. (2008). The what, when, whether model of intentional action. *Neuroscientist*, 14(4):319–25.

Braver, T. S. (2012). The variable nature of cognitive control: a dual mechanisms framework. *Trends in Cognitive Sciences*, 16(2):106–113.

Braver, T. S., Gray, J. R., and Burgess, G. C. (2007). Explaining the many varieties of working memory variation: Dual mechanisms of cognitive control. *Variation in working memory*, 75:106.

Breese, B. B. (1899). On inhibition. *The Psychological Review: Monograph Supplements*, 3(1):i–65.

Brinkman, C. and Porter, R. (1979). Supplementary motor area in the monkey: activity of neurons during performance of a learned motor task. *Journal of Neurophysiology*, 42(3):681–709.

Brockett, A. T., Hricz, N. W., Tennyson, S. S., Bryden, D. W., and Roesch, M. R. (2020). Neural signals in red nucleus during reactive and proactive adjustments in behavior. *Journal of Neuroscience*, 40(24):4715–4726.

Bromberg-Martin, E. S., Matsumoto, M., and Hikosaka, O. (2010). Dopamine in motivational control: Rewarding, aversive, and alerting. *Neuron*, 68(5):815–34.

Brotchie, P., Iansek, R., and Horne, M. K. (1991). Motor function of the monkey globus pallidus. 1. Neuronal discharge and parameters of movement. *Brain*, 114 (Pt 4):1667–83.

Brown, J. W. and Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, 307(5712):1118–1121.

Brown, J. W., Hanes, D. P., Schall, J. D., and Stuphorn, V. (2008). Relation of frontal eye field activity to saccade initiation during a countermanding task. *Experimental Brain Research*, 190(2):135–51.

Brown, S. D. and Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cognitive Psychology*, 57(3):153–78.

Bruce, C. J. and Goldberg, M. E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *Journal of Neurophysiology*, 53(3):603–635.

Bruce, C. J., Goldberg, M. E., Bushnell, M. C., and Stanton, G. B. (1985). Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *Journal of Neurophysiology*, 54(3):714–734.

Brunamonti, E., Ferraina, S., and Paré, M. (2012). Controlled movement processing: evidence for a common inhibitory control of finger, wrist, and arm movements. *Neuroscience*, 215:69–78.

Buhusi, C. V. and Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience*, 6(10):755–65.

Burman, D. D. and Bruce, C. J. (1997). Suppression of task-related saccades by electrical stimulation in the primate's frontal eye field. *Journal of Neurophysiology*, 77(5):2252–2267.

Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420.

Buzsáki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929.

Buzsáki, G., Logothetis, N., and Singer, W. (2013). Scaling brain size, keeping timing: evolutionary preservation of brain rhythms. *Neuron*, 80(3):751–764.

Caan, W., Perrett, D. I., and Rolls, E. T. (1984). Responses of striatal neurons in the behaving monkey. 2. Visual processing in the caudal neostriatum. *Brain Research*, 290(1):53–65.

Cabel, D. W., Armstrong, I. T., Reingold, E., and Munoz, D. P. (2000). Control of saccade initiation in a countermanding task using visual and auditory stop signals. *Experimental Brain Research*, 133(4):431–41.

Cai, X. and Padoa-Schioppa, C. (2012). Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *Journal of Neuroscience*, 32(11):3791–3808.

Camalier, C. R., Gotler, A., Murthy, A., Thompson, K. G., Logan, G. D., Palmeri, T. J., and Schall, J. D. (2007). Dynamics of saccade target selection: race model analysis of double step and search step saccade production in human and macaque. *Vision Research*, 47(16):2187–211.

Carpenter, R. H. S. (1988). *Movements of the eyes, 2nd rev. enlarged ed*. Pion Limited, London, England.

Carter, C. S., Mintun, M., and Cohen, J. D. (1995). Interference and facilitation effects during selective attention: an H215O PET study of Stroop task performance. *NeuroImage*, 2(4):264–272.

Castellar, E. n., Kühn, S., Fias, W., and Notebaert, W. (2010). Outcome expectancy and not accuracy determines posterror slowing: ERP support. *Cognitive, Affective, & Behavioral Neuroscience*, 10(2):270–278.

Castiglione, A., Wagner, J., Anderson, M., and Aron, A. R. (2019). Preventing a thought from coming to mind elicits increased right frontal beta just as stopping action does. *Cerebral Cortex*, 29(5):2160–2172.

Cavada, C. and Goldman-Rakic, P. S. (1989a). Posterior parietal cortex in rhesus monkey: I. Parcellation of areas based on distinctive limbic and sensory corticocortical connections. *Journal of Comparative Neurology*, 287(4):393–421.

Cavada, C. and Goldman-Rakic, P. S. (1989b). Posterior parietal cortex in rhesus monkey: II. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *Journal of Comparative Neurology*, 287(4):422–45.

Cavanagh, J. F., Cohen, M. X., and Allen, J. J. B. (2009). Prelude to and Resolution of an Error: EEG Phase Synchrony Reveals Cognitive Control Dynamics during Action Monitoring. *Journal of Neuroscience*, 29(1):98–105.

Cavanagh, J. F. and Shackman, A. J. (2015). Frontal midline theta reflects anxiety and cognitive control: Meta-analytic evidence. *Journal of Physiology-Paris*, 109(1-3):3–15.

Cerkevich, C. M., Lyon, D. C., Balaram, P., and Kaas, J. H. (2014). Distribution of cortical neurons projecting to the superior colliculus in macaque monkeys. *Eye and Brain*, 6(Suppl 1):121.

Chang, S. W., Dickinson, A. R., and Snyder, L. H. (2008). Limb-specific representation for reaching in the posterior parietal cortex. *Journal of Neuroscience*, 28(24):6128–40.

Chen, L. L. and Wise, S. P. (1995). Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *Journal of Neurophysiology*, 73(3):1101–21.

Chen, W., Hemptinne, C. d., Miller, A. M., Leibbrand, M., Little, S. J., Lim, D. A., Larson, P. S., and Starr, P. A. (2020). Prefrontal-subthalamic hyperdirect pathway modulates movement inhibition in humans. *Neuron*, 106(4):579–588 e3.

Cheney, P. D. and Fetz, E. E. (1980). Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of Neurophysiology*, 44(4):773–91.

Choudhury, S., Roy, A., Mondal, B., Singh, R., Halder, S., Chatterjee, K., Baker, M. R., Kumar, H., and Baker, S. N. (2019). Slowed movement stopping in parkinson's disease and focal dystonia is improved by standard treatment. *Scientific Reports*, 9(1):1–9.

Claassen, D. O., van den Wildenberg, W. P., Harrison, M. B., van Wouwe, N. C., Kanoff, K., Neimat, J. S., and Wylie, S. A. (2015). Proficient motor impulse control in parkinson disease patients with impulsive and compulsive behaviors. *Pharmacology Biochemistry and Behavior*, 129:19–25.

Clark, J. M. (1996). Contributions of inhibitory mechanisms to unified theory in neuroscience and psychology. *Brain and Cognition*, 30(1):127–152.

Cohen, M. X. (2014). A neural microcircuit for cognitive conflict detection and signaling. *Trends in Neurosciences*, 37(9):480–490.

Cohen, M. X. (2017). Where Does EEG Come From and What Does It Mean? *Trends in Neurosciences*, 40(4):208–218.

Cohen, M. X., Ridderinkhof, K. R., Haupt, S., Elger, C. E., and Fell, J. (2008). Medial frontal cortex and response conflict: Evidence from human intracranial EEG and medial frontal cortex lesion. *Brain Research*, 1238:127–142.

Cole, M. W., Yeung, N., Freiwald, W. A., and Botvinick, M. (2009). Cingulate cortex: diverging data from humans and monkeys. *Trends in Neuroscience*, 32(11):566–74.

Colonius, H. (1990). A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review*, 97(2):309.

Colonius, H., Özyurt, J., and Arndt, P. A. (2001). Countermanding saccades with auditory stop signals: testing the race model. *Vision Research*, 41(15):1951–1968.

Colzato, L. S., van den Wildenberg, W. P., van Wouwe, N. C., Pannebakker, M. M., and Hommel, B. (2009). Dopamine and inhibitory action control: evidence from spontaneous eye blink rates. *Experimental brain research*, 196(3):467–474.

Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., and Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *Journal of Neuroscience*, 11(8):2383–2402.

Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3):201–15.

Corneil, B. D. and Elsley, J. K. (2005). Countermanding eye-head gaze shifts in humans: marching orders are delivered to the head first. *Journal of Neurophysiology*, 94(1):883–895.

Correa, and Nobre, A. C. (2008). Neural modulation by regularity and passage of time. *Journal of Neurophysiology*, 100(3):1649–1655.

Cosman, J. D., Lowe, K. A., Zinke, W., Woodman, G. F., and Schall, J. D. (2018). Prefrontal control of visual distraction. *Current Biology*, 28(3):414–420.

Coull, J. and Nobre, A. (2008). Dissociating explicit timing from temporal expectation with fMRI. *Current Opinion in Neurobiology*, 18(2):137–44.

Coxon, J. P., Stinear, C. M., and Byblow, W. D. (2007). Selective inhibition of movement. *Journal of Neurophysiology*, 97(3):2480–9.

Crammond, D. J. and Kalaska, J. F. (2000). Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *Journal of Neurophysiology*, 84(2):986–1005.

Crutcher, M. D. and Alexander, G. E. (1990). Movement-related neuronal activity selectively coding either direction or muscle pattern in three motor areas of the monkey. *Journal of Neurophysiology*, 64(1):151–63.

Crutcher, M. D. and DeLong, M. R. (1984). Single cell studies of the primate putamen. II. Relations to direction of movement and pattern of muscular activity. *Experimental Brain Research*, 53(2):244–58.

Dash, S., Peel, T. R., Lomber, S. G., and Corneil, B. D. (2020). Impairment but not abolishment of express saccades after unilateral or bilateral inactivation of the frontal eye fields. *Journal of Neurophysiology*, 123(5):1907–1919.

De Jong, R., Coles, M. G., Logan, G. D., and Gratton, G. (1990). In search of the point of no return: the control of response processes. *Journal of Experimental Psychology: Human Perception and Performance*, 16(1):164.

Deecke, L. and Kornhuber, H. H. (1978). An electrical sign of participation of the mesial 'supplementary' motor cortex in human voluntary finger movement. *Brain Research*, 159(2):473–6.

Deffains, M., Legallet, E., and Apicella, P. (2010). Modulation of neuronal activity in the monkey putamen associated with changes in the habitual order of sequential movements. *Journal of Neurophysiology*, 104(3):1355–1369.

Deiber, M. P., Honda, M., Ibanez, V., Sadato, N., and Hallett, M. (1999). Mesial motor areas in self-initiated versus externally triggered movements examined with fMRI: effect of movement type and rate. *Journal of Neurophysiology*, 81(6):3065–77.

DeLong, M. R. (1973). Putamen: activity of single units during slow and rapid arm movements. *Science*, 179(4079):1240–2.

DeLong, M. R., Crutcher, M. D., and Georgopoulos, A. P. (1985). Primate globus pallidus and subthalamic nucleus: functional organization. *Journal of Neurophysiology*, 53(2):530–43.

Dennett, D. C. (2015). *Elbow Room: The Varieties of Free Will Worth Wanting*. The MIT Press, new edition.

Di Caprio, V., Modugno, N., Mancini, C., Olivola, E., and Mirabella, G. (2020). Early-stage parkinson's patients show selective impairment in reactive but not proactive inhibition. *Movement Disorders*, 35(3):409–418.

Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, 64(1):135–168.

Diamond, A. and Gilbert, J. (1989). Development as progressive inhibitory control of action: Retrieval of a contiguous object. *Cognitive Development*, 4(3):223–249.

Dias, E. C., Kiesau, M., and Segraves, M. A. (1995). Acute activation and inactivation of macaque frontal eye field with GABA-related drugs. *Journal of Neurophysiology*, 74(6):2744–2748.

Dias, E. C. and Segraves, M. A. (1999). Muscimol-induced inactivation of monkey frontal eye field: effects on visually and memory-guided saccades. *Journal of Neurophysiology*, 81(5):2191–2214.

Diesburg, D. A., Greenlee, J. D., and Wessel, J. R. (2021). Cortico-subcortical burst dynamics underlying movement cancellation in humans. *eLife*, 10:e70270.

Diesburg, D. A. and Wessel, J. R. (2021). The Pause-then-Cancel model of human action-stopping: Theoretical considerations and empirical evidence. *Neuroscience Biobehavioral Reviews*, 129:17–34.

Dimoska, A. and Johnstone, S. J. (2008). Effects of varying stop-signal probability on erps in the stop-signal task: do they reflect variations in inhibitory processing or simply novelty effects? *Biological Psychology*, 77(3):324–336.

Ding, L. and Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cerebral cortex*, 22(5):1052–1067.

Donahue, C., Seo, H., and Lee, D. (2013). Cortical signals for rewarded actions and strategic exploration. *Neuron*, 80(1):223–234.

Donders, F. C. (1969). On the speed of mental processes. *Acta Psychologica*, 30:412–431.

Douglas, R. J., Martin, K. A., and Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural computation*, 1(4):480–488.

Drazin, D. H. (1961). Effects of foreperiod, foreperiod variability, and probability of stimulus occurrence on simple reaction time. *Journal of Experimental Psychology*, 62(1):43–50.

Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F., and Ostojic, S. (2021). The role of population structure in computations through neural dynamics. *bioRxiv*, page 2020.07.03.185942.

Dum, R. P. and Strick, P. L. (1991). The origin of corticospinal projections from the premotor areas in the frontal lobe. *Journal of Neuroscience*, 11(3):667–89.

Dum, R. P. and Strick, P. L. (1996). Spinal cord terminations of the medial wall motor areas in macaque monkeys. *Journal of Neuroscience*, 16(20):6513–25.

Duncan, H., Johnson, R., Swales, M., and Freer, C. (1997). Frontal lobe deficits after head injury: Unity and diversity of function. *Cognitive Neuropsychology*, 14(5):713–741.

Dutra, I. C., Waller, D. A., and Wessel, J. R. (2018). Perceptual surprise improves action stopping by nonselectively suppressing motor activity via a neural mechanism for motor inhibition. *Journal of Neuroscience*, 38(6):1482–1492.

Eagle, D. M., Bari, A., and Robbins, T. W. (2008). The neuropsychopharmacology of action inhibition: cross-species translation of the stop-signal and go/no-go tasks. *Psychopharmacology*, 199(3):439–456.

Ebbesen, C. L. and Brecht, M. (2017). Motor cortex - to act or not to act? *Nature Reviews Neuroscience*, 18(11):694–705.

Ebitz, R. B. and Platt, M. L. (2015). Neuronal activity in primate dorsal anterior cingulate cortex signals task conflict and predicts adjustments in pupil-linked arousal. *Neuron*, 85(3):628–640.

Egger, S. W., Le, N. M., and Jazayeri, M. (2020). A neural circuit model for human sensorimotor timing. *Nat Commun*, 11(1):3933.

Egner, T. and Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience*, 8(12):1784–1790.

Elias, S., Joshua, M., Goldberg, J. A., Heimer, G., Arkadir, D., Morris, G., and Bergman, H. (2007). Statistical properties of pauses of the high-frequency discharge neurons in the external segment of the globus pallidus. *Journal of Neuroscience*, 27(10):2525–2538.

Emeric, E. E., Brown, J. W., Boucher, L., Carpenter, R. H., Hanes, D. P., Harris, R., Logan, G. D., Mashru, R. N., Paré, M., Pouget, P., Stuphorn, V., Taylor, T. L., and Schall, J. D. (2007). Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vision Research*, 47(1):35–49.

Emeric, E. E., Brown, J. W., Leslie, M., Pouget, P., Stuphorn, V., and Schall, J. D. (2008). Performance monitoring local field potentials in the medial frontal cortex of primates: Anterior cingulate cortex. *Journal of Neurophysiology*, 99(2):759–772.

Emeric, E. E., Leslie, M., Pouget, P., and Schall, J. D. (2010). Performance monitoring local field potentials in the medial frontal cortex of primates: supplementary eye field. *Journal of Neurophysiology*, 104(3):1523–37.

Emmons, E., Tunes-Chiuffa, G., Choi, J., Bruce, R. A., Weber, M. A., Kim, Y., and Narayanan, N. S. (2020). Temporal learning among prefrontal and striatal ensembles. *Cerebral Cortex Communications*, 1(1):tgaa058.

Enriquez-Geppert, S., Konrad, C., Pantev, C., and Huster, R. J. (2010). Conflict and inhibition differentially affect the N200/P300 complex in a combined go/nogo and stop-signal task. *NeuroImage*, 51(2):877–87.

Enticott, P. G., Ogloff, J. R., and Bradshaw, J. L. (2008). Response inhibition and impulsivity in schizophrenia. *Psychiatry Research*, 157(1-3):251–254.

Enz, N., Ruddy, K. L., Rueda-Delgado, L. M., and Whelan, R. (2021). Volume of -Bursts, But Not Their Rate, Predicts Successful Response Inhibition. *Journal of Neuroscience*, 41(23):5069–5079.

Enz, N., Schmidt, J., Nolan, K., Mitchell, M., Alvarez Gomez, S., Alkayyali, M., Cambay, P., Gippert, M., Whelan, R., and Ruddy, K. (2022). Self-regulation of the brain's right frontal beta rhythm using a brain-computer interface. *Psychophysiology*, page e14115.

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

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

Eskandar, E. N. and Assad, J. A. (1999). Dissociation of visual, motor and predictive signals in parietal cortex during visual guidance. *Nature Neuroscience*, 2(1):88–93.

Feng, S., Holmes, P., Rorie, A., and Newsome, W. T. (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Computational Biology*, 5(2):e1000284.

Ferrera, V. P., Yanike, M., and Cassanello, C. (2009). Frontal eye field neurons signal changes in decision criteria. *Nature Neuroscience*, 12(11):1458–1462.

Filevich, E., Kuhn, S., and Haggard, P. (2012). Negative motor phenomena in cortical stimulation: implications for inhibitory control of human action. *Cortex*, 48(10):1251–61.

Findlay, J. M. (1981). Spatial and temporal factors in the predictive generation of saccadic eye movements. *Vision Research*, 21(3):347–354.

Fiorillo, C. D., Newsome, W. T., and Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. *Nature Neuroscience*, 11(8):966–73.

Fischer, A. G., Danielmeier, C., Villringer, A., Klein, T. A., and Ullsperger, M. (2016). Gender influences on brain responses to errors and post-error adjustments. *Scientific Reports*, 6(1):24435.

Fischer, B. and Boch, R. (1983). Saccadic eye movements after extremely short reaction times in the monkey. *Brain Research*, 260(1):21–26.

Fischer, B., Boch, R., and Ramsperger, E. (1984). Express-saccades of the monkey: Effect of daily training on probability of occurrence and reaction time. *Experimental Brain Research*, 55(2):232–242.

Flaherty, A. W. and Graybiel, A. M. (1991). Corticostriatal transformations in the primate somatosensory system. Projections from physiologically mapped body-part representations. *Journal of Neurophysiology*, 66(4):1249–63.

Floden, D. and Stuss, D. T. (2006). Inhibitory control is slowed in patients with right superior medial frontal damage. *Journal of Cognitive Neuroscience*, 18(11):1843–1849.

Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, 19(8):1120–36.

Friedman, N. P. and Miyake, A. (2004). The relations among inhibition and interference control functions: A latent-variable analysis. *Journal of Experimental Psychology: General*, 133(1):101–135.

Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Science*, 9(10):474–80.

Fries, W. (1984). Cortical projections to the superior colliculus in the macaque monkey: a retrograde study using horseradish peroxidase. *Journal of Comparative Neurology*, 230(1):55–76.

Fu, Z., Beam, D., Chung, J. M., Reed, C. M., Mamelak, A. N., Adolphs, R., and Rutishauser, U. (2022a). The geometry of domain-general performance monitoring in the human medial frontal cortex. *Science*, 376(6593):eabm9922.

Fu, Z., Sajad, A., Errington, S. P., Schall, J. D., and Rutishauser, U. (2022b). Neurophysiological mechanisms of error monitoring in human and nonhuman primates. *Nature Reviews Neuroscience*, In press.

Fu, Z., Wu, D.-A. J., Ross, I., Chung, J. M., Mamelak, A. N., Adolphs, R., and Rutishauser, U. (2019). Single-neuron correlates of error monitoring and post-error adjustments in human medial frontal cortex. *Neuron*, 101(1):165–177.e5.

Fuchs, A. (1967). Saccadic and smooth pursuit eye movements in the monkey. *Journal of Physiology*, 191(3):609.

Fujii, N. and Graybiel, A. M. (2003). Representation of action sequence boundaries by macaque prefrontal cortical neurons. *Science*, 301(5637):1246–1249.

Fuster, J. M. (2001). The prefrontal cortex—an update: time is of the essence. *Neuron*, 30(2):319–333.

Gauggel, S., Rieger, M., and Feghoff, T. (2004). Inhibition of ongoing responses in patients with parkinson's disease. *Journal of Neurology, Neurosurgery & Psychiatry*, 75(4):539–544.

Gehring, W. J. and Fencsik, D. E. (2001). Functions of the medial frontal cortex in the processing of conflict and errors. *Journal of Neuroscience*, 21(23):9430–7.

Gehring, W. J., Goss, B., Coles, M. G., Meyer, D. E., and Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science*, 4(6):385–390.

Gehring, W. J. and Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science*, 295(5563):2279–2282.

Genovesio, A., Brasted, P. J., and Wise, S. P. (2006). Representation of future and previous spatial goals by separate neural populations in prefrontal cortex. *Journal of Neuroscience*, 26(27):7305–7316.

Georgopoulos, A. P., DeLong, M. R., and Crutcher, M. D. (1983). Relations between parameters of step-tracking movements and single cell discharge in the globus pallidus and subthalamic nucleus of the behaving monkey. *Journal of Neuroscience*, 3(8):1586–98.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 2(11):1527–37.

Georgopoulos, A. P., Kalaska, J. F., and Massey, J. T. (1981). Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty, and change in target location. *Journal of Neurophysiology*, 46(4):725–43.

Geyer, S., Ledberg, A., Schleicher, A., Kinomura, S., Schormann, T., Burgel, U., Klingberg, T., Larsson, J., Zilles, K., and Roland, P. E. (1996). Two different areas within the primary motor cortex of man. *Nature*, 382(6594):805–7.

Geyer, S., Matelli, M., Luppino, G., and Zilles, K. (2000). Functional neuroanatomy of the primate isocortical motor system. *Anatomy and embryology*, 202(6):443–474.

Giamundo, M., Giarrocco, F., Brunamonti, E., Fabbrini, F., Pani, P., and Ferraina, S. (2021). Neuronal activity in the premotor cortex of monkeys reflects both cue salience and motivation for action generation and inhibition. *Journal of Neuroscience*, 41(36):7591–7606.

Giarrocco, F., Bardella, G., Giamundo, M., Fabbrini, F., Brunamonti, E., Pani, P., and Ferraina, S. (2021). Neuronal dynamics of signal selective motor plan cancellation in the macaque dorsal premotor cortex. *Cortex*, 135:326–340.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 84(3):279.

Gibbon, J., Church, R. M., and Meck, W. H. (1984). Scalar timing in memory. *Annals of the New York Academy of Sciences*, 423:52–77.

Gidon, A. and Segev, I. (2012). Principles governing the operation of synaptic inhibition in dendrites. *Neuron*, 75(2):330–41.

Gillies, A. J. and Willshaw, D. J. (1998). A massively connected subthalamic nucleus leads to the generation of widespread pulses. *Proceedings of the Royal Society B: Biological Sciences*, 265(1410):2101–9.

Gnanalingham, K. K., Byrne, E. J., Thornton, A., Sambrook, M. A., and Bannister, P. (1997). Motor and cognitive function in lewy body dementia: comparison with alzheimer's and parkinson's diseases. *Journal of Neurology, Neurosurgery & Psychiatry*, 62(3):243–252.

Godlove, D. C., Emeric, E. E., Segovis, C. M., Young, M. S., Schall, J. D., and Woodman, G. F. (2011). Event-Related Potentials Elicited by Errors during the Stop-Signal Task. I. Macaque Monkeys. *Journal of Neuroscience*, 31(44):15640–15649.

Godlove, D. C., Maier, A., Woodman, G. F., and Schall, J. D. (2014). Microcircuitry of agranular frontal cortex: Testing the generality of the canonical cortical microcircuit. *Journal of Neuroscience*, 34(15):5355–5369.

Godlove, D. C. and Schall, J. D. (2016). Microsaccade production during saccade cancelation in a stop-signal task. *Vision Research*, 118:5–16.

Goldman-Rakic, P. S. and Porrino, L. J. (1985). The primate mediodorsal (MD) nucleus and its projection to the frontal lobe. *Journal of Comparative Neurology*, 242(4):535–60.

Goonetilleke, S. C., Wong, J. P., and Corneil, B. D. (2012). Validation of a within-trial measure of the oculomotor stop process. *Journal of Neurophysiology*, 108(3):760–770.

Grace, A. A. and Bunney, B. S. (1980). Nigral dopamine neurons: Intracellular recording and identification with L-dopa injection and histofluorescence. *Science*, 210(4470):654–6.

Gratton, G., Coles, M. G., and Donchin, E. (1992). Optimizing the use of information: strategic control of activation of responses. *Journal of Experimental Psychology: General*, 121(4):480.

Greenhouse, I., Oldenkamp, C. L., and Aron, A. R. (2012). Stopping a response has global or nonglobal effects on the motor system depending on preparation. *Journal of Neurophysiology*, 107(1):384–92.

Grefkes, C. and Fink, G. R. (2005). The functional organization of the intraparietal sulcus in humans and monkeys. *Journal of Anatomy*, 207(1):3–17.

Griggs, W. S., Kim, H. F., Ghazizadeh, A., Costello, M. G., Wall, K. M., and Hikosaka, O. (2017). Flexible and stable value coding areas in caudate head and tail receive anatomically distinct cortical and subcortical inputs. *Frontiers in Neuroanatomy*, 11:106.

Grosbras, M.-H., Lobel, E., Van de Moortele, P.-F., LeBihan, D., and Berthoz, A. (1999). An anatomical landmark for the supplementary eye fields in human revealed with functional magnetic resonance imaging. *Cerebral Cortex*, 9(7):705–711.

Grossman, C. D. and Cohen, J. Y. (2022). Neuromodulation and neurophysiology on the timescale of learning and decision-making. *Annual Review of Neuroscience*, 45(1):317–337.

Hafed, Z. M., Goffart, L., and Krauzlis, R. J. (2008). Superior colliculus inactivation causes stable offsets in eye position during tracking. *Journal of Neuroscience*, 28(32):8124–37.

Hafed, Z. M. and Krauzlis, R. J. (2012). Similarity of superior colliculus involvement in microsaccade and saccade generation. *Journal of Neurophysiology*, 107(7):1904–16.

Haggard, P. (2008). Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, 9(12):934–46.

Hajcak, G., McDonald, N., and Simons, R. F. (2003). To err is autonomic: Error-related brain potentials, ANS activity, and post-error compensatory behavior. *Psychophysiology*, 40(6):895–903.

Hammond, D. C. (2005). Neurofeedback treatment of depression and anxiety. *Journal of Adult Development*, 12(2):131–137.

Hampshire, A., Chamberlain, S. R., Monti, M. M., Duncan, J., and Owen, A. M. (2010). The role of the right inferior frontal gyrus: inhibition and attentional control. *NeuroImage*, 50(3):1313–1319.

Handel, A. and Glimcher, P. W. (1999). Quantitative analysis of substantia nigra pars reticulata activity during a visually guided saccade task. *Journal of Neurophysiology*, 82(6):3458–75.

Hanes, D. P. and Carpenter, R. (1999). Countermanding saccades in humans. *Vision Research*, 39(16):2777–2791.

Hanes, D. P., Patterson, W. F., and Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *Journal of Neurophysiology*, 79(2):817–834.

Hanes, D. P. and Schall, J. D. (1995). Countermanding saccades in macaque. *Visual Neuroscience*, 12(5):929–37.

Hanes, D. P. and Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, 274(5286):427–430.

Hannah, R., Muralidharan, V., Sundby, K. K., and Aron, A. R. (2020). Temporally-precise disruption of prefrontal cortex informed by the timing of beta bursts impairs human action-stopping. *NeuroImage*, 222:117222.

Hanslmayr, S., Matuschek, J., and Fellner, M.-C. (2014). Entrainment of prefrontal beta oscillations induces an endogenous echo and impairs memory formation. *Current Biology*, 24(8):904–909.

Harnishfeger, K. K. (1995). *The development of cognitive inhibition: Theories, definitions, and research evidence*, pages 175–204. Academic Press, San Diego, CA, US.

Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., and Platt, M. L. (2011a). Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *Journal of Neuroscience*, 31(11):4178–4187.

Hayden, B. Y., Pearson, J. M., and Platt, M. L. (2011b). Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, 14(7):933–939.

Haynes, W. I. and Haber, S. N. (2013). The organization of prefrontal-subthalamic inputs in primates provides an anatomical substrate for both functional specificity and integration: implications for Basal Ganglia models and deep brain stimulation. *Journal of Neuroscience*, 33(11):4804–14.

He, S. Q., Dum, R. P., and Strick, P. L. (1995). Topographic organization of corticospinal projections from the frontal lobe: motor areas on the medial surface of the hemisphere. *Journal of Neuroscience*, 15(5 Pt 1):3284–306.

Heathcote, A., Popiel, S. J., and Mewhort, D. J. (1991). Analysis of response time distributions: an example using the Stroop task. *Psychological Bulletin*, 109(2):340.

Heilbronner, S. R. and Hayden, B. Y. (2015). Dorsal anterior cingulate cortex: A bottom-up view. *Annual Review of Neuroscience*, 39(1):1–22.

Heitz, R. P., Cohen, J. Y., Woodman, G. F., and Schall, J. D. (2010). Neural correlates of correct and errant attentional selection revealed through n2pc and frontal eye field activity. *Journal of Neurophysiology*, 104(5):2433–2441.

Heitz, R. P. and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, 76(3):616–628.

Herrera, B., Sajad, A., Woodman, G. F., Schall, J. D., and Riera, J. J. (2020). A minimal biophysical model of neocortical pyramidal cells: Implications for frontal cortex microcircuitry and field potential generation. *Journal of Neuroscience*, 40(44):8513–8529.

Herrera, B., Westerberg, J. A., Schall, M. S., Maier, A., Woodman, G. F., Schall, J. D., and Riera, J. J. (2022). Resolving the mesoscopic missing link: Biophysical modeling of EEG from cortical columns in primates. *bioRxiv*.

Herrmann, C. S., Strüber, D., Helfrich, R. F., and Engel, A. K. (2016). EEG oscillations: from correlation to causality. *International Journal of Psychophysiology*, 103:12–21.

Heukelum, S. v., Mars, R. B., Guthrie, M., Buitelaar, J. K., Beckmann, C. F., Tiesinga, P. H., Vogt, B. A., Glennon, J. C., and Havenith, M. N. (2020). Where is cingulate cortex? a cross-species view. *Trends in Neurosciences*, 43(5):285–299.

Hikosaka, O. and Isoda, M. (2010). Switching from automatic to controlled behavior: cortico-basal ganglia mechanisms. *Trends in Cognitive Science*, 14(4):154–61.

Hikosaka, O. and Wurtz, R. H. (1983a). Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. *Journal of Neurophysiology*, 49(5):1230–53.

Hikosaka, O. and Wurtz, R. H. (1983b). Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *Journal of Neurophysiology*, 49(5):1268–84.

Hikosaka, O. and Wurtz, R. H. (1983c). Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. *Journal of Neurophysiology*, 49(5):1285–301.

Hikosaka, O. and Wurtz, R. H. (1985). Modification of saccadic eye movements by GABA-related substances. II. Effects of muscimol in monkey substantia nigra pars reticulata. *Journal of Neurophysiology*, 53(1):292–308.

Holroyd, C. B. and Krigolson, O. E. (2007). Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*, 44(6):913–917.

Holroyd, C. B., Yeung, N., Coles, M. G. H., and Cohen, J. D. (2005). A mechanism for error detection in speeded response time tasks. *Journal of Experimental Psychology: General*, 134(2):163–191.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558.

Hu, S. and Li, C.-S. R. (2012). Neural processes of preparatory control for stop signal inhibition. *Human Brain Mapping*, 33(12):2785–2796.

Hubel, D. H. and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195(1):215–43.

Huerta, M. F. and Kaas, J. H. (1990). Supplementary eye field as defined by intracortical microstimulation: Connections in macaques. *Journal of Comparative Neurology*, 293(2):299–330.

Husain, M., Parton, A., Hodgson, T. L., Mort, D., and Rees, G. (2003). Self-control during response conflict by human supplementary eye field. *Nature Neuroscience*, 6(2):117–118.

Huster, R. J., Enriquez-Geppert, S., Lavallee, C. F., Falkenstein, M., and Herrmann, C. S. (2013). Electroencephalography of response inhibition tasks: Functional networks and cognitive contributions. *International Journal of Psychophysiology*, 87(3):217–33.

Huster, R. J., Messel, M. S., Thunberg, C., and Raud, L. (2020). The P300 as marker of inhibitory control - fact or fiction? *Cortex*, 132:334–348.

Ito, S., Stuphorn, V., Brown, J. W., and Schall, J. D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science*, 302(5642):120–122.

Iwamuro, H., Tachibana, Y., Ugawa, Y., Saito, N., and Nambu, A. (2017). Information processing from the motor cortices to the subthalamic nucleus and globus pallidus and their somatotopic organizations revealed electrophysiologically in monkeys. *European Journal of Neuroscience*, 46(11):2684–2701.

Izawa, Y., Suzuki, H., and Shinoda, Y. (2004a). Suppression of visually and memory-guided saccades induced by electrical stimulation of the monkey frontal eye field. I. Suppression of ipsilateral saccades. *Journal of Neurophysiology*, 92(4):2248–60.

Izawa, Y., Suzuki, H., and Shinoda, Y. (2004b). Suppression of visually and memory-guided saccades induced by electrical stimulation of the monkey frontal eye field. II. Suppression of bilateral saccades. *Journal of Neurophysiology*, 92(4):2261–73.

Izawa, Y., Suzuki, H., and Shinoda, Y. (2009). Response properties of fixation neurons and their location in the frontal eye field in the monkey. *Journal of Neurophysiology*, 102(4):2410–22.

Jacobsen, C. (1936). Studies of cerebral function in primates. i. the functions of the frontal association areas in monkeys. comp. *Psychological Monographs*, 13.

Jahanshahi, M., Jenkins, I. H., Brown, R. G., Marsden, C. D., Passingham, R. E., and Brooks, D. J. (1995). Self-initiated versus externally triggered movements. I. An investigation using measurement of regional cerebral blood flow with PET and movement-related potentials in normal and Parkinson's disease subjects. *Brain*, 118 (Pt 4):913–33.

Jana, S., Hannah, R., Muralidharan, V., and Aron, A. R. (2020). Temporal cascade of frontal, motor and muscle processes underlying human action-stopping. *eLife*, 9:e50371.

Janer, K. W. and Pardo, J. V. (1991). Deficits in selective attention following bilateral anterior cingulotomy. *Journal of Cognitive Neuroscience*, 3(3):231–241.

Janssen, P. and Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience*, 8(2):234–41.

Jasper, H. and Penfield, W. (1949). Electrocorticograms in man: Effect of voluntary movement upon the electrical activity of the precentral gyrus. *Archiv für Psychiatrie und Nervenkrankheiten*, 183(1-2):163–174.

Jay, M. F. and Sparks, D. L. (1987). Sensorimotor integration in the primate superior colliculus. i. motor convergence. *Journal of Neurophysiology*, 57(1):22–34.

Jeannerod, M. (1984). The timing of natural prehension movements. *Journal of Motor Behavior*, 16(3):235–54.

Jeannerod, M., Arbib, M. A., Rizzolatti, G., and Sakata, H. (1995). Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends in Neuroscience*, 18(7):314–20.

Jha, A., Nachev, P., Barnes, G., Husain, M., Brown, P., and Litvak, V. (2015). The frontal control of stopping. *Cerebral Cortex*, 25(11):4392–4406.

Jiang, J., Wagner, A. D., and Egner, T. (2018). Integrated externally and internally generated task predictions jointly guide cognitive control in prefrontal cortex. *eLife*, 7:e39497.

Jóhannesson, Ó. I., Edelman, J. A., Sigurórsson, B. D., and Kristjánsson, Á. (2018). Effects of saccade training on express saccade proportions, saccade latencies, and peak velocities: an investigation of nasal/temporal differences. *Experimental Brain Research*, 236(5):1251–1262.

Joiner, W. M., Lee, J.-E., and Shelhamer, M. (2007). Behavioral analysis of predictive saccade tracking as studied by countermanding. *Experimental brain research*, 181(2):307–320.

Jong, R. D., Coles, M. G. H., and Logan, G. D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology: Human Perception and Performance*, 21(3):498–511.

Kawagoe, R., Takikawa, Y., and Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neuroscience*, 1(5):411–6.

Kawaguchi, N., Sakamoto, K., Saito, N., Furusawa, Y., Tanji, J., Aoki, M., and Mushiake, H. (2015). Surprise signals in the supplementary eye field: Rectified prediction errors drive exploration-exploitation transitions. *Journal of Neurophysiology*, 113(3):1001–14.

Kawai, T., Yamada, H., Sato, N., Takada, M., and Matsumoto, M. (2015). Roles of the lateral habenula and anterior cingulate cortex in negative outcome monitoring and behavioral adjustment in nonhuman primates. *Neuron*, 88(4):792–804.

Kennerley, S. W., Behrens, T. E., and Wallis, J. D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience*, 14(12):1581–1589.

Kennerley, S. W., Dahmubed, A. F., Lara, A. H., and Wallis, J. D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. *Journal of Cognitive Neuroscience*, 21(6):1162–1178.

Kerns, J. G., Cohen, J. D., III, A. W. M., Cho, R. Y., Stenger, V. A., and Carter, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. *Science*, 303(5660):1023–1026.

Khalighinejad, N., Garrett, N., Priestley, L., Lockwood, P., and Rushworth, M. F. (2021). A habenula-insular circuit encodes the willingness to act. *Nature Communications*, 12(1):1–12.

Kim, H. F. and Hikosaka, O. (2013). Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron*, 79(5):1001–10.

Kim, J., Ghim, J. W., Lee, J. H., and Jung, M. W. (2013). Neural correlates of interval timing in rodent prefrontal cortex. *Journal of Neuroscience*, 33(34):13834–47.

Klein-Flügge, M. C., Bongioanni, A., and Rushworth, M. F. (2022). Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron*.

Knox, P. C. and Wolohan, F. D. A. (2015). Temporal stability and the effects of training on saccade latency in "express saccade makers". *PLoS ONE*, 10(3):e0120437.

Kok, A., Ramautar, J. R., Ruiter, M. B. D., Band, G. P., and Ridderinkhof, K. R. (2004). ERP components associated with successful and unsuccessful stopping in a stop-signal task. *Psychophysiology*, 41(1):9–20.

Kolling, N. and O'Reilly, J. X. (2018). State-change decisions and dorsomedial prefrontal cortex: the importance of time. *Current Opinion in Behavioral Sciences*, 22(55):152–160.

Kolling, N., Wittmann, M. K., Behrens, T. E., Boorman, E. D., Mars, R. B., and Rushworth, M. F. (2016). Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, 19(10):1280–5.

Kornylo, K., Dill, N., Saenz, M., and Krauzlis, R. J. (2003). Canceling of pursuit and saccadic eye movements in humans and monkeys. *Journal of Neurophysiology*, 89(6):2984–2999.

Koyama, M., Hasegawa, I., Osada, T., Adachi, Y., Nakahara, K., and Miyashita, Y. (2004). Functional magnetic resonance imaging of macaque monkeys performing visually guided saccade tasks: comparison of cortical eye fields with humans. *Neuron*, 41(5):795–807.

Kozak, R. A. and Corneil, B. D. (2021). High-contrast, moving targets in an emerging target paradigm promote fast visuomotor responses during visually guided reaching. *Journal of Neurophysiology*, 126(1):68–81.

Krauzlis, R. J., Goffart, L., and Hafed, Z. M. (2017). Neuronal control of fixation and fixational eye movements. *Proceedings of the Royal Society B: Biological Sciences*, 372(1718).

Krämer, U. M., Solbakk, A.-K., Funderud, I., Løvstad, M., Endestad, T., and Knight, R. T. (2013). The role of the lateral prefrontal cortex in inhibitory motor control. *Cortex*, 49(3):837–849.

Kudo, K. and Ohtsuki, T. (1998). Functional modification of agonist-antagonist electromyographic activity for rapid movement inhibition. *Experimental Brain Research*, 122(1):23–30.

Ladouceur, C. D., Dahl, R. E., and Carter, C. S. (2007). Development of action monitoring through adolescence into adulthood: ERP and source localization. *Developmental Science*, 10(6):874–891.

Lappin, J. S. and Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology*, 72(6):805.

Lear, A., Baker, S. N., Clarke, H. F., Roberts, A. C., Schmid, M. C., and Jarrett, W. (2022). Understanding them to understand ourselves: The importance of nhp research for translational neuroscience. *Current Research in Neurobiology*, 3:100049.

Lecas, J. C., Requin, J., Anger, C., and Vitton, N. (1986). Changes in neuronal activity of the monkey precentral cortex during preparation for movement. *Journal of Neurophysiology*, 56(6):1680–702.

Leigh, R. J. and Zee, D. S. (2015). *The neurology of eye movements*. Contemporary Neurology.

Lemon, R. N., Baker, S. N., and Kraskov, A. (2021). Classification of cortical neurons by spike shape and the identification of pyramidal neurons. *Cerebral Cortex*, 31(11):5131–5138.

Lenartowicz, A., Verbruggen, F., Logan, G. D., and Poldrack, R. A. (2011). Inhibition-related activation in the right inferior frontal gyrus in the absence of inhibitory cues. *Journal of Cognitive Neuroscience*, 23(11):3388–3399.

Libet, B., Gleason, C. A., Wright, E. W., and Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain*, (106):623–42.

Lidow, M. S., Goldman-Rakic, P. S., Gallager, D. W., and Rakic, P. (1991). Distribution of dopaminergic receptors in the primate cerebral cortex: Quantitative autoradiographic analysis using [3H]raclopride, [3H]spiperone and [3H]SCH23390. *Neuroscience*, 40(3):657–71.

Lijffijt, M., Kenemans, J. L., Verbaten, M. N., and van Engeland, H. (2005). A meta-analytic review of stopping performance in attention-deficit/hyperactivity disorder: deficient inhibitory motor control? *Journal of Abnormal Psychology*, 114(2):216.

Lim, S. and Goldman, M. S. (2013). Balanced cortical microcircuitry for maintaining information in working memory. *Nature Neuroscience*, 16(9):1306–14.

Littman, R. and Takács, (2017). Do all inhibitions act alike? A study of go/no-go and stop-signal paradigms. *PloS one*, 12(10):e0186774.

Llinas, R. R. (2001). *I of the vortex : from neurons to self*. MIT Press, Cambridge, Massachusetts, 1st ed. edition.

Lo, C. C., Boucher, L., Paré, M., Schall, J. D., and Wang, X. J. (2009). Proactive inhibitory control and attractor dynamics in countermanding action: A spiking neural circuit model. *Journal of Neuroscience*, 29(28):9059–71.

Lo, C. C. and Wang, X. J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience*, 9(7):956–63.

Logan, G. D. (2003). Executive control of thought and action: In search of the wild homunculus. *Current Directions in Psychological Science*, 12(2):45–48.

Logan, G. D. and Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91(3):295.

Logan, G. D. and Irwin, D. E. (2000). Don't look! Don't touch! Inhibitory control of eye and hand movements. *Psychonomic Bulletin & Review*, 7(1):107–12.

Logan, G. D., Schachar, R. J., and Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science*, 8(1):60–64.

Logan, G. D., Yamaguchi, M., Schall, J. D., and Palmeri, T. J. (2015). Inhibitory control in mind and brain 2.0: Blocked-input models of saccadic countermanding. *Psychological Review*, 122(2):115–47.

Lowe, K. A. and Schall, J. D. (2018). Functional categories of visuomotor neurons in macaque frontal eye field. *eNeuro*, 5(5).

Lowe, K. A. and Schall, J. D. (2019). Sequential operations revealed by serendipitous feature selectivity in frontal eye field. *bioRxiv*, page 683144.

Lu, M. T., Preston, J. B., and Strick, P. L. (1994). Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. *Journal of Comparative Neurology*, 341(3):375–92.

Luce, R. D. (1986). *Response times : their role in inferring elementary mental organization*. Oxford psychology series. Oxford University Press, New York.

Luck, S. J. (2014). *An Introduction to the Event-Related Potential Technique, Second Edition*. MIT Press, Cambridge.

Luders, H. O., Dinner, D. S., Morris, H. H., Wyllie, E., and Comair, Y. G. (1995). Cortical electrical stimulation in humans. The negative motor areas. *Adv Neurol*, 67:115–29.

Luk, C.-H. and Wallis, J. D. (2013). Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. *Journal of Neuroscience*, 33(5):1864–1871.

Lundqvist, M., Rose, J., Herman, P., Brincat, S. L., Buschman, T. J., and Miller, E. K. (2016). Gamma and beta bursts underlie working memory. *Neuron*, 90(1):152–164.

Luppino, G., Matelli, M., Camarda, R., and Rizzolatti, G. (1993). Corticocortical connections of area F3 (SMA-proper) and area F6 (pre-SMA) in the macaque monkey. *Journal of Comparative Neurology*, 338(1):114–40.

Luu, P., Tucker, D. M., and Makeig, S. (2004). Frontal midline theta and the error-related negativity: neurophysiological mechanisms of action regulation. *Clinical Neurophysiology*, 115(8):1821–1835.

Lynch, J. C., Mountcastle, V. B., Talbot, W. H., and Yin, T. C. (1977). Parietal lobe mechanisms for directed visual attention. *Journal of Neurophysiology*, 40(2):362–389.

MacDonald, H. J., Coxon, J. P., Stinear, C. M., and Byblow, W. D. (2014). The fall and rise of corticomotor excitability with cancellation and reinitiation of prepared action. *Journal of Neurophysiology*, 112(11):2707–17.

MacLeod, C. M. (2007). *The concept of inhibition in cognition*, pages 3–23. American Psychological Association, Washington, DC, US.

Maier, A., Adams, G. K., Aura, C., and Leopold, D. A. (2010). Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation. *Frontiers in System Neuroscience*, 4.

Majid, D. S., Cai, W., George, J. S., Verbruggen, F., and Aron, A. R. (2012). Transcranial magnetic stimulation reveals dissociable mechanisms for global versus selective corticomotor suppression underlying the stopping of action. *Cerebral Cortex*, 22(2):363–71.

Mann, L. G., Hay, K. R., Song, A. K., Errington, S. P., Trujillo, P., Zald, D. H., Yan, Y., Kang, H., Logan, G. D., and Claassen, D. O. (2021). D2-like receptor expression in the hippocampus and amygdala informs performance on the stop-signal task in Parkinson's disease. *Journal of Neuroscience*, 41(48):10023–10030.

Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84.

Mar, K., Townes, P., Pechlivanoglou, P., Arnold, P., and Schachar, R. (2022). Obsessive compulsive disorder and response inhibition: Meta-analysis of the stop-signal task. *Journal of Psychopathology and Clinical Science*, 131(2):152.

Marco-Pallars, J., Camara, E., Mnte, T. F., and Rodrguez-Fornells, A. (2008). Neural mechanisms underlying adaptive actions after slips. *Journal of Cognitive Neuroscience*, 20(9):1595–1610.

Marino, R. A., Rodgers, C. K., Levy, R., and Munoz, D. P. (2008). Spatial relationships of visuomotor transformations in the superior colliculus map. *Journal of Neurophysiology*, 100(5):2564–2576.

Martinez-Trujillo, J. C., Medendorp, W. P., Wang, H., and Crawford, J. D. (2004). Frames of reference for eye-head gaze commands in primate supplementary eye fields. *Neuron*, 44(6):1057–1066.

Matelli, M., Camarda, R., Glickstein, M., and Rizzolatti, G. (1986). Afferent and efferent projections of the inferior area 6 in the macaque monkey. *Journal of Comparative Neurology*, 251(3):281–298.

Matelli, M., Govoni, P., Galletti, C., Kutz, D. F., and Luppino, G. (1998). Superior area 6 afferents from the superior parietal lobule in the macaque monkey. *Journal of Comparative Neurology*, 402(3):327–352.

Matelli, M., Luppino, G., Fogassi, L., and Rizzolatti, G. (1989). Thalamic input to inferior area 6 and area 4 in the macaque monkey. *Journal of Comparative Neurology*, 280(3):468–488.

Matelli, M., Luppino, G., and Rizzolatti, G. (1991). Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. *Journal of Comparative Neurology*, 311(4):445–62.

Matsumura, M., Kojima, J., Gardiner, T. W., and Hikosaka, O. (1992). Visual and oculomotor functions of monkey subthalamic nucleus. *Journal of Neurophysiology*, 67(6):1615–32.

Matzke, D., Dolan, C. V., Logan, G. D., Brown, S. D., and Wagenmakers, E. J. (2013a). Bayesian parametric estimation of stop-signal reaction time distributions. *Journal of Experimental Psychology: General*, 142(4):1047–73.

Matzke, D., Hughes, M., Badcock, J. C., Michie, P., and Heathcote, A. (2017a). Failures of cognitive control or attention? The case of stop-signal deficits in schizophrenia. *Attention, Perception, & Psychophysics*, 79(4):1078–1086.

Matzke, D., Love, J., and Heathcote, A. (2017b). A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behavior Research Methods*, 49(1):267–281.

Matzke, D., Love, J., Wiecki, T. V., Brown, S. D., Logan, G. D., and Wagenmakers, E.-J. (2013b). Release the BEESTS: Bayesian Estimation of Ex-Gaussian STop-Signal reaction time distributions. *Frontiers in Psychology*, 4:918.

Matzke, D. and Wagenmakers, E.-J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review*, 16(5):798–817.

McPeek, R. M. and Schiller, P. H. (1994). The effects of visual scene composition on the latency of saccadic eye movements of the rhesus monkey. *Vision Research*, 34(17):2293–2305.

Medalla, M. and Barbas, H. (2009). Synapses with inhibitory neurons differentiate anterior cingulate from dorsolateral prefrontal pathways associated with cognitive control. *Neuron*, 61(4):609–20.

Mendoza-Halliday, D., Major, A. J., Lee, N., Lichtenfeld, M., Carlson, B., Mitchell, B., Meng, P. D., Xiong, Y. S., Westerberg, J. A., Maier, A., et al. (2022). A ubiquitous spectrolaminar motif of local field potential power across the primate cortex. *bioRxiv*.

Mesulam, M.-M. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology*, 10(4):309–325.

Meyer, G. (1987). Forms and spatial arrangement of neurons in the primary motor cortex of man. *Journal of Comparative Neurology*, 262(3):402–28.

Meyer, H. C. and Bucci, D. J. (2016). Neural and behavioral mechanisms of proactive and reactive inhibition. *Learning & Memory*, 23(10):504–514.

Middlebrooks, P. G., Zandbelt, B. B., Logan, G. D., Palmeri, T. J., and Schall, J. D. (2020). Countermanding Perceptual Decision-Making. *iScience*, 23(1):100777.

Miller, E. K. and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1):167–202.

Miller, G. A., Galanter, E., and Pribram, K. H. (1960). *The integration of plans*, pages 95–102. Henry Holt and Co, New York, NY, US.

Milstein, D. M. and Dorris, M. C. (2007). The influence of expected value on saccadic preparation. *Journal of Neuroscience*, 27(18):4810–4818.

Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50(4):381–425.

Mirabella, G., Pani, P., and Ferraina, S. (2011). Neural correlates of cognitive control of reaching movements in the dorsal premotor cortex of rhesus monkeys. *Journal of Neurophysiology*, 106(3):1454–1466.

Mitchell, S. J., Richardson, R. T., Baker, F. H., and DeLong, M. R. (1987). The primate globus pallidus: neuronal activity related to direction of movement. *Experimental Brain Research*, 68(3):491–505.

Mitz, A. and Wise, S. (1987). The somatotopic organization of the supplementary motor area: intracortical microstimulation mapping. *Journal of Neuroscience*, 7(4):1010–1021.

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., and Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, 41(1):49–100.

Monakow, K. H., Akert, K., and Kunzle, H. (1978). Projections of the precentral motor cortex and other cortical areas of the frontal lobe to the subthalamic nucleus in the monkey. *Experimental Brain Research*, 33(3-4):395–403.

Monosov, I. E. (2017). Anterior cingulate is a source of valence-specific information about value and uncertainty. *Nature Communications*, 8(1):134.

Monosov, I. E., Haber, S. N., Leuthardt, E. C., and Jezzini, A. (2020). Anterior cingulate cortex and the control of dynamic behavior in primates. *Current Biology*, 30(23):R1442–R1454.

Monosov, I. E. and Thompson, K. G. (2009). Frontal eye field activity enhances object identification during covert visual search. *Journal of Neurophysiology*, 102(6):3656–3672.

Morein-Zamir, S. and Kingstone, A. (2006). Fixation offset and stop signal intensity effects on saccadic countermanding: a crossmodal investigation. *Experimental Brain Research*, 175(3):453–462.

Mosher, C. P., Mamelak, A. N., Malekmohammadi, M., Pouratian, N., and Rutishauser, U. (2021). Distinct roles of dorsal and ventral subthalamic neurons in action selection and cancellation. *Neuron*, 109(5):869–881 e6.

Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H., and Acuna, C. (1975). Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *Journal of Neurophysiology*, 38(4):871–908.

Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., and O'Reilly, R. C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Sciences*, 15(10):453–459.

Munoz, D. P. and Istvan, P. J. (1998). Lateral inhibitory interactions in the intermediate layers of the monkey superior colliculus. *Journal of Neurophysiology*, 79(3):1193–209.

Munoz, D. P. and Wurtz, R. H. (1993). Fixation cells in monkey superior colliculus. II. Reversible activation and deactivation. *Journal of Neurophysiology*, 70(2):576–89.

Munoz, D. P. and Wurtz, R. H. (1995). Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. *Journal of Neurophysiology*, 73(6):2313–33.

Muralidharan, V., Aron, A. R., and Schmidt, R. (2022). Transient beta modulates decision thresholds during human action-stopping. *NeuroImage*, 254:119145.

Murata, A., Gallese, V., Luppino, G., Kaseda, M., and Sakata, H. (2000). Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. *Journal of Neurophysiology*, 83(5):2580–601.

Murphy, J. T., Wong, Y. C., and Kwan, H. C. (1985). Sequential activation of neurons in primate motor cortex during unrestrained forelimb movement. *Journal of Neurophysiology*, 53(2):435–45.

Murthy, A., Ray, S., Shorter, S. M., Schall, J. D., and Thompson, K. G. (2009). Neural control of visual search by frontal eye field: Effects of unexpected target displacement on visual selection and saccade preparation. *Journal of Neurophysiology*, 101(5):2485–2506.

Näätänen, R. (1970). The diminishing time-uncertainty with the lapse of time after the warning signal in reaction-time experiments with varying fore-periods. *Acta Psychologica*, 34:399–419.

Nachev, P., Kennard, C., and Husain, M. (2008). Functional role of the supplementary and pre-supplementary motor areas. *Nature Reviews Neuroscience*, 9(11):856–69.

Nachev, P., Rees, G., Parton, A., Kennard, C., and Husain, M. (2005). Volition and conflict in human medial frontal cortex. *Current Biology*, 15(2):122–8.

Nachev, P., Wydell, H., O'Neill, K., Husain, M., and Kennard, C. (2007). The role of the pre-supplementary motor area in the control of action. *NeuroImage*, 36(3-3):T155–T163.

Næss, S., Halnes, G., Hagen, E., Hagler Jr, D. J., Dale, A. M., Einevoll, G. T., and Ness, T. V. (2021). Biophysically detailed forward modeling of the neural origin of EEG and MEG signals. *NeuroImage*, 225:117467.

Nakamura, K., Roesch, M. R., and Olson, C. R. (2005). Neuronal activity in macaque sef and acc during performance of tasks involving conflict. *Journal of Neurophysiology*, 93(2):884–908.

Nakamura, K., Sakai, K., and Hikosaka, O. (1998). Neuronal activity in medial frontal cortex during learning of sequential procedures. *Journal of Neurophysiology*, 80(5):2671–87.

Nambu, A., Tokuno, H., Inase, M., and Takada, M. (1997). Corticosubthalamic input zones from forelimb representations of the dorsal and ventral divisions of the premotor cortex in the macaque monkey: comparison with the input zones from the primary motor cortex and the supplementary motor area. *Neuroscience letters*, 239(1):13–16.

Nambu, A., Tokuno, H., and Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal 'hyperdirect' pathway. *Neuroscience Research*, 43(2):111–7.

Nambu, A., Yoshida, S., and Jinnai, K. (1990). Discharge patterns of pallidal neurons with input from various cortical areas during movement in the monkey. *Brain Research*, 519(1-2):183–91.

Nelson, M. J., Boucher, L., Logan, G. D., Palmeri, T. J., and Schall, J. D. (2010). Nonindependent and nonstationary response times in stopping and stepping saccade tasks. *Attention Perception Psychophysics*, 72(7):1913–29.

Neubert, F.-X., Mars, R. B., Sallet, J., and Rushworth, M. F. (2015). Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proceedings of the National Academy of Sciences*, 112(20):E2695–E2704.

Niemi, P. and Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, 89(1):133.

Nieuwenhuis, S., Schweizer, T. S., Mars, R. B., Botvinick, M. M., and Hajcak, G. (2007). Error-likelihood prediction in the medial frontal cortex: a critical evaluation. *Cerebral Cortex*, 17(7):1570–1581.

Nigbur, R., Ivanova, G., and Stürmer, B. (2011). Theta power as a marker for cognitive interference. *Clinical Neurophysiology*, 122(11):2185–2194.

Nigg, J. T. (2000). On inhibition/disinhibition in developmental psychopathology: Views from cognitive and personality psychology and a working inhibition taxonomy. *Psychological Bulletin*, 126(2):220–246.

Nimchinsky, E. A., Vogt, B. A., Morrison, J. H., and Hof, P. R. (1995). Spindle neurons of the human anterior cingul. ate cortex. *Journal of Comparative Neurology*, 355(1):27–37.

Ninomiya, T., Dougherty, K., Godlove, D. C., Schall, J. D., and Maier, A. (2015). Microcircuitry of agranular frontal cortex: Contrasting laminar connectivity between occipital and frontal areas. *Journal of Neurophysiology*, 113(9):3242–55.

Nobre, A., Correa, A., and Coull, J. (2007). The hazards of time. *Current Opinion in Neurobiology*, 17(4):465–70.

Oemisch, M., Westendorff, S., Azimi, M., Hassani, S. A., Ardid, S., Tiesinga, P., and Womelsdorf, T. (2019). Feature-specific prediction errors and surprise across macaque fronto-striatal circuits. *Nature Communications*, 10(1):1–15.

Ogasawara, T., Nejime, M., Takada, M., and Matsumoto, M. (2018). Primate nigrostriatal dopamine system regulates saccadic response inhibition. *Neuron*, 100(6):1513–1526 e4.

Ohmae, S., Lu, X., Takahashi, T., Uchida, Y., and Kitazawa, S. (2008). Neuronal activity related to anticipated and elapsed time in macaque supplementary eye field. *Experimental Brain Research*, 184(4):593–8.

Ollman, R. T. (1973). *Simple reactions with random countermanding of the "go" signal*, pages 571–581. Academic Press, New York, NY, US.

O'Reilly, J. X., Schuffelgen, U., Cuell, S. F., Behrens, T. E., Mars, R. B., and Rushworth, M. F. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 110(38):E3660–9.

Ottes, F. P., Gisbergen, J. A. V., and Eggermont, J. J. (1986). Visuomotor fields of the superior colliculus: a quantitative model. *Vision Research*, 26(6):857–73.

Owens, D. F. and Kriegstein, A. R. (2002). Is there more to GABA than synaptic inhibition? *Nature Reviews Neuroscience*, 3(9):715–727.

O'Dowd, S., Schumacher, J., Burn, D. J., Bonanni, L., Onofrj, M., Thomas, A., and Taylor, J.-P. (2019). Fluctuating cognition in the lewy body dementias. *Brain*, 142(11):3338–3350.

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5(5):1–1.

Pani, P., Giarrocco, F., Bardella, G., Brunamonti, E., and Ferraina, S. (2022). Action-stopping models must consider the role of the dorsal premotor cortex. *Cortex*.

Papez, J. W. (1937). A proposed mechanism of emotion. *Archives of Neurology & Psychiatry*, 38(4):725–743.

Pardo, J. V., Pardo, P. J., Janer, K. W., and Raichle, M. E. (1990). The anterior cingulate cortex mediates processing selection in the stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences*, 87(1):256–259.

Pare, M. and Hanes, D. P. (2003). Controlled movement processing: superior colliculus activity associated with countermanded saccades. *Journal of Neuroscience*, 23(16):6480–9.

Pare, M. and Munoz, D. P. (1996). Saccadic reaction time in the monkey: advanced preparation of oculomotor programs is primarily responsible for express saccade occurrence. *Journal of Neurophysiology*, 76(6):3666–3681.

Parthasarathy, H. B., Schall, J. D., and Graybiel, A. M. (1992). Distributed but convergent ordering of corticostriatal projections: analysis of the frontal eye field and the supplementary eye field in the macaque monkey. *Journal of Neuroscience*, 12(11):4468–88.

Pasquereau, B. and Turner, R. S. (2017). A selective role for ventromedial subthalamic nucleus in inhibitory control. *eLife*, 6:e31627.

Passingham, R. (2009). How good is the macaque monkey model of the human brain? *Current Opinion in Neurobiology*, 19(1):6–11.

Passingham, R. E. (2007). *Two Cortical Systems for Directing Movement*. Novartis Foundation Symposia.

Passingham, R. E., Bengtsson, S. L., and Lau, H. C. (2010). Medial frontal cortex: from self-generated action to reflection on one's own performance. *Trends in Cognitive Science*, 14(1):16–21.

Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive and cognition interface. *Nature Reviews Neuroscience*, 2(6):417–424.

Peak, H. (1933). An evaluation of the concepts of reflex and voluntary action. *Psychological Review*, 40(1):71.

Penades, R., Catalan, R., Rubia, K., Andres, S., Salamero, M., and Gasto, C. (2007). Impaired response inhibition in obsessive compulsive disorder. *European Psychiatry*, 22(6):404–410.

Penfield, W. (1950). The supplementary motor area in the cerebral cortex of man. *Archiv für Psychiatrie und Nervenkrankheiten*, 185(6-7):670–674.

Perry, V. H. and Cowey, A. (1984). Retinal ganglion cells that project to the superior colliculus and pretectum in the macaque monkey. *Neuroscience*, 12(4):1125–37.

Pessoa, L., Kastner, S., and Ungerleider, L. G. (2003). Neuroimaging studies of attention: from modulation of sensory processing to top-down control. *Journal of Neuroscience*, 23(10):3990–8.

Peterson, N. N., Schroeder, C. E., and Arezzo, J. C. (1995). Neural generators of early cortical somatosensory evoked potentials in the awake monkey. *Electroencephalography and Clinical Neurophysiology*, 96(3):248–260.

Petter, E. A., Gershman, S. J., and Meck, W. H. (2018). Integrating models of interval timing and reinforcement learning. *Trends in Cognitive Sciences*, 22(10):911–922.

Polti, I., Martin, B., and Wassenhove, V. v. (2018). The effect of attention and working memory on the estimation of elapsed time. *Scientific Reports*, 8(1):6690.

Posner, M. I. and Dehaene, S. (1994). Attentional networks. *Trends in neurosciences*, 17(2):75–79.

Posner, M. I., Petersen, S. E., Fox, P. T., and Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, 240(4859):1627–1631.

Posner, M. I., Snyder, C. R., and Solso, R. (2004). Attention and cognitive control. *Cognitive Psychology: Key readings*, 205:55–85.

Postle, B. R., Brush, L. N., and Nick, A. M. (2004). Prefrontal cortex and the mediation of proactive interference in working memory. *Cognitive, Affective, & Behavioral Neuroscience*, 4(4):600–608.

Pouget, P., Emeric, E. E., Stuphorn, V., Reis, K., and Schall, J. D. (2005). Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *Journal of Neurophysiology*, 94(3):2086–2092.

Pouget, P., Logan, G. D., Palmeri, T. J., Boucher, L., Paré, M., and Schall, J. D. (2011). Neural basis of adaptive response time adjustment during saccade countermanding. *Journal of Neuroscience*, 31(35):12604–12.

Procyk, E., Tanaka, Y., and Joseph, J.-P. (2000). Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience*, 3(5):502–508.

Procyk, E., Wilson, C. R. E., Stoll, F. M., Faraut, M. C. M., Petrides, M., and Amiez, C. (2016). Midcingulate motor map and feedback detection: Converging data from humans and monkeys. *Cerebral Cortex*, 26(2):467–476.

Purcell, B. A. and Palmeri, T. J. (2017). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, 76(B):156–171.

Purcell, B. A., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2012a). From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *Journal of Neuroscience*, 32(10):3433–3446.

Purcell, B. A., Schall, J. D., and Woodman, G. F. (2013). On the origin of event-related potentials indexing covert attentional selection during visual search: timing of selection by macaque frontal eye field and event-related potentials during pop-out search. *Journal of Neurophysiology*, 109(2):557–569.

Purcell, B. A., Weigand, P. K., and Schall, J. D. (2012b). Supplementary eye field during visual search: salience, cognitive control, and performance monitoring. *Journal of Neuroscience*, 32(30):10273–10285.

Rabbitt, P. A. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology*, 71(2):264.

Rae, C. L., Hughes, L. E., Anderson, M. C., and Rowe, J. B. (2015). The prefrontal cortex achieves inhibitory control by facilitating subcortical motor pathway connectivity. *Journal of Neuroscience*, 35(2):786–794.

Ramakrishnan, A., Byun, Y. W., Rand, K., Pedersen, C. E., Lebedev, M. A., and Nicolelis, M. A. L. (2017). Cortical neurons multiplex reward-related signals along with sensory and motor information. *Proceedings of the National Academy of Sciences*, 114(24):E4841–E4850.

Ramautar, J., Kok, A., and Ridderinkhof, K. (2004). Effects of stop-signal probability in the stop-signal paradigm: the n2/p3 complex further validated. *Brain and Cognition*, 56(2):234–252.

Rao, H. M., Mayo, J. P., and Sommer, M. A. (2016). Circuits for presaccadic visual remapping. *Journal of Neurophysiology*, 116(6):2624–2636.

Rapan, L., Froudist-Walsh, S., Niu, M., Xu, T., Funck, T., Zilles, K., and Palomero-Gallagher, N. (2021). Multimodal 3D atlas of the macaque monkey motor and premotor cortex. *NeuroImage*, 226:117574.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59.

Ratcliff, R. and Smith, P. L. (2010). Perceptual discrimination in static and dynamic noise: the temporal relation between perceptual encoding and decision making. *Journal of Experimental Psychology: General*, 139(1):70–94.

Ratcliff, R., Smith, P. L., Brown, S. D., and McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Science*, 20(4):260–281.

Raud, L., Westerhausen, R., Dooley, N., and Huster, R. J. (2020). Differences in unity: The go/no-go and stop signal tasks rely on different mechanisms. *NeuroImage*, 210:116582.

Ray, N., Jenkinson, N., Brittain, J., Holland, P., Joint, C., Nandi, D., Bain, P., Yousif, N., Green, A., Stein, J., et al. (2009). The role of the subthalamic nucleus in response inhibition: evidence from deep brain stimulation for parkinson's disease. *Neuropsychologia*, 47(13):2828–2834.

Redgrave, P. and Gurney, K. (2006). The short-latency dopamine signal: A role in discovering novel actions? *Nature Reviews Neuroscience*, 7(12):967–75.

Reinhart, R. M., Heitz, R. P., Purcell, B. A., Weigand, P. K., Schall, J. D., and Woodman, G. F. (2012a). Homologous mechanisms of visuospatial working memory maintenance in macaque and human: properties and sources. *Journal of Neuroscience*, 32(22):7711–7722.

Reinhart, R. M. G., Carlisle, N. B., Kang, M.-S., and Woodman, G. F. (2012b). Event-related potentials elicited by errors during the stop-signal task. II: Human effector-specific error responses. *Journal of Neurophysiology*, 107(10):2794–2807.

Reppert, T. R., Servant, M., Heitz, R. P., and Schall, J. D. (2018). Neural mechanisms of speed-accuracy tradeoff of visual search: saccade vigor, the origin of targeting errors, and comparison of the superior colliculus and frontal eye field. *Journal of Neurophysiology*, 120(7):372–384.

Rey-Mermet, A. and Gade, M. (2018). Inhibition in aging: What is preserved? What declines? A meta-analysis. *Psychonomic Bulletin & Review*, 25(5):1695–1716.

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., and Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, 306(5695):443–447.

Riehle, A. and Requin, J. (1989). Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *Journal of Neurophysiology*, 61(3):534–49.

Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–90.

Rizzolatti, G., Luppino, G., and Matelli, M. (1998). The organization of the cortical motor system: new concepts. *Electroencephalography and Clinical Neurophysiology*, 106(4):283–96.

Robinson, D. (1964). The mechanics of human saccadic eye movement. *Journal of Physiology*, 174(2):245.

Robinson, D. A. (1972). Eye movements evoked by collicular stimulation in the alert monkey. *Vision Research*, 12(11):1795–808.

Rodriguez-Fornells, A., Kurzbuch, A. R., and Munte, T. F. (2002). Time course of error detection and correction in humans: neurophysiological evidence. *Journal of Neuroscience*, 22(22):9990–6.

Ros, T., Munneke, M., Parkinson, L., and Gruzelier, J. (2014). Neurofeedback facilitation of implicit motor learning. *Biological Psychology*, 95:54–58.

Roy, A., Choudhury, S., Basu, P., Baker, M. R., Baker, S. N., and Kumar, H. (2020). Stop signal reaction time measured with a portable device validates optimum stn-dbs programming. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, 13(6):1609–1611.

Rushworth, M., Hadland, K., Gaffan, D., and Passingham, R. (2003). The effect of cingulate cortex lesions on task switching and working memory. *Journal of Cognitive Neuroscience*, 15(3):338–353.

Rushworth, M., Walton, M. E., Kennerley, S. W., and Bannerman, D. (2004). Action sets and decisions in the medial frontal cortex. *Trends in Cognitive Sciences*, 8(9):410–417.

Saint-Cyr, J. A., Ungerleider, L. G., and Desimone, R. (1990). Organization of visual cortical inputs to the striatum and subsequent outputs to the pallido-nigral complex in the monkey. *Journal of Comparative Neurology*, 298(2):129–56.

Saito, N., Mushiake, H., Sakamoto, K., Itoyama, Y., and Tanji, J. (2005). Representation of immediate and final behavioral goals in the monkey prefrontal cortex during an instructed delay period. *Cerebral Cortex*, 15(10):1535–1546.

Sajad, A., Errington, S. P., and Schall, J. D. (2022). Functional architecture of executive control and associated event-related potentials in macaques. *Nature Communications*, 13(1):1–19.

Sajad, A., Godlove, D. C., and Schall, J. D. (2019). Cortical microcircuitry of performance monitoring. *Nature Neuroscience*, 22(2):265–274.

Sakai, K., Hikosaka, O., Miyauchi, S., Sasaki, Y., Fujimaki, N., and Putz, B. (1999). Presupplementary motor area activation during sequence learning reflects visuo-motor association. *Journal of Neuroscience*, 19(10):RC1.

Sallet, J., Mars, R. B., Noonan, M. P., Neubert, F.-X., Jbabdi, S., O'Reilly, J. X., Filippini, N., Thomas, A. G., and Rushworth, M. F. (2013). The organization of dorsal frontal cortex in humans and macaques. *Journal of Neuroscience*, 33(30):12255–12274.

Sandhaeger, F., Nicolai, C. v., Miller, E. K., and Siegel, M. (2019). Monkey EEG links neuronal color and motion information across species and scales. *eLife*, 8:e45645.

Saslow, M. G. (1967). Latency for saccadic eye movement. *Journal of the Optical Society of America*, 57(8):1030–3.

Sauerbier, A., Jenner, P., Todorova, A., and Chaudhuri, K. R. (2016). Non motor subtypes and parkinson's disease. *Parkinsonism & Related Disorders*, 22:S41–S46.

Scangos, K. W. and Stuphorn, V. (2010). Medial frontal cortex motivates but does not control movement initiation in the countermanding task. *Journal of Neuroscience*, 30(5):1968–82.

Schachar, R. and Logan, G. D. (1990). Impulsivity and inhibitory control in normal development and childhood psychopathology. *Developmental Psychology*, 26(5):710.

Schachar, R., Logan, G. D., Robaey, P., Chen, S., Ickowicz, A., and Barr, C. (2007). Restraint and cancellation: multiple inhibition deficits in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology*, 35(2):229–238.

Schachar, R. J., Chen, S., Logan, G. D., Ornstein, T. J., Crosbie, J., Ickowicz, A., and Pakulak, A. (2004). Evidence for an error monitoring deficit in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology*, 32(3):285–293.

Schaeffer, D. J., Hori, Y., Gilbert, K. M., Gati, J. S., Menon, R. S., and Everling, S. (2020). Divergence of rodent and primate medial frontal cortex functional connectivity. *Proceedings of the National Academy of Sciences*, 117(35):21681–21689.

Schall, J. (1988). Saccade latency and preparatory neuronal activity in the supplementary and frontal eye fields. In *Society of Neuroscience Abstract*, volume 159.

Schall, J. D. (1991a). Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *Journal of Neurophysiology*, 66(2):559–79.

Schall, J. D. (1991b). Neuronal activity related to visually guided saccadic eye movements in the supplementary motor area of rhesus monkeys. *Journal of Neurophysiology*, 66(2):530–558.

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience*, 2(1):33–42.

Schall, J. D. (2019). Accumulators, neurons, and response time. *Trends in Neurosciences*, 42(12):848–860.

Schall, J. D. and Boucher, L. (2007). Executive control of gaze by the frontal lobes. *Cognitive Affective Behavioral Neuroscience*, 7(4):396–412.

Schall, J. D. and Emeric, E. E. (2010). Conflict in cingulate cortex function between humans and macaque monkeys: More apparent than real. *Brain, Behavior and Evolution*, 75(4):237–8.

Schall, J. D. and Hanes, D. P. (1993). Neural basis of saccade target selection in frontal eye field during visual search. *Nature*, 366(6454):467–469.

Schall, J. D., Morel, A., King, D. J., and Bullier, J. (1995). Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. *Journal of Neuroscience*, 15(6):4464–4487.

Schall, J. D., Stuphorn, V., and Brown, J. W. (2002). Monitoring and control of action by the frontal lobes. *Neuron*, 36(2):309–322.

Schaum, M., Pinzuti, E., Sebastian, A., Lieb, K., Fries, P., Mobascher, A., Jung, P., Wibral, M., and Tuscher, O. (2021). Right inferior frontal gyrus implements motor inhibitory control via beta-band oscillations in humans. *eLife*, 10:e61679.

Schiller, P., Haushofer, J., and Kendall, G. (2004). An examination of the variables that affect express saccade generation. *Visual Neuroscience*, 21(2):119–127.

Schiller, P. H. and Chou, I.-h. (1998). The effects of frontal eye field and dorsomedial frontal cortex lesions on visually guided eye movements. *Nature Neuroscience*, 1(3):248–253.

Schiller, P. H., True, S. D., and Conway, J. L. (1980). Deficits in eye movements following frontal eye-field and superior colliculus ablations. *Journal of Neurophysiology*, 44(6):1175–1189.

Schlag, J. and Schlag-Rey, M. (1984). Visuomotor functions of central thalamus in monkey. II. Unit activity related to visual events, targeting, and fixation. *Journal of Neurophysiology*, 51(6):1175–95.

Schlag, J. and Schlag-Rey, M. (1986). Role of the central thalamus in gaze control. *Prog Brain Research*, 64:191–201.

Schlag, J. and Schlag-Rey, M. (1987). Evidence for a supplementary eye field. *Journal of Neurophysiology*, 57(1):179–200.

Schmidgen, H. (2002). Of frogs and men: the origins of psychophysiological time experiments, 1850-1865. *Endeavour*, 26(4):142–8.

Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., and Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature Neuroscience*, 16(8):1118–24.

Schneider, D. W. and Logan, G. D. (2009). Selecting a response in task switching: testing a model of compound cue retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(1):122–36.

Schultz, W., Carelli, R. M., and Wightman, R. M. (2015). Phasic dopamine signals: From subjective reward value to formal economic utility. *Current Opinion in Behavioral Sciences*, 5:147–154.

Scudder, C. A., Kaneko, C. S., and Fuchs, A. F. (2002). The brainstem burst generator for saccadic eye movements: a modern synthesis. *Experimental Brain Research*, 142(4):439–62.

Servant, M., Tillman, G., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2019). Neurally constrained modeling of speed-accuracy tradeoff during visual search: gated accumulation of modulated evidence. *Journal of Neurophysiology*, 121(4):1300–1314.

Sharp, D. J., Bonnelle, V., Boissezon, X. D., Beckmann, C. F., James, S. G., Patel, M. C., and Mehta, M. A. (2010). Distinct frontal systems for response inhibition, attentional capture, and error processing. *Proceedings of the National Academy of Sciences*, 107(13):6106–6111.

Shenhav, A., Cohen, J. D., and Botvinick, M. M. (2016a). Dorsal anterior cingulate cortex and the value of control. *Nature Neuroscience*, 19(10):1286–91.

Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., and Botvinick, M. M. (2016b). Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience*, 40(1):1–26.

Shenhav, A., Straccia, M. A., Cohen, J. D., and Botvinick, M. M. (2014). Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value. *Nature Neuroscience*, 17(9):1249–1254.

Shenoy, K. V., Sahani, M., Churchland, M. M., et al. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annual Review of Neuroscience*, 36(1):337–359.

Sherman, M. A., Lee, S., Law, R., Haegens, S., Thorn, C. A., Hämäläinen, M. S., Moore, C. I., and Jones, S. R. (2016). Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. *Proceedings of the National Academy of Sciences*, 113(33):E4885–E4894.

Shidara, M. and Richmond, B. J. (2002). Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science*, 296(5573):1709–1711.

Shiffrin, R. M. and Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84(2):127.

Shima, K. and Tanji, J. (2000). Neuronal activity in the supplementary and presupplementary motor areas for temporal organization of multiple movements. *Journal of Neurophysiology*, 84(4):2148–60.

Shin, H., Law, R., Tsutsui, S., Moore, C. I., and Jones, S. R. (2017). The rate of transient beta frequency events predicts behavior across tasks and species. *eLife*, 6:e29086.

Shin, S. and Sommer, M. A. (2010). Activity of neurons in monkey globus pallidus during oculomotor behavior compared with that in substantia nigra pars reticulata. *Journal of Neurophysiology*, 103(4):1874–87.

Shipp, S. (2005). The importance of being agranular: A comparative account of visual and motor cortex. *Proceedings of the Royal Society B: Biological Sciences*, 360(1456):797–814.

Shook, B. L., Schlag-Rey, M., and Schlag, J. (1990). Primate supplementary eye field: I. Comparative aspects of mesencephalic and pontine connections. *Journal of Comparative Neurology*, 301(4):618–42.

Sitaram, R., Ros, T., Stoeckel, L., Haller, S., Scharnowski, F., Lewis-Peacock, J., Weiskopf, N., Blefari, M. L., Rana, M., Oblak, E., et al. (2017). Closed-loop brain training: the science of neurofeedback. *Nature Reviews Neuroscience*, 18(2):86–100.

Skippen, P., Fulham, W. R., Michie, P. T., Matzke, D., Heathcote, A., and Karayanidis, F. (2020). Reconsidering electrophysiological markers of response inhibition in light of trigger failures in the stop-signal task. *Psychophysiology*, (53):e13619.

Skippen, P., Matzke, D., Heathcote, A., Fulham, W. R., Michie, P., and Karayanidis, F. (2019). Reliability of triggering inhibitory process is a better predictor of impulsivity than SSRT. *Acta Psychologica*, 192:104–117.

Slater-Hammel, A. T. (1960). Reliability, accuracy, and refractoriness of a transit reaction. *Research Quarterly. American Association for Health, Physical Education and Recreation*, 31(2):217–228.

Smith, J. L., Mattick, R. P., Jamadar, S. D., and Iredale, J. M. (2014). Deficits in behavioural inhibition in substance abuse and addiction: a meta-analysis. *Drug and Alcohol Dependence*, 145:1–33.

Smith, J. L., Smith, E. A., Provost, A. L., and Heathcote, A. (2010). Sequence effects support the conflict theory of N2 and P3 in the Go/NoGo task. *International Journal of Psychophysiology*, 75(3):217–26.

Smith, M. A. and Sommer, M. A. (2013). Spatial and temporal scales of neuronal correlation in visual area v4. *Journal of Neuroscience*, 33(12):5422–5432.

Snyder, L. H., Batista, A. P., and Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature*, 386(6621):167–70.

So, N. Y. and Stuphorn, V. (2010). Supplementary eye field encodes option and action value for saccades with variable reward. *Journal of Neurophysiology*, 104(5):2634–53.

Sommer, M. A. and Tehovnik, E. J. (1997). Reversible inactivation of macaque frontal eye field. *Experimental Brain Research*, 116(2):229–249.

Sommer, M. A. and Wurtz, R. H. (2002). A pathway in primate brain for internal monitoring of movements. *Science*, 296(5572):1480–2.

Sommer, M. A. and Wurtz, R. H. (2004a). What the brain stem tells the frontal cortex. I. Oculomotor signals sent from superior colliculus to frontal eye field via mediodorsal thalamus. *Journal of Neurophysiology*, 91(3):1381–402.

Sommer, M. A. and Wurtz, R. H. (2004b). What the brain stem tells the frontal cortex. II. Role of the SC-MD-FEF pathway in corollary discharge. *Journal of Neurophysiology*, 91(3):1403–23.

Sommer, M. A. and Wurtz, R. H. (2006). Influence of the thalamus on spatial visual processing in frontal cortex. *Nature*, 444(7117):374–7.

Sotero, R. C., Bortel, A., Naaman, S., Mocanu, V. M., Kropf, P., Villeneuve, M. Y., and Shmuel, A. (2015). Laminar distribution of phase-amplitude coupling of spontaneous current sources and sinks. *Frontiers in Neuroscience*, 9:454.

Sparks, D. L., Holland, R., and Guthrie, B. L. (1976). Size and distribution of movement fields in the monkey superior colliculus. *Brain Research*, 113(1):21–34.

Stahl, J. and 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. *Clinical Neurophysiology*, 118(3):581–596.

Stanford, T. R., Freedman, E. G., and Sparks, D. L. (1996). Site and parameters of microstimulation: evidence for independent effects on the properties of saccades evoked from the primate superior colliculus. *Journal of Neurophysiology*, 76(5):3360–81.

Stanton, G., Goldberg, M., and Bruce, C. (1988). Frontal eye field efferents in the macaque monkey: Ii. topography of terminal fields in midbrain and pons. *Journal of Comparative Neurology*, 271(4):493–506.

Starkweather, C. K., Babayan, B. M., Uchida, N., and Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, 20(4):581–589.

Stevenson, S. A., Elsley, J. K., and Corneil, B. D. (2009). A "gap effect" on stop signal reaction times in a human saccadic countermanding task. *Journal of Neurophysiology*, 101(2):580–590.

Stoll, F. M., Fontanier, V., and Procyk, E. (2016). Specific frontal neural dynamics contribute to decisions to check. *Nature Communications*, 7(1):11990.

Stuphorn, V., Brown, J. W., and Schall, J. D. (2010). Role of supplementary eye field in saccade initiation: executive, not direct, control. *Journal of Neurophysiology*, 103(2):801–16.

Stuphorn, V. and Schall, J. D. (2006). Executive control of countermanding saccades by the supplementary eye field. *Nature Neuroscience*, 9(7):925–31.

Stuphorn, V., Taylor, T. L., and Schall, J. D. (2000). Performance monitoring by the supplementary eye field. *Nature*, 408(6814):857–60.

Subramanian, D., Alers, A., and Sommer, M. A. (2019). Corollary discharge for action and cognition. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 4(9):782–790.

Sumner, P., Nachev, P., Morris, P., Peters, A. M., Jackson, S. R., Kennard, C., and Husain, M. (2007). Human medial frontal cortex mediates unconscious inhibition of voluntary action. *Neuron*, 54(5):697–711.

Swann, N., Tandon, N., Canolty, R., Ellmore, T. M., McEvoy, L. K., Dreyer, S., DiSano, M., and Aron, A. R. (2009). Intracranial EEG reveals a time- and frequency-specific role for the right inferior frontal gyrus and primary motor cortex in stopping initiated responses. *Journal of Neuroscience*, 29(40):12675–85.

Swann, N. C., Cai, W., Conner, C. R., Pieters, T. A., Claffey, M. P., George, J. S., Aron, A. R., and 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(3):2860–70.

Swick, D., Ashley, V., and Turken, U. (2011). Are the neural correlates of stopping and not going identical? Quantitative meta-analysis of two response inhibition tasks. *NeuroImage*, 56(3):1655–1665.

Takada, M., Tokuno, H., Nambu, A., and Inase, M. (1998). Corticostriatal input zones from the supplementary motor area overlap those from the contra-rather than ipsilateral primary motor cortex. *Brain research*, 791(1-2):335–340.

Takikawa, Y., Kawagoe, R., Itoh, H., Nakahara, H., and Hikosaka, O. (2002). Modulation of saccadic eye movements by predicted reward outcome. *Experimental Brain Research*, 142(2):284–291.

Tallot, L. and Doyere, V. (2020). Neural encoding of time in the animal brain. *Neuroscience Biobehavioral Reviews*, 115:146–163.

Tanibuchi, I. and Goldman-Rakic, P. S. (2003). Dissociation of spatial-, object-, and sound-coding neurons in the mediodorsal nucleus of the primate thalamus. *Journal of Neurophysiology*, 89(2):1067–77.

Tanji, J. and Kurata, K. (1982). Comparison of movement-related activity in two cortical motor areas of primates. *Journal of Neurophysiology*, 48(3):633–53.

Tanji, J. and Shima, K. (1994). Role for supplementary motor area cells in planning several movements ahead. *Nature*, 371(6496):413–6.

Teller, D. Y. (1984). Linking propositions. *Vision Research*, 24(10):1233–46.

Thakkar, K. N., Heiligenberg, F. M. v. d., Kahn, R. S., and Neggers, S. F. (2014). Frontal-subcortical circuits involved in reactive control and monitoring of gaze. *Journal of Neuroscience*, 34(26):8918–8929.

Thakkar, K. N. and Rolfs, M. (2019). Disrupted corollary discharge in schizophrenia: Evidence from the oculomotor system. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 4(9):773–781.

Thakkar, K. N., Schall, J. D., Boucher, L., Logan, G. D., and Park, S. (2011). Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biol Psychiatry*, 69(1):55–62.

Thakkar, K. N., Schall, J. D., Logan, G. D., and Park, S. (2015). Cognitive control of gaze in bipolar disorder and schizophrenia. *Psychiatry Research*, 225(3):254–262.

Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica*, 135(2):77–99.

Thomaschke, R., Kiesel, A., and Hoffmann, J. (2011). Response specific temporal expectancy: Evidence from a variable foreperiod paradigm. *Attention, Perception, & Psychophysics*, 73(7):2309.

Thompson, K. G., Bichot, N. P., and Schall, J. D. (1997). Dissociation of visual discrimination from saccade programming in macaque frontal eye field. *Journal of Neurophysiology*, 77(2):1046–1050.

Thompson, K. G., Hanes, D. P., Bichot, N. P., and Schall, J. D. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *Journal of Neurophysiology*, 76(6):4040–4055.

Tiego, J., Testa, R., Bellgrove, M. A., Pantelis, C., and Whittle, S. (2018). A hierarchical model of inhibitory control. *Frontiers in Psychology*, 9:1339.

Tinkhauser, G., Pogosyan, A., Tan, H., Herz, D. M., Kühn, A. A., and Brown, P. (2017). Beta burst dynamics in parkinson's disease off and on dopaminergic medication. *Brain*, 140(11):2968–2981.

Tremblay, R., Lee, S., and Rudy, B. (2016). GABAergic interneurons in the neocortex: From cellular properties to circuits. *Neuron*, 91(2):260–92.

Trujillo, L. T. and Allen, J. J. (2007). Theta EEG dynamics of the error-related negativity. *Clinical Neurophysiology*, 118(3):645–668.

Ullsperger, M., Danielmeier, C., and Jocham, G. (2014). Neurophysiology of performance monitoring and adaptive behavior. *Physiological Reviews*, 94(1):35–79.

van den Wildenberg, W. P., Burle, B., Vidal, F., Molen, M. W. v. d., Ridderinkhof, K. R., and Hasbroucq, T. (2010). Mechanisms and dynamics of cortical motor inhibition in the stop-signal paradigm: a TMS study. *Journal of Cognitive Neuroscience*, 22(2):225–39.

van den Wildenberg, W. P., van Boxtel, G. J., van der Molen, M. W., Bosch, D. A., Speelman, J. D., and Brunia, C. H. (2006). Stimulation of the subthalamic region facilitates the selection and inhibition of motor responses in parkinson's disease. *Journal of Cognitive Neuroscience*, 18(4):626–636.

van Veen, V. and Carter, C. S. (2002). The anterior cingulate as a conflict monitor: fMRI and ERP studies. *Physiology & Behavior*, 77(4-5):477–482.

Varela, F., Lachaux, J.-P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2(4):229–239.

Verbruggen, F., Aron, A. R., Band, G. P., Beste, C., Bissett, P. G., Brockett, A. T., Brown, J. W., Chamberlain, S. R., Chambers, C. D., and Colonius, H. (2019a). A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *elife*, 8:e46323.

Verbruggen, F., Aron, A. R., Band, G. P., Beste, C., Bissett, P. G., Brockett, A. T., Brown, J. W., Chamberlain, S. R., Chambers, C. D., Colonius, H., Colzato, L. S., Corneil, B. D., Coxon, J. P., Dupuis, A., Eagle, D. M., Garavan, H., Greenhouse, I., Heathcote, A., Huster, R. J., Jahfari, S., Kenemans, J. L., Leunissen, I., Li, C. R., Logan, G. D., Matzke, D., Morein-Zamir, S., Murthy, A., Paré, M., Poldrack, R. A., Ridderinkhof, K. R., Robbins, T. W., Roesch, M., Rubia, K., Schachar, R. J., Schall, J. D., Stock, A. K., Swann, N. C., Thakkar, K. N., Molen, M. W. v. d., Vermeylen, L., Vink, M., Wessel, J. R., Whelan, R., Zandbelt, B. B., and Boehler, C. N. (2019b). A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife*, 8:e46323.

Verbruggen, F., Chambers, C. D., and Logan, G. D. (2013). Fictitious inhibitory differences: how skewness and slowing distort the estimation of stopping latencies. *Psychological Science*, 24(3):352–362.

Verbruggen, F. and Logan, G. D. (2008). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences*, 12(11):418–424.

Verbruggen, F. and Logan, G. D. (2009a). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience & Biobehavioral Reviews*, 33(5):647–661.

Verbruggen, F. and Logan, G. D. (2009b). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3):835–854.

Verbruggen, F., Logan, G. D., Liefooghe, B., and Vandierendonck, A. (2008). Short-term aftereffects of response inhibition: repetition priming or between-trial control adjustments? *Journal of Experimental Psychology: Human Perception and Performance*, 34(2):413.

Vince, M. A. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology*, 38(3):149.

Vogt, B. A. (2016). Midcingulate cortex: Structure, connections, homologies, functions and diseases. *Journal of Chemical Neuroanatomy*, 74:28–46.

Vogt, B. A., Berger, G. R., and Derbyshire, S. W. G. (2003). Structural and functional dichotomy of human midcingulate cortex. *European Journal of Neuroscience*, 18(11):3134–3144.

Vogt, B. A., Nimchinsky, E. A., Vogt, L. J., and Hof, P. R. (1995). Human cingulate cortex: Surface features, flat maps, and cytoarchitecture. *Journal of Comparative Neurology*, 359(3):490–506.

Vogt, B. A., Vogt, L., Farber, N. B., and Bush, G. (2005). Architecture and neurocytology of monkey cingulate gyrus. *Journal of Comparative Neurology*, 485(3):218–239.

Vogt, C. and Vogt, O. (1919). *Allgemeine ergebnisse unserer hirnforschung*, volume 25. JA Barth.

Voloh, B., Oemisch, M., and Womelsdorf, T. (2020). Phase of firing coding of learning variables across the fronto-striatal network during feature-based learning. *Nature Communications*, 11(1):1–16.

Vosskuhl, J., Strüber, D., and Herrmann, C. S. (2018). Non-invasive brain stimulation: a paradigm shift in understanding brain oscillations. *Frontiers in Human Neuroscience*, 12:211.

Vullings, C. and Madelain, L. (2018). Control of saccadic latency in a dynamic environment: allocation of saccades in time follows the matching law. *Journal of Neurophysiology*, 119(2):413–421.

Vullings, C. and Madelain, L. (2019). Discriminative control of saccade latencies. *Journal of Vision*, 19(3):16.

Waller, D. A., Hazeltine, E., and Wessel, J. R. (2021). Common neural processes during action-stopping and infrequent stimulus detection: The frontocentral P3 as an index of generic motor inhibition. *International Journal of Psychophysiology*, 163:11–21.

Walton, M. M. G. and Gandhi, N. J. (2006). Behavioral evaluation of movement cancellation. *Journal of Neurophysiology*, 96(4):2011–2024.

Wang, J., Narain, D., Hosseini, E. A., and Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience*, 21(1):102–110.

Wang, X.-J. (2020). Macroscopic gradients of synaptic excitation and inhibition in the neocortex. *Nature Reviews Neuroscience*, 21(3):169–178.

Wang, X. J., Tegner, J., Constantinidis, C., and Goldman-Rakic, P. S. (2004). Division of labor among distinct subtypes of inhibitory neurons in a cortical microcircuit of working memory. *Proceedings of the National Academy of Sciences*, 101(5):1368–73.

Watanabe, Y. and Funahashi, S. (2004). Neuronal activity throughout the primate mediodorsal nucleus of the thalamus during oculomotor delayed-responses. I. Cue-, delay-, and response-period activity. *Journal of Neurophysiology*, 92(3):1738–55.

Wattiez, N., Poitou, T., Rivaud-Péchoux, S., and Pouget, P. (2016). Evidence for spatial tuning of movement inhibition. *Experimental Brain Research*, 234(7):1957–1966.

Wessel, J. R. (2020). beta-bursts reveal the trial-to-trial dynamics of movement initiation and cancellation. *Journal of Neuroscience*, 40(2):411–423.

Wessel, J. R. and Aron, A. R. (2013). Unexpected events induce motor slowing via a brain mechanism for action-stopping with global suppressive effects. *Journal of Neuroscience*, 33(47):18481–91.

Wessel, J. R. and Aron, A. R. (2015). It's not too late: the onset of the frontocentral P3 indexes successful response inhibition in the stop-signal paradigm. *Psychophysiology*, 52(4):472–80.

Wessel, J. R., Ghahremani, A., Udupa, K., Saha, U., Kalia, S. K., Hodaie, M., Lozano, A. M., Aron, A. R., and Chen, R. (2016). Stop-related subthalamic beta activity indexes global motor suppression in Parkinson's disease. *Movement Disorders*, 31(12):1846–1853.

Wessel, J. R., Reynoso, H. S., and Aron, A. R. (2013). Saccade suppression exerts global effects on the motor system. *Journal of Neurophysiology*, 110(4):883–90.

West, R. and Travers, S. (2008). Tracking the temporal dynamics of updating cognitive control: An examination of error processing. *Cerebral Cortex*, 18(5):1112–1124.

Westerberg, J. A., Cox, M. A., Dougherty, K., and Maier, A. (2019). V1 microcircuit dynamics: altered signal propagation suggests intracortical origins for adaptation in response to visual repetition. *Journal of Neurophysiology*, 121(5):1938–1952.

Westerberg, J. A., Maier, A., Woodman, G. F., and Schall, J. D. (2020). Performance monitoring during visual priming. *Journal of Cognitive Neuroscience*, 32(3):515–526.

Westerberg, J. A., Schall, M. S., Maier, A., Woodman, G. F., and Schall, J. D. (2022). Laminar microcircuitry of visual cortex producing attention-associated electric fields. *eLife*, 11:e72139.

Whittingstall, K. and Logothetis, N. K. (2009). Frequency-band coupling in surface EEG reflects spiking activity in monkey visual cortex. *Neuron*, 64(2):281–289.

Wichmann, T., Bergman, H., and DeLong, M. R. (1994). The primate subthalamic nucleus. I. Functional properties in intact animals. *Journal of Neurophysiology*, 72(2):494–506.

Wiecki, T. V. and Frank, M. J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychological Review*, 120(2):329–55.

Williams, S. M. and Goldman-Rakic, P. S. (1998). Widespread origin of the primate mesofrontal dopamine system. *Cerebral Cortex*, 8(4):321–45.

Wise, S. P. (1996). Corticospinal efferents of the supplementary sensorimotor area in relation to the primary motor area. *Advances in Neurology*, 70:57–69.

Wittmann, M. K., Fouragnan, E., Folloni, D., Klein-Flügge, M. C., Chau, B. K. H., Khamassi, M., and Rushworth, M. F. S. (2020). Global reward state affects learning and activity in raphe nucleus and anterior insula in monkeys. *Nature Communications*, 11(1):3771.

Woodman, G. F. (2012). *Homologues of human ERP components in nonhuman primates*, pages 611–626. Oxford University Press, New York, NY, US.

Wurtz, R. H. and Goldberg, M. E. (1972). Activity of superior colliculus in behaving monkey. 3. Cells discharging before eye movements. *Journal of Neurophysiology*, 35(4):575–86.

Wyder, M. T., Massoglia, D. P., and Stanford, T. R. (2003). Quantitative assessment of the timing and tuning of visual-related, saccade-related, and delay period activity in primate central thalamus. *Journal of Neurophysiology*, 90(3):2029–52.

Xing, D., Yeh, C.-I., Burns, S., and Shapley, R. M. (2012). Laminar analysis of visually evoked activity in the primary visual cortex. *Proceedings of the National Academy of Sciences*, 109(34):13871–13876.

Yamamoto, S., Monosov, I. E., Yasuda, M., and Hikosaka, O. (2012). What and where information in the caudate tail guides saccades to visual objects. *Journal of Neuroscience*, 32(32):11005–16.

Yarbus, A. L. (1967). *Eye movements during perception of complex objects*. Springer, New York, NY, United States.

Yelnik, J. and Percheron, G. (1979). Subthalamic neurons in primates: a quantitative and comparative analysis. *Neuroscience*, 4(11):1717–43.

Yeterian, E. H. and Hoesen, G. W. V. (1978). Cortico-striate projections in the rhesus monkey: the organization of certain cortico-caudate connections. *Brain Research*, 139(1):43–63.

Yoshida, S., Nambu, A., and Jinnai, K. (1993). The distribution of the globus pallidus neurons with input from various cortical areas in the monkeys. *Brain Research*, 611(1):170–4.

Yu, A. J., Dayan, P., and Cohen, J. D. (2009). Dynamics of attentional selection under conflict: toward a rational bayesian account. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3):700.

Yu, Y., Sanabria, D. E., Wang, J., Hendrix, C. M., Zhang, J., Nebeck, S. D., Amundson, A. M., Busby, Z. B., Bauer, D. L., Johnson, M. D., et al. (2021). Parkinsonism alters beta burst dynamics across the basal ganglia–motor cortical network. *Journal of Neuroscience*, 41(10):2274–2286.

Zhang, K., Chen, C. D., and Monosov, I. E. (2019). Novelty, salience, and surprise timing are signaled by neurons in the basal forebrain. *Current Biology*, 29(1):134–142 e3.

Zimnik, A. J., Lara, A. H., and Churchland, M. M. (2019). Perturbation of macaque supplementary motor area produces context-independent changes in the probability of movement initiation. *Journal of Neuroscience*, 39(17):3217–3233.