Neural basis of self-control

Ka Eun Lee[1,3], Jacob Elsey[1,3], Jaewon Hwang[2,3,4], Erik E. Emeric[2,3], Veit Stuphorn[1,2,3,*]

Author affiliations:

[1]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD 21218-2685, USA

[2]Department of Neuroscience, Johns Hopkins University School of Medicine and Zanvyl Krieger Mind/Brain Institute, Baltimore, MD 21218-2685, USA

[3]The Zanvyl Krieger Mind/Brain Institute, Baltimore, MD 21218-2685, USA

[4]Current address: Section on Neurobiology of Learning and Memory, Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, MD, USA

[*]Corresponding author.

Corresponding author and lead contact:
Veit Stuphorn
Johns Hopkins University
338 Krieger Hall
3400 N. Charles St.
Baltimore, MD 21218
USA
Tel: (410) 516-7964
Fax: (410) 516-8648
Email: veit@jhu.edu

Abstract

Humans are often tempted by small, but immediate rewards that seem more attractive than a larger, delayed reward, even when such choices are clearly against one's own best interest. Resisting this temptation and remaining committed to a long-term goal requires self-control, the ability to inhibit self-defeating behavior. Here, we report evidence that neurons in the Supplementary Eye Field encode self-control of oculomotor behavior. We designed a modified version of the intertemporal choice task that distinguishes states of high and low self-control. Supplementary Eye Field neurons encodes distinct levels of self-control during all stages of the trial, even before any choice targets or a later temptation were presented. Fluctuations in Supplementary Eye Field activity were predictive of the monkeys' behavioral response to temptation, suggesting Supplementary Eye Field is critical for self-control behaviors. A partially overlapping population predicted the initial choice, revealing a common neuronal mechanism underlying the ability to make prudent choices that are more beneficial in the long run, and the ability to adhere to those choices by exercising self-control and resisting temptation. Our findings suggest that Supplementary Eye Field is part of a neuronal circuit that underlies the capacity for self-control, which is crucial for maintaining and achieving long-term goals.

Introduction

The achievement of long-term goals often requires humans to exert self-control, foregoing small immediate pleasure to achieve larger rewards in the future. Examples span everyday situations such as dieting, or exercising, to those involving larger goals such as working overtime for a promotion instead of relaxing in front of a TV. A lack of self-control is associated with conduct disorders, drug addiction, and obesity, and it predicts early-onset criminal conduct. However, despite its importance in everyday life, the neuronal mechanisms of self-control remain elusive.

Self-control has traditionally been studied under the framework of intertemporal choice (1–3), which involves a choice between a delayed larger reward (L) and an immediate smaller reward (S). Choices in this task can be understood as the outcome of a cost-benefit evaluation of the offered rewards (4). The effect of self-control on choice is therefore confounded with that of other behavioral factors, such as effort and opportunity costs associated with delay or erroneous decisions due to stochastic value representations.

To dissociate self-control from the process of value-based choice, we designed a modified version of the intertemporal choice task that allows us to identify states of high and low self-control (**Figure 1A**). Four male rhesus monkeys (*macaca mulatta*) performed the self-control task (monkey A: 47 sessions; monkey I: 17 sessions; monkey G: 38 sessions; monkey K: 17 sessions). In each trial, they indicated their preferences between a smaller sooner reward (S) and a larger later reward (L), with explicitly cued amounts and delays (**Figure 1B**). Because this choice was binding in the vast majority of 'no temptation' trials, the monkeys had an incentive to indicate their true preferences. However, in a small and unpredictable subset of 'temptation' trials, the initial choice was non-binding and the monkeys were allowed to switch to the unchosen option. Such a preference reversal would be irrational from an economic point of view, because the reward was initially deemed inferior. Nevertheless, in some temptation trials where the monkey initially chooses L (the more prudent option in our setup), the monkeys were tempted by the initially unchosen S option and switched. Thus, the behavioral response to this temptation reveals the momentary level of self-control on that trial. If the monkey is not affected by temptation and maintains the initially chosen L reward, the monkey shows high self-control. In contrast, if the monkey gives in to temptation by switching to S, he shows low self-control. In this way, the same value-based decision (i.e., choosing L) can be associated with different levels of self-control, allowing us to identify neuronal activity associated with differing levels of self-control. Our findings suggest that SEF is part of a neuronal circuit underlying the capacity for self-control, which is crucial for maintaining and achieving prudent long-term goals.

We recorded from the supplementary eye field (SEF), a higher-order oculomotor area that has been shown to be involved in the executive control of value-based decisions in the oculomotor domain (**Figure 1C**) (5). SEF neurons were found to carry self-control signals from early in the trial and their activity predicted the monkeys' behavioral responses to temptation.

Results

Behavioral results

On each trial of the self-control task (**Figure 1**), the monkeys chose initially between a smaller sooner (S option) and a larger later reward (L option). In no-temptation trials, the unchosen option disappeared once the monkey made its choice. In temptation trials, the unchosen option remained available on the screen, indicating to the monkey that it was possible to switch to the initially unchosen option. Whether the monkeys switched or stayed determined their final choice.

The monkeys' behavioral response to temptation also defined self-control categorically. We only considered trials in which monkeys initially chose the L option indicating that it was preferred despite the associated cost. Those trials, where the monkeys overcame temptation and stayed were labelled as 'high self-control' (LL trials). On the other hand, trials in which monkeys failed to exercise self-control and switched to the S option were labelled as 'low self-control' (LS trials). All other trials (i.e., no temptation trials where the monkey chose the L [L trials] or S option [S trials], temptation trials where the monkey initially chose the S option and stayed [SS trials] or switched [SL trials]) were labelled as 'other', because the effect of self-control on choice in these trials were confounded with that of value-based preference (e.g., the subjective value of L can be sufficiently higher than that of the S option, thereby not requiring high self-control to choose L).

Overall, four monkeys completed 102,218 trials over 119 behavioral sessions (no temptation trials: 64,651; temptation trials: 37,567). We fitted a logistic regression function to the choice data and quantified the monkeys' attitude toward delayed rewards by computing the indifference delay (i.e., the L delay at which the monkey is indifferent between L and S options). All four monkeys discounted the value of options proportional to their delay (**Figure 2A**). For initial choice, monkeys should be unable to anticipate whether the unchosen alternative will remain available, as temptation trials occurred infrequently (36.8%) and unpredictably. Consequently, in theory, choice functions (and their indifference delays) for no temptation and temptation trials were not expected to exhibit significant differences. We found a small (absolute mean difference in indifference delay: 275ms), but significant (paired t-test: t=4.9796, p=1.9118e-06) difference in choice functions (**Figure 2D**). This was likely an artifact of our task design (see **Supplementary Figure 1**).

The final choice on temptation trials indicated a significant and consistent leftward shift in preferences towards shorter delays relative to the initial choice (absolute mean difference in indifference delay: 2,026ms; paired t-test: t=-11.8974, p=9.0434e-23). This shift indicates an asymmetrical influence of temptation on the monkeys' preferences. When the monkey initially chose the L option, they stayed in 48.7% of temptation trials (LL trials) and switched to the S option in 51.3% of trials (LS trials). In contrast, when the monkey initially chose the S option, they stayed in 82.6% of trials (SS trials) and switched to L in only 17.4% of trials (SL trials) (**Figure 2B**) . This leftward shift indicated an unidirectional influence of temptation on the monkeys' preferences, causing monkeys to more likely make LS (85.9%) than SL (14.1%) switches (Two-sample t-test t=15.1, p=1.4585e-36). Thus, monkeys were more likely to ultimately choose the S option when they were tempted.

This asymmetry in the propensity to switch for trials with initial L- and S-choices was not the result of a physical incapability of the monkeys to switch in both directions, as 78.6% of switches occurred before 500ms (median=312ms, mean=396ms, SD=270ms), the shortest S delay used in the task (**Figure 2F**). The

imbalance also indicated that most switches were not due to error correction, as error correction should not result in changes of preference. LS and SL switches should therefore occur equally often around the indifference delay and result in a steeper choice function. Neither effect was present in the choice functions for initial and final choice in temptation trials (**Figure 2A**).

Together, choice data revealed that monkeys made initial choices according to the delay-discounted subjective values of options, and choices were unidirectionally affected (i.e., switching from L to S) by the presence of temptation. This asymmetric effect indicates that monkeys need self-control to maintain delayed gratification in the face of temptation, but are not always successful.

SEF neurons encode distinct levels of self-control

We recorded from 310 SEF neurons in two monkeys while they performed the self-control task (monkey A: 208 neurons; monkey I: 102 neurons). We hypothesized that SEF neurons influence behavioral responses to temptation by carrying a signal that reflects the monkey's level of self-control.

We identified SEF neurons representing self-control and other behaviorally relevant variables using generalized linear models (GLM) (see **Methods**). We examined SEF activity in three time periods of interest (central fixation, target onset, before choice), aligned to three events (fixation onset, target onset, initial choice) in overlapping time windows (100-ms windows shifted by 25ms) (**Figure 3A**). For each time period, we found the set of variables that best explained each neuron's firing activity. A neuron was considered to robustly encode a variable in a given time bin only if it was included in the best fitting model for 3 consecutive time bins (i.e., present for at least 150ms, accounting for bin overlap).

We identified neurons encoding self-control in all 3 time periods (**Figure 3B**,**C)**. Neurons could have both a positive (i.e., higher firing rate for high self-control) or negative (i.e., higher firing rate for low self-control) correlation with self-control (**Table 1**). The proportion of each type was similar in each time period (one proportion z test: central fixation: z=-1.3270, p=0.1845; after target: z=1.1209, p=0.2623; before choice: z=-0.2673, p=0.7893).

At fixation (-100 to 700ms around fixation), the majority of neurons (59%; 183/310) encoded self-control. On average, 16.8% of neurons (mean=52.2, SD=2.14 neurons) encoded self-control in each 100-ms time bin during fixation (**Figure 3C**). Choice targets were presented at least 731ms after fixation onset (mean=989ms, median=994ms, SD=146ms), Thus, self-control signals in this time period were unlikely to be related to saccade preparation. We next examined SEF activity in the time period between target onset and initial choice (mean choice reaction time=254ms, median=233ms, SD=80ms). In this time period, the S and L target options are presented, starting the decision process culminating in the initial choice. Because SEF activity could reflect both sensory responses to the targets or cognitive and motor-related processes involved in choice selection, we analyzed the activity both aligned on target onset and initial choice. First, we aligned the activity to target onset (-50 to 300ms). We found that the SEF carried various task-related signals (**Figure 3C**). Independent of these other signals, 39/310 neurons (12.6%) encoded self-control, 30 of which had also been identified to encode self-control during fixation. On average, 5.2% of neurons (mean=16.2 neurons, SD=6.9) encoded self-control after target onset in each 100-ms time bin, increasing slightly over this time period (**Figure 3C**). Next, we aligned the activity to initial choice (-300 to 50ms around initial choice). We found 56/310 neurons (18.1%) encoding self-control. On average, 7.45% of neurons (mean=23.7 neurons, SD=7.15) encoded self-control in each 100-ms time bin, with the proportion increasing slightly over time (**Figure 3C**). Ignoring the robustness criterion, the average length of self-

control signals in the SEF was 145ms (2.78 time bins) after fixation, 119ms (1.77 time bins) after target onset, and 120ms (1.8 time bins) before initial choice. Overall, a total of 202/310 neurons (65%) were identified to robustly encode self-control in at least one of the three time periods.

In addition to self-control, a substantial proportion of SEF neurons encoded the saccade direction of the initial choice (*initial direction*; aligned to target onset: 89/310, 28.7%; aligned to choice: 92/310, 29.7%) and the delay-discounted subjective value of the initially chosen option (*initial chosen value*; target onset: 55/310, 17.7%; choice: 62/310, 20%). A moderate proportion encoded the identity of the chosen option (S or L; *initial choice*; target onset: 41/310, 13.2%; choice: 57/310, 18.4%). Only a small proportion encoded unchosen value (*initial unchosen value*; target onset: 25/310, 8.1%; choice: 36/310, 11.6%), choice difficulty (i.e., absolute value difference; target onset: 24/310, 7.7%; choice: 45/310, 14.5%), and relative L value (*SVL-SVS*; target onset: 17/310, 5.5%; choice: 17/310, 5.5%).

Multiple neurons encoded interactions of self-control with other variables in at least one time bin (target onset: 28/39, 71.8%; choice: 40/56, 71.4%) (**Table 2**). The most instances of interactions were found with *initial direction* (target onset**:** 20/39 neurons, 50 time bins; choice: 28/56 neurons, 84 time bins), indicating that these self-control signals were spatially tuned. SEF neurons encoded therefore both an abstract non-directional self-control signal (target onset: 29/39 neurons, 102 time bins; choice: 42/56 neurons, 132 time bins), and a directional self-control signal that could directly control motor activity. The second most frequent interaction was with *initial chosen value* (target onset**:** 13 neurons, 39 time bins; choice: 19 neurons, 63 time bins). Frequency of other interactions are displayed in **Table 2**.

Fluctuations in SEF activity predict monkeys' response to temptation.

The preceding analysis demonstrated that SEF neurons differentiate amongst different levels of self-control, which was defined *post hoc* based on the monkey's response to temptation. Since these signals were present before initial and final choice, they are in a position to potentially influence the monkey's response to temptation. We therefore determined how well the final choice of the monkey could be predicted by trial-to-trial variations in the neuronal activity of SEF neurons. Only temptation trials were used for this analysis.

We collected neuronal activity in overlapping 50-ms time bins shifted by 10ms in two time periods: (1) -200ms to 300ms around target onset, and (2) -200ms to 300ms around initial choice (**Figure 4A**). For each time bin in a given neuron, we fitted a GLM that predicted final choice (L or S) based on L delay, final chosen direction, and the number of spikes recorded in the time bin as potential predictors. We searched for neurons and time bins for which the best fitting model included the neuron's spike count in explaining the variance in final choice, as this would indicate that SEF activity predicts final choice over and above that of other choice-relevant variables.

We found that trial-to-trial fluctuations in the activity of a subset of SEF neurons (aligned to target onset: 76/310, 24.5%, **Figure 4B**; aligned to initial choice: 97/310, 31.3%, **Figure 4H**) significantly correlated with final choice for at least three consecutive time bins (70ms, accounting for bin overlap). A mean of 9.43 (SD=4.1) and 12.6 (SD=6.2) neurons robustly predicted final choice at each given 50-ms time window when aligned to target onset and choice, respectively (**Figure 4E, K**). To estimate the magnitude of the influence of neuronal activity on final choice, we determined the difference in the probability of choosing the L option in response to temptation (i.e., final L choice; LL, SL trials) for maximum and minimum spike rates in time bins with significant correlation of activity and choice. **Figure 4C, I** shows the simulated choice functions for final L choice when an SEF neuron's activity was at its maximum (thick green line) or

minimum (thick red line). The estimated indifference delay as a function of spike rate is shown to the right (**Figure 4D, J**). In both neurons, increasing firing rates in the relevant time bin (grey background) resulted in a final choice function that was strongly shifted towards longer L delays. This indicates that increases in spike rate for this neuron substantially increased the probability of finally choosing the L option. Neurons started predicting final choice early in the trial (target aligned: mean=-88ms, SD=13.3; choice aligned: mean=-89.9ms, SD=10.9s), and could be either positively or negatively correlated with finally choosing L in response to temptation (**Figure 4F, L**).

Overall, a total of 122/310 neurons (39.4%) were predicted final choice in response to temptation in at least one of the two time periods. 80 neurons were found to both encode self-control and predict the monkey's final choice in response to temptation. This is 39.6% (80/202) of neurons encoding self-control and 65.6% (80/122) of neurons predicting final choice.

Thus, fluctuations in the activity of SEF neurons are predictive of the willingness to delay gratification in the face of temptation. This provides additional strong evidence for an active role of SEF in exercising self-control in the context of delayed rewards.

SEF activity and graded self-control

Our task design only allows us to distinguish high or low self-control levels and we treated it accordingly as a categorical variable. However, self-control might be a continuous variable with graded levels. Graded differences in self-control would not lead to behavioral differences in LL trials, since weaker or stronger self-control above a threshold level all lead to rejection of the temptation. In contrast, in LS trials the strength of self-control should be proportional to the length of time the monkey is able to withstand temptation before giving in to it (i.e., switch reaction time; switch RT) (**Figure 5A**). Indeed, there was considerable variation in the length of time the monkey waited before ultimately switching to the S option (**Figure 2F**), which might reveal finer-graded differences in self-control levels.

We tested if there are any SEF neurons whose activity predicted graded levels of self-control, defined by the reaction time for switching from the L to S option. Only low self-control trials (i.e., temptation trials where the monkey initially chooses the L and switches to the S option) were used for analysis. We collected neuronal activity in overlapping 100-ms time bins shifted by 25ms aligned to two events: (1) -300ms to 150ms around initial choice, and (2) -1000ms to 0ms before switching (**Figure 5B**).

We found a substantial proportion of neurons (aligned to choice: 173 neurons, 55.8%; aligned to switch: 244 neurons, 78.7%) robustly predicting switch time (i.e., predicting switch time for 3 consecutive time bins, or 150ms, considering bin overlap). A mean of 79.9 (SD=12.1) and 101.1 (SD=11) neurons predicted switch RT at each given 100-ms time window when aligned to choice (**Figure 5G)** and switch **(Figure 5L**). Example neuronal activity that differentiates among different switch RTs is shown aligned to initial choice (**Figure 5C-F)** and switch **(Figure 5H-K)**. Thus, SEF activity not only differentiated between high and low self-control, but also reflected much finer-grained differences in self-control. A high proportion of neurons that predicted switch RT was directionally tuned (choice: 81/173, 46.8%; switch: 158/244, 64.8%).

Overall, a total of 248/310 neurons (80%) were identified to robustly predict the timing of switch in at least one of the two alignments. 178 neurons encoded both self-control and predicted graded self-control. This is 88.1% (178/202) of neurons encoding self-control and 71.8 (178/248) of neurons predicting graded self-

control. 111 neurons predicted both final choice and graded self-control levels. This is 91% (111/122) of neurons predicting final choice and 44.8% (111/248) of neurons predicting graded self-control,

Discussion

By designing a novel task that explicitly reveals self-control levels, we identified SEF activity indicating whether monkeys will resist (high self-control; LL trials) or give in to temptation (low self-control; LS trials). SEF neuronal activity was also directly predictive of the final choice in the presence of temptation, as well as reflective of more fine-grained, continuous differences in self-control.

In this paper, we concentrate on self-control signals that emerge before the onset of temptation. This form of self-control represents proactive control (6,7). Proactive control refers to the maintenance of goal-relevant information before the occurrence of a cognitively demanding event, thereby anticipating and preventing interference to enhance performance. Such maintenance of goal-relevant information has been proposed to give rise to cognitive control by providing bias signals (8).

A recent line of work has argued that self-control is just a form of value-based decision making and does not require specific control circuits independent of decision making (4,9–11). In that case, the activity differences on LL and LS trials would be caused by random fluctuations in the neural encoding of action value that are correlated with choice. However, we would expect similar variability in action values for initial-S choices, and likewise the same rate of SL switches. Instead, we observed dissimilar proportions of LS and SL switches. Therefore, the difference in neural activity between LL and LS trials do not seem to be primarily driven by random fluctuations in action value.

More generally, our data clearly reveals that self-control is a process that is distinct from, and thus not simply a form of, value-based choice. The GLM analysis clearly showed independent neural signals reflecting both self-control and value-based choice (i.e., chosen value, unchosen value, SVL-SVS). However, we do not propose that self-control and value-based choice processes are completely separated. We found a high proportion of neurons with interactions between self-control and chosen value. It seems likely that the self-control circuitry can provide modulatory input to the value-based choice process. Thus, our results demonstrate that self-control is a process that is separable from value-based choice, although self-control can influence the choice process by modulating the valuation process.

Are there other possible interpretations of the SEF activity than self-control signals? The fact that visual stimuli, choices, and motor activity is the same across LL and LS trials argues strongly against visual or motor differences causing the self-control signals. Likewise, the presence of spatially specific self-control signals argues against these signals simply reflecting general arousal, or motivation level. Lastly, it is unlikely that LS switches and the related neural activity are caused by a failure to suppress attention to interference from the unexpected visual cue. Such failures of inhibitory control of attention should result equally often in LS and SL switches.

Previous human neuroimaging studies and neurophysiology studies in monkeys identified a distributed network of brain regions involved in self-control (12–16). The role of SEF within this network is not well understood, but SEF is interconnected with this network (17). Its connections allow it to connect value-related and cognitive areas with areas more directly involved with initiating eye movements. While the SEF has been found not to provide direct input in initiating a saccade, it is hypothesized to play a modulatory, biasing role to adjust saccade production according to one's goals (18). In a risky decision-making task, SEF inactivation was found to reduce the ability to suppress automatic, maladaptive responses (i.e., aborting the trial before reward delivery) to losing a gamble (19). Therefore, SEF does seem anatomically and

functionally situated to encode and utilize self-control signals to causally influence behavior in response to temptation.

These signals could bias towards the prudent option to influence behavior for when temptation appears. Previous works have reported biasing effects of pre-trial neural activity on choices later in the trial (20–22). In the SEF, neuronal activity early in the trial was found to vary with sequential adjustments in response latency in the stop signal task, and it has been proposed that SEF proactively regulates movement initiation by adjusting the level of excitation and inhibition of the oculomotor systems based on anticipated task requirements (18). Similarly, in our task, the self-control signals and their interaction with *initial direction* (i.e., the spatial location of the initially chosen option) could bias towards or against the withholding of potential switches. Indeed, SEF neurons were not only encoding self-control but other value-related variables necessary for initial choice. Thus, it seems that SEF is involved both in explicit delay-discounted valuation judgments and also self-control, which interact to influence initial and final choice.

Methods

Four adult male rhesus macaque monkeys were trained to perform the self-control task (**Figure 1**). For two monkeys, each animal was surgically implanted with titanium head posts and a hexagonal titanium recording chamber. All animal care and experimental procedures were approved by Johns Hopkins University Institutional Animal Care and Use Committee.

Experimental Design

Self-control task. The search for self-control signals required a behavioral indicator that estimates the level of self-control on individual trials, independent from other factors that influence choice. Our self-control task provides such a behavioral indicator.

The self-control task (**Figure 1**) included two trial types: (1) no temptation (47-90% of trials, mean = 63.2%) and 2) temptation (10-53% of trials, mean = 36.8%). We presented both trial types interleaved in blocks consisting of sets including both trial types and all delay contingencies. Within a block, the order of appearance was randomized and a particular trial was never repeated, so that the monkeys could not decide before the targets were shown. Each target location was randomized and counterbalanced across trials to prevent the monkeys from preparing a movement toward a certain direction before target appearance.

Delay length of the S option (S delay) was fixed within each behavioral session (500 or 700ms), and delay length of L options (L delay) varied from trial to trial (500-10,000ms; equal to or longer than S delay). Monkeys indicated their initial choice by making a saccade to one of the options and were required to remain fixated on the chosen option until the end of the delay in order to be rewarded. If the S option was chosen, the intertrial interval was adjusted to match what the total duration of the trial would have been had the monkey chosen the L option. Therefore, to maximize total reward in a given session the monkeys should always choose the L option. However, choice behavior indicates that the probability of choosing the L option is dependent on the length of the L delay (**Figure 2A**), suggesting that the subjective valuation of the L option is discounted by the waiting time.

<u>No temptation trials</u>: No temptation trials were structured like the classic intertemporal choice task. A trial started with the monkey fixating a central fixation point. After a variable fixation period (700-1200ms), two annulus-shaped stimuli were presented which served as saccade targets. Simultaneously, the fixation point disappeared, which was a signal for the monkey to choose one of the targets. Each annulus had a colored section of variable length on it. The color and length of the section indicated the amount of reward and the delivery delay associated with the target, respectively. The reward amount and the delay of the targets were adjusted so that one was always a small, short-delayed target (S option) and the other, a large, long-delayed one (L option). Once the monkey chose a target by making a saccade towards it and fixating it, the colored annulus section of the chosen target began to decrease with an angular velocity of 60˚/sec. At the same time, the unchosen target disappeared. On each trial, the initial angle of the band and the direction in which it decreases were randomized independently for each stimulus. If the monkey fixated the chosen target until its colored section was extinguished, fluid reward in the amount indicated by the color was given. If the monkey chose the S option, the intertrial interval was adjusted corresponding to the non-chosen L option, so that the total length of a trial was the same irrespective of the choice of the monkey. This adjustment ensured that there was no opportunity cost of choosing the L option. The reward rate over all trials was always maximized when the monkey exclusively chose the L reward.

The monkey was always required to maintain fixation of the fixation spot until it disappears and then of the target until reward delivery ('complete trials'). If the monkey broke fixation in either one of the two time periods, the trial was aborted and he received no reward ('incomplete error trials'). After the usual

inter-trial-interval, a new trial started. In some experimental sessions, an aborted trial was followed by a 'repeat-trial', which used the same reward options as in the aborted trial. If the monkey had made a choice before breaking fixation, only the chosen target was presented, otherwise both targets were presented. The location of the targets however was randomized, so that the monkey could not prepare a saccade in advance. Such 'repeat-trials' were not included in analysis.

Temptation trials: The initial sequence of events in temptation trials was identical to no temptation trials. However, after the monkey made a saccade to one of the targets, both targets continued to stay on the screen and the bands of both targets started to decrease simultaneously. This continued presence of the alternative target, while the monkey waited for delivery of the initially chosen reward, indicated that the monkey was allowed to switch targets once. At this stage, the monkey could either maintain fixation of the initially chosen target ("stay trial"), or he could change his choice by making a saccade to the other target ("switch trial"). If the monkey switched, the initially chosen target disappeared, indicating that the monkey now had to fixate his final target until reward delivery. In stay trials, the monkey continued to fixate the initial target until reward delivery. When the monkey had chosen the L target, the band of the S target might extinguish, while he still waited for the more delayed L reward. In that case, the S target changed to a filled circle indicating immediate reward availability. The intertrial interval was adjusted in the same way as in no temptation trials depending on the final choice, ensuring that switching targets did not change overall reward rate over all trials. Depending on which options were selected on the initial and final choice there were four different possible types of temptation trials: LL, LS, SS, and SL trials. On LL trials, the initial and final choices were the L option. On LS trials, the initial choice was the L option but the final choice was the S option. On SS trials, the initial and final choices were the S option. On SL trials, the initial choice was the S option but the final choice was the L option.

Data Acquisition

Neurophysiological Recording Procedures. Single neuron activities were recorded extracellularly with single tungsten microelectrodes. Electrodes were inserted through a guide tube positioned just above the surface of the dura mater and were lowered into the cortex under the control of a self-built Microdrive system. The electrodes penetrated the cortex perpendicular to the surface of the cortex. Electrophysiological data were collected using the TDT system (Tucker-Davis technologies, Alachua, FL) at a sampling rate of 1000 Hz. Action potentials were amplified, filtered, and discriminated conventionally with a time-amplitude window discriminator. Spikes were isolated online if the amplitude of the action potential was sufficiently above a background threshold to reliably trigger a time-amplitude window discriminator and the waveform of the action potential was invariant and sustained throughout the experimental recording. Spikes were then identified using principal component analysis (PCA) and the time stamps were collected at a sampling rate of 1000 Hz.

Behavioral Recording Procedures. MonkeyLogic software (https://monkeylogic.nimh.nih.gov/) was used to control task events, stimuli, and reward, as well as to monitor and store behavioral events. During the experimental sessions, the monkey was seated in an electrically insulated enclosure with its head restrained, facing a video monitor. Eye positions were monitored with an infrared corneal reflection system, a 500 Hz video eye tracker (EyeLink 1000; SR Research, Ottawa, Canada) at a sampling rate of 1000 Hz. Reinforcement was delivered using a stepper motor driven syringe pump fluid delivery system that is highly accurate across the entire range of fluid amounts used in the experiment. All analyses were performed using self-written Matlab code, unless noted otherwise.

Behavioral Analysis

Choice functions. The probability that the monkey chooses the S option was plotted as a function of the delay associated with the larger later reward (L delay).

When the L delay was short, monkeys are more likely to choose the L option. As the delay of the L option increases, monkeys increasingly choose the S option. We used a generalized linear model to estimate the probability of choosing the S option, p(S), as a continuous function of the value of the L delay (i.e., the choice function):

$$\log\left[\frac{p(S)}{1-p(S)}\right] = b0 + b1 * L\ delay$$

For no temptation trials, the choice function was estimated for only the initial choice because once an option was selected and the monkey fixated on its choice for the required amount of time, the trial ended. For temptation trials, the choice function was estimated for both the initial and the final choice. The L delay at which choosing either the L or S option is equal [p(S) = 0.5] is known as the indifference point (23). By definition, at this point, the subjective value of the two options must be equal. We called this the indifference delay (i.e., the L delay at which the monkey is indifferent between the S and L options), D_{indiff}, with which we computed the temporal discounting parameter and the subjective value of each option.

Estimation of Subjective Value (SV). SV of options was calculated by delay-discounting the final reward amount, using the following hyberbolic discounting formula (2):

$$SV_{option} = \frac{AV_{option}}{1 + k * D_{option}}$$

SV_{option} is the calculated SV of the option, AV_{option} is the actual reward amount the monkey is to receive, k is the delay discounting parameter which reflects how much the animal discounts later rewards, and D_{option} is the length of delay of that option.

The discounting parameter, k, was calculated by solving for the following:

$$\frac{AV_S}{1 + k * D_{indiff}} = \frac{AV_L}{1 + k * D_{indiff}}$$

AV_S and AV_L indicate actual reward amount of the S and L options, respectively, and D_{indiff} is the indifference Delay length of the L option, at which the probability of choosing the S is equal to 0.5 (i.e. the delay length at which S and L becomes indifferent to the monkey). Thus, the equation solves for k so that $SV_S = SV_L$ at the given indifference delay.

Neurophysiological Analyses

Only completed (i.e., trials where the monkey did not break fixation from the chosen option until the end of the delay and thus received reinforcement) no-temptation and temptation trials were used in analyses.

LASSO regularization and Poisson-GLM Analysis. To quantitatively characterize how the individual SEF neuron is modulated by task-related variables in different time periods throughout the task, we performed Poisson-GLM analyses. Because the Poisson distribution is used to model the number of times an event occurs in a given interval of time, it can also be used to characterize the firing activity of a neuron. Under the assumption that neuronal firing (i.e., the number of spikes) follows a Poisson distribution, the activity of a neuron in a given time period can be represented as the exponential of the linear combination of predictor variables. LASSO regularization, when performed before Poisson-GLM, reduces the number of predictors to test. The LASSO regularization stage described below was only performed for target onset and before choice periods of encoding models. Only the Poisson-GLM analysis, without LASSO

regularization, was performed for encoding models aligned to fixation, decoding models, and graded self-control models.

LASSO regularization. In the first stage of model fitting, we performed a cross-validated LASSO regularization of a Poisson-GLM to select a subgroup of the 7 predictors from which we would fit every nested model and find the best fitting model. LASSO can be used in generalized linear modeling to reduce the number of predictors by identifying and selecting important predictors over redundant predictors. The LASSO optimization method shrinks regression coefficients towards zero, such that the relatively unimportant predictors are set to zero earlier than the more important predictors. We used the MATLAB lassoglm function, which solves the following for a given nonnegative lambda:

$$\min_{\beta_0,\beta}\left(\frac{1}{N}Deviance(\beta_0,\beta) + \lambda\sum_{j=1}^{p}|\beta_0|\right)$$

Here, lambda is the regularization parameter, N is the number of observations, and B0 and B are the intercept and the vector of predictor coefficients of length p, respectively. Deviance is a function which calculates the deviance of the fitted model applied to new data. We selected predictors with non-zero coefficients at the lambda value with minimum expected lambda-penalized deviance, which is equivalent to maximizing the lambda-penalized loglikelihood, as calculated by cross-validation. This first stage was essential as even with 7 predictors, there are over a billion different nested models that must be tested, which was unrealistic given the current computational and time constraints. We included all predictors and select 2-way interactions (see **Table 3**) for LASSO optimization.

Poisson-GLM. After LASSO optimization, we moved onto the second stage of model fitting, in which we performed a series of Poisson-GLM to select the best fitting model for each neuron in a given time window. Given a set of predictors and interactions selected by LASSO, any predictors included in the interaction were included in the full model for Poisson-GLM (e.g., if *Trial identity* and *initial CV*initial direction* were selected by LASSO, *Trial identity*, *initial CV* and *initial direction* were all included in the full model in the second stage, such that they were also considered as potential predictors). Then, the full model and every nested model of the full model (i.e., every combination of predictors and with and without their two-way interactions), including the null model, were tested. The Akaike Information Criterion (AIC) value was extracted from each model fit. The model with the lowest AIC was selected as the neuron's best fit model in the given time period, and the predictors included in the best fitting model were considered to be represented by the neuron in that time period.

Encoding models

Encoding models were designed to predict SEF neuronal activity (*Spike count*) from a series of task-related variables in three different time periods: (1) central fixation: -100~700ms around fixation onset (i.e., the time the monkey fixated on the central fixation point), (2) target onset: -50~300ms around target onset (i.e., when choice targets appeared), (3) before choice: -300~50ms around initial choice. Within each time period, the full model including all variables and each of its nested models were fit in 100-ms moving windows, shifted by 25ms. For target onset and before choice, we allowed for select 2-way interactions (**Table 3**). The model with the lowest AIC was selected as the best fit model. The variables included in this best fit model were regarded as the set of variables the neuron encodes.

Central fixation. During fixation, only the central fixation point is on the screen and no information about the upcoming choice is available. Therefore, the only variable we tested for was *Self-control*. The full model thus included a constant term and *Self-control*:

$$\log[\textit{Spike count}] = b0 + b1 * \textit{Self-control}$$

The full model and the null model were fit to the data. The Akaike Information Criterion (AIC) value was extracted from each model fit. The model with the lowest AIC was selected as the neuron's best fit model in the given time period, and if *Spike count* was included in the best fitting model, the neuron was considered to encode self-control at fixation.

The response variable, *Spike count*, was defined as the number of action potentials fired by the neuron within the given time period. For visualization, spike count was later converted to firing rate (*FR*; spikes/sec).

Target onset, before choice. After target onset and before initial choice, both choice targets are available to the monkey (but no information on whether the trial will be a temptation trial). Therefore, the full model for time periods 'target aligned' and 'before choice' included a constant term, 6 choice-related variables, the variable *Self-control*, and select 2-way interactions (**Table 3**):

$$\log[\textit{Spike count}] = b0 + b1 * \textit{initial direction} + b2 * \textit{initial choice} + b3 * \textit{initial chosen value} + b4 * \textit{initial unchosen value} + b5 * \textit{SVL-SVS} + b6 * \textit{Choice difficulty} + b7 * \textit{Self-control}$$

Two non-value categorical variables and four subjective value-related variables were considered as potential predictors. *Initial direction* was a categorical variable that indicated the location of the chosen target to which the saccade was directed. There were four possible target locations, one on each quadrant. *Initial choice* was another categorical variable which indicated the identity of the chosen option, i.e., S or L options. There were four value variables: *Initial chosen value (Initial CV)* was defined as the subjective delay-discounted value of the chosen option in the given trial, *Initial unchosen value* (*Initial UV*) was the subjective value of the unchosen option in the given trial, *SVL-SVS* is the subjective value difference in the L and S options. Because the delay and amount and therefore the subjective value of the S option was kept constant throughout an experimental session, *SVL-SVS* is directly proportional to the value of the L option, and is directly anticorrelated with L delay length (see *Estimation of Subjective Value* above for details). Finally, *Choice difficulty* was calculated by taking the absolute value of SVL-SVS, to capture neuronal activity that reflects the absolute difference in subjective values of the available options, irrespective of which was preferred. Smaller *Choice difficulty* values would indicate a more difficult trial, as smaller value differences between the available options would make the choice harder.

Self-control was a categorical variable which distinguished amongst trials with high self-control, low self-control, and other trials. High self-control trials were defined as temptation trials in which the monkey initially chooses the L option and stays with it (LL trials) and low self-control trials were defined as temptation trials in which the monkey initially chooses the L option but switches to the S option (LS trials). No temptation trials, temptation trials in which the monkey chose and stayed with the S option, and temptation trials in which the monkey chose the S option but switched to the L option were considered as other trials, as it was either impossible to distinguish amongst high and low self-control levels, or self-control was not required to reach the end of the trial and receive the reward.

Decoding models.

Decoding models were designed to predict choice from SEF neuronal activity in two time periods: (1) -200 ~ 300ms around target onset, and (2) -200 ~ 300ms around initial choice. Within each time period, the full model including all variables and each of its nested models were fit in moving windows of 50-ms time bins, shifted by 10ms. For both time periods, the predictors other than *spike count* were *L delay* and *initial direction*, where *L delay* refers to the L option's delay length and *initial direction* refers to the spatial location of the initially chosen option. This was done both for predicting initial choice and final choice.

Since *L delay* is closely correlated with the subjective value of the L option and is known behaviorally to influence choice, the activity of neurons that represents the decision process would not be expected to explain variance in final choice beyond the influence of *L delay*. In contrast, a neuron whose activity explains a significant proportion of the variance in final choice beyond the influence of *L delay* should result in time bins where the model includes the predictor *spike count*. For such a neuron, increasing firing rate should be correlated with either a rightward shift or leftward shift in the choice functions. The inclusion of saccade direction as a predictor allowed us to determine if the influence of neuronal activity was spatially tuned.

For predicting initial choice, *spike count*, *L delay*, and *initial direction* from both no-temptation and temptation trials (L, S, LL, LS, SL, SS trials) were used as predictors of the model. For predicting final choice, *spike count*, *L delay*, and *initial direction* from only the temptation trials (LL, LS, SL, SS trials) were used as predictors of the model. The full model includes all of the predictors, a constant term, and all two-way and three-way interactions of the predictors:

$$\log\left[\frac{p(S)}{1-p(S)}\right] = b0 + b1 * L\ delay + b2 * initial\ direction + b3 * spike\ count$$

For each time bin in a given neuron, the full model and its nested models were fit. The model with the lowest AIC was taken as the neuron's best fitting model. A given neuron was interpreted to influence choice over and above the behavioral variables only if the predictor *spike count* was included in the best fitting model.

Graded self-control model.

The graded self-control model was designed to predict the latency for switching from *spike count* of the SEF neuron, controlling for other task-related variables that can influence switch latency in the given trial. We examined -300 ~ 150ms around initial choice in moving windows of 100-ms time bins shifted by 25ms. Switch latency, which is the time between initial choice and time of switch, was the response variable. Predictors other than *spike count* (number of spikes by a given neuron in the given time bin) included *L delay* (L option's delay length), *CV* (subjective value of the chosen option), *UV* (subjective value of the unchosen option), and *initial direction* (location of the chosen option). This aimed to capture neurons whose activity around initial choice significantly predicted the time the monkey waits after realizing it is a temptation trial (i.e., the unchosen option does not disappear and remains available) before ultimately switching. Only temptation trials where the monkey initially chose the L option (i.e., LS trials) were included for analysis.

The full model includes the above predictor variables, a constant term, and all two-way and three-way interactions, producing the following formula:

Switch latency = b0 + b1**L delay* + b2**CV* + b3**UV* + b4**initial direction* + b4**spike count*

The *L delay*, *CV*, and *UV* can all influence the decision of whether and when to switch. Therefore, a neuron whose activity around initial choice significantly influences variance in switch latency beyond the influence of *L delay*, *CV*, and *UV*, should have a best fitting model (defined as the model with the lowest AIC) that includes *spike count* as a predictor.

References

1. Frederick S, Loewenstein G, O'Donoghue T. Time Discounting and Time Preference: A Critical Review. J Econ Lit. 2002 Jun;40(2):351–401.

2. Green L, Myerson J. A Discounting Framework for Choice With Delayed and Probabilistic Rewards. Psychol Bull. 2004 Sep;130(5):769–92.

3. Crockett MJ, Braams BR, Clark L, Tobler PN, Robbins TW, Kalenscher T. Restricting temptations: neural mechanisms of precommitment. Neuron. 2013 Jul 24;79(2):391–401.

4. Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. Nat Neurosci. 2007 Dec;10(12):1625–33.

5. Stuphorn V. The role of supplementary eye field in goal-directed behavior. J Physiol Paris. 2015;109(1–3):118–28.

6. Braver TS. The variable nature of cognitive control: a dual mechanisms framework. Trends Cogn Sci. 2012 Feb 1;16(2):106–13.

7. Stuphorn V, Emeric E. Proactive and reactive control by the medial frontal cortex. Front Neuroengineering [Internet]. 2012 [cited 2023 Jun 19];5. Available from: https://www.frontiersin.org/articles/10.3389/fneng.2012.00009

8. Miller EK, Cohen JD. An Integrative Theory of Prefrontal Cortex Function. Annu Rev Neurosci. 2001;24(1):167–202.

9. McGuire JT, Kable JW. Medial prefrontal cortical activity reflects dynamic re-evaluation during voluntary persistence. Nat Neurosci. 2015 May;18(5):760–6.

10. Berkman ET, Hutcherson CA, Livingston JL, Kahn LE, Inzlicht M. Self-Control as Value-Based Choice. Curr Dir Psychol Sci. 2017 Oct 1;26(5):422–8.

11. Bulley A, Lempert KM, Conwell C, Irish M, Schacter DL. Intertemporal choice reflects value comparison rather than self-control: insights from confidence judgements. Philos Trans R Soc Lond B Biol Sci. 2022 Dec 19;377(1866):20210338.

12. McClure SM, Laibson DI, Loewenstein G, Cohen JD. Separate neural systems value immediate and delayed monetary rewards. Science. 2004 Oct 15;306(5695):503–7.

13. Hare TA, Camerer CF, Rangel A. Self-control in decision-making involves modulation of the vmPFC valuation system. Science. 2009 May 1;324(5927):646–8.

14. Kim S, Hwang J, Lee D. Prefrontal Coding of Temporally Discounted Values during Intertemporal Choice. Neuron. 2008 Jul 10;59(1):161–72.

15. Figner B, Knoch D, Johnson EJ, Krosch AR, Lisanby SH, Fehr E, et al. Lateral prefrontal cortex and self-control in intertemporal choice. Nat Neurosci. 2010 May;13(5):538–9.

16. Turner BM, Rodriguez CA, Liu Q, Molloy MF, Hoogendijk M, McClure SM. On the Neural and Mechanistic Bases of Self-Control. Cereb Cortex. 2019 Feb 1;29(2):732–50.

17. Huerta MF, Kaas JH. Supplementary eye field as defined by intracortical microstimulation: connections in macaques. J Comp Neurol. 1990 Mar 8;293(2):299–330.

18. Stuphorn V, Brown JW, Schall JD. Role of supplementary eye field in saccade initiation: executive, not direct, control. J Neurophysiol. 2010 Feb;103(2):801–16.

19. Chen X, Stuphorn V. Inactivation of Medial Frontal Cortex Changes Risk Preference. Curr Biol CB. 2018 Nov 19;28(22):3709.

20. Mochizuki K, Funahashi S. Prefrontal spatial working memory network predicts animal's decision making in a free choice saccade task. J Neurophysiol. 2016 Jan;115(1):127–42.

21. Padoa-Schioppa C. Neuronal origins of choice variability in economic decisions. Neuron. 2013 Dec 4;80(5):10.1016/j.neuron.2013.09.013.

22. Shadlen MN, Newsome WT. Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. J Neurophysiol. 2001 Oct;86(4):1916–36.

23. Luce, R. Duncan. Utility of gains and losses : measurement- theoretical, and experimental approaches. Mahwah, NJ ; London: L. Erlbaum; 2000. 331 p.

Tables and figures

Table 1.

Neurons positively or negatively correlated with self-control in different time periods.

		Fixation	Target	Choice
Positively correlated	Mostly positive	22	10	9
	Purely positive	61	13	18
	Total	83	23	27
Negatively correlated	Mostly negative	29	9	10
	Purely negative	72	7	19
	Total	101	16	29
Total		**184**	**39**	**56**

A neuron is counted as mostly positive if it was positively correlated in the majority of time bins that included self-control in the best fitting model. A neuron is counted as purely positive if it was positively correlated in all time bins that included self-control in the best fitting model.

Table 2.

Instances of interactions with other variables.

Interactions	Number of neurons		Number of time bins	
	Target	Choice	Target	Choice
Initial direction * self-control	20	28	50	84
Initial CV * self-control	13	19	39	63
Choice difficulty * self-control	9	19	21	47
SVL-SVS * self-control	9	7	16	10

Table 3.

Types of interactions considered for LASSO-GLM

Interactions
Initial choice * self-control
Initial direction * self-control
Initial CV * self-control
Initial UV * self-control
SVL-SVS * self-control
Choice difficulty * self-control
Initial direction * initial CV

Interactions included for encoding analysis, for time periods aligned to target onset and initial choice.

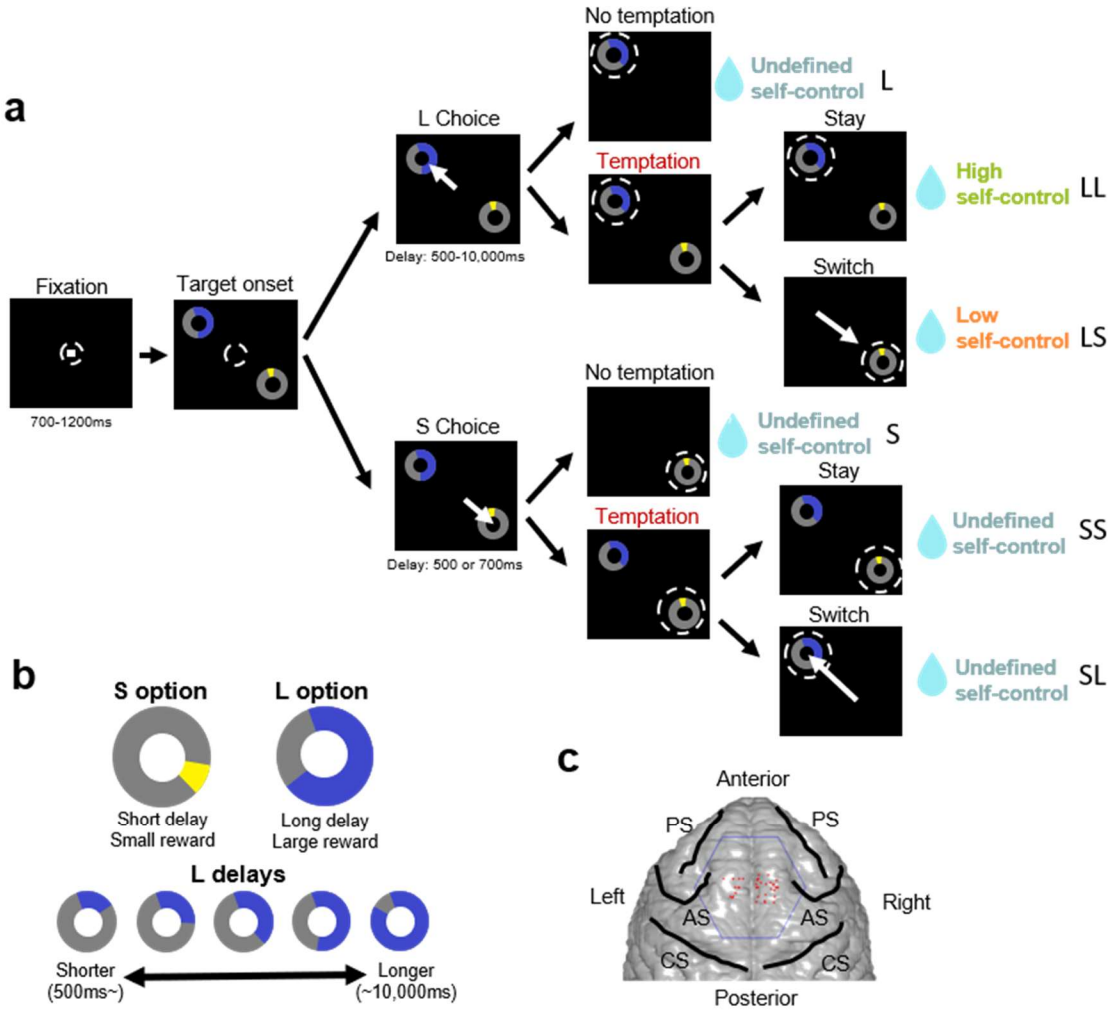

Figure 1 Recording locations and self-control task design. **a** Task design. After target presentation, the monkey chooses between S and L options by making a saccade to an option and remaining fixated until the end of the delay period. Dotted circles indicate the location of fixation by the monkey. Arrows indicate eye movements of the monkey. On temptation trials, the unchosen option remains available to switch from the initial choice. The finial decision to switch from or stay with the initially chosen option is referred to as final choice. **b** Choice targets. The delay to reward is indicated by the area of space colored within the circular cue, and reward amount is indicated by the color. S delay was fixed within each session. L delay varied from trial to trial, ranging between 500ms and 10,000ms. **c** Example recording sites for a monkey. PS: principal sulcus, AS: arcuate sulcus, CS: central sulcus.

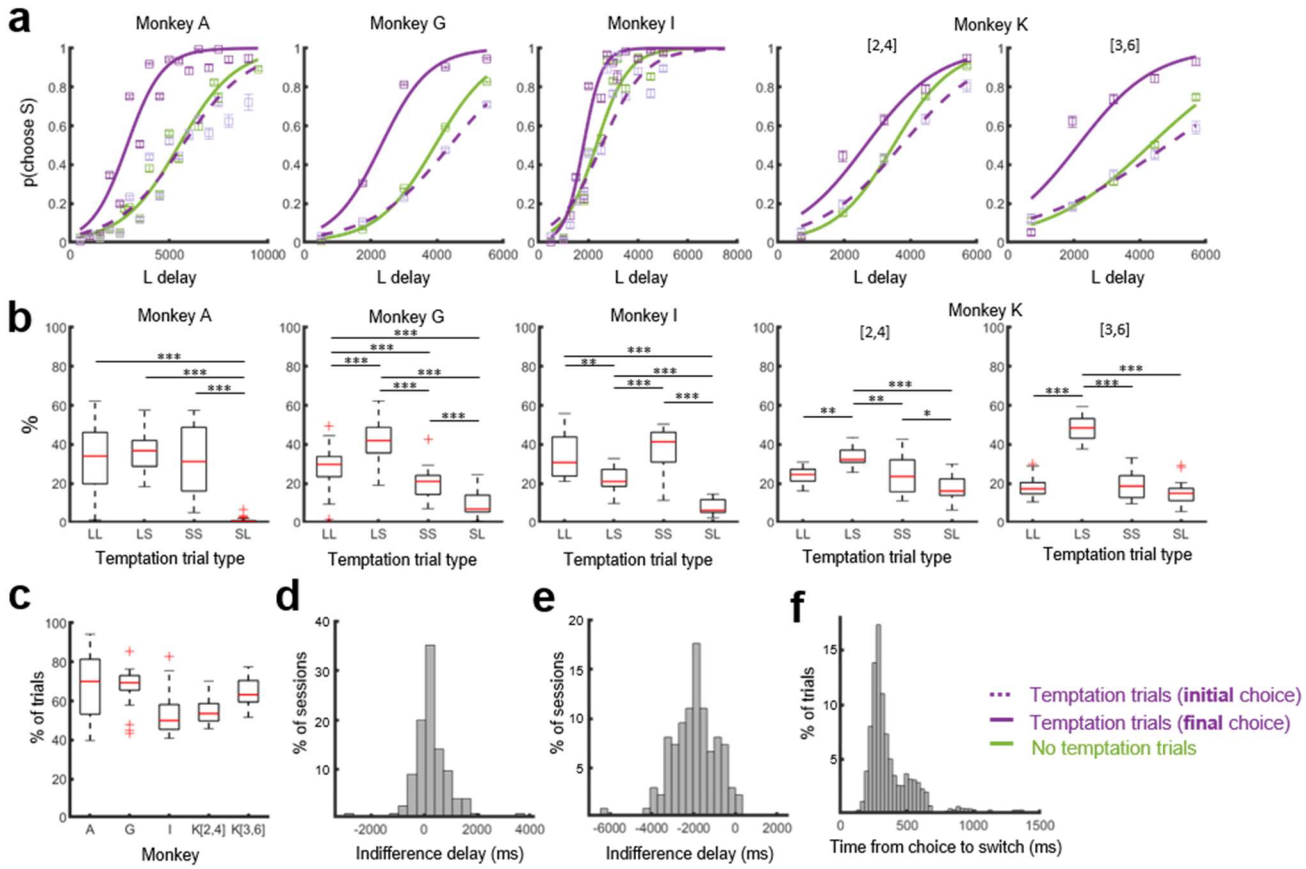

Figure 2 Behavioral results. **a** Choice functions for different trial types. The proportion of trials where the S option was chosen was plotted as a function of L delay. Solid green: no temptation trials, dotted purple: temptation trials initial choice, solid purple: temptation trials final choice. Different S and L reward combinations were interleaved within each session for monkey K, and therefore are visualized in separate plots. Numbers in square brackets indicate S reward and L reward amounts. **b** Proportion of temptation trial types. Difference in trial type proportions was statistically significant in all four monkeys (one-way ANOVA, monkey A: F=82.25, p=8.5358e-34; monkey G: F=103.75, p=3.2668e-36; monkey I: F=32.73, p=6.0707e-13; monkey K[2,4]: F=14.35, p=2.9678e-07; monkey[3,6]: F=93.27, p=2.5272e-23). Stars represent the result of multiple comparison of group means (*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$). **c** Proportion of trials where the L option was initially chosen. **d** Distribution of the difference in indifference delay for initial choice in temptation and no temptation trials. The mean was significantly different from zero (paired t-test, t = 4.9796, p = 1.9118e-06). **e** Distribution of the difference in indifference delay for final choice and initial choice in temptation trials. The mean was significantly different from zero (paired t-test, t=-11.8974, p=9.0434e-23). Two data points, one less than -7000 and one above 2000ms, were excluded from visualization. **f** Distribution of switch reaction time (switch RT). Switch RT indicates the time from choice until switch, in no temptation switch trials (LS, SL). 176/17152 data points which were above 1500ms are not shown.

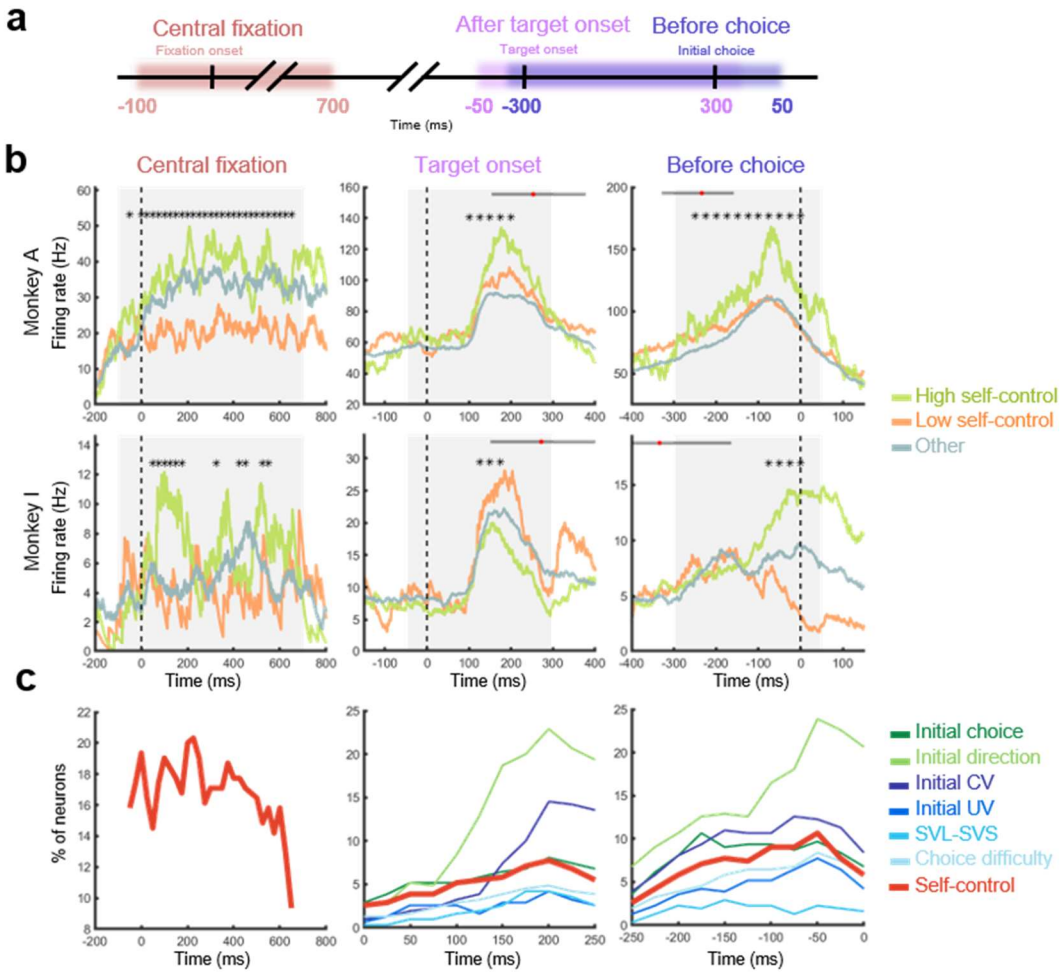

Figure 3 Encoding of self-control in SEF neurons. **a** Analyzed time periods and their alignments. Within each time period, SEF activity was analyzed in 100-ms time bins shifted by 25ms. Neuronal activity for central fixation, target onset and before choice was aligned to the time of fixation, target onset, and initial choice, respectively. **b** Spike density functions of example single neurons that encode different levels of self-control. Each panel is a different neuron, from monkey A (top row) and monkey I (bottom row). Stars indicate time bins where self-control was included as a predictor in the neuron's best fitting model. Black bars for target onset and before choice indicate the range of time of choice and target onset, respectively, in trials where the given neuron was active. Red dots indicate the median. **c** Proportion of neurons encoding each variable in a given time bin.

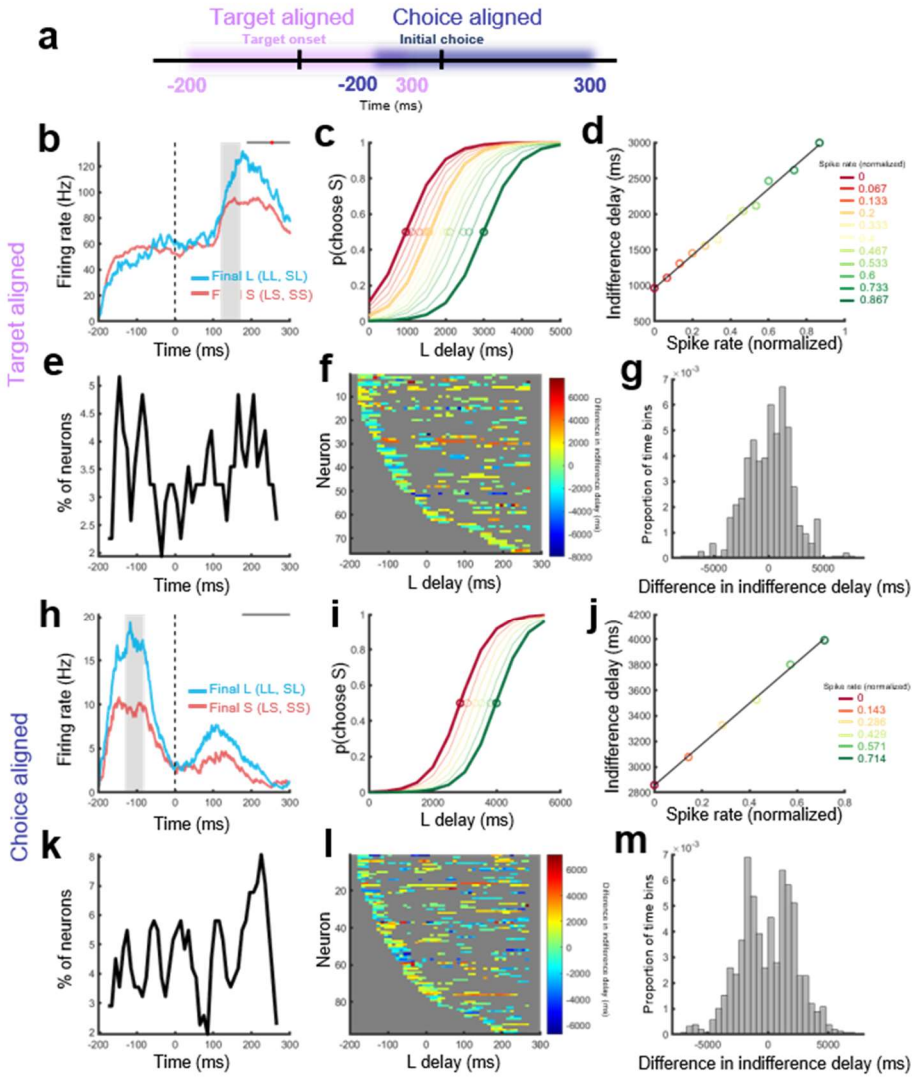

Figure 4 SEF neurons predict final choice in response to temptation. **a** Analyzed time periods and their alignments. Within each time period, SEF activity was analyzed in 50-ms time bins shifted by 10ms. Neuronal activity was aligned to target onset and initial choice. **b** Spike density function of an example SEF neuron, whose activity aligned to target onset predicted final choice. Black bar indicates the range of timing of choice, and red dot indicates the median. Gray box indicates the time bin from which panels **c** and **d** were created. **c** Simulated choice functions from using the best fitting model in the grayed out time bin from **b**. Each choice function is the expected choice function for a given activity level observed from the neuron. Red and green choice functions indicate the expected choice function for when the neuron is firing at its minimum and maximum firing rate, respectively. The thick orange choice function indicates that the medium normalized firing rate for this neuron was 0.2. **e** Proportion of SEF neurons whose activity predicted final choice when aligned to target onset, in each time bin. **f** The difference in the estimated indifference delay for maximum and minimum firing rates of each given neuron, plotted over time for all neurons that robustly predicted final choice, sorted by the onset of predictive activity. **g** Distribution of the difference in estimated indifference delay for minimum and maximum firing rates of neurons that robustly predict final choice. **h-m** Same as in **b-g**, but for the choice period.

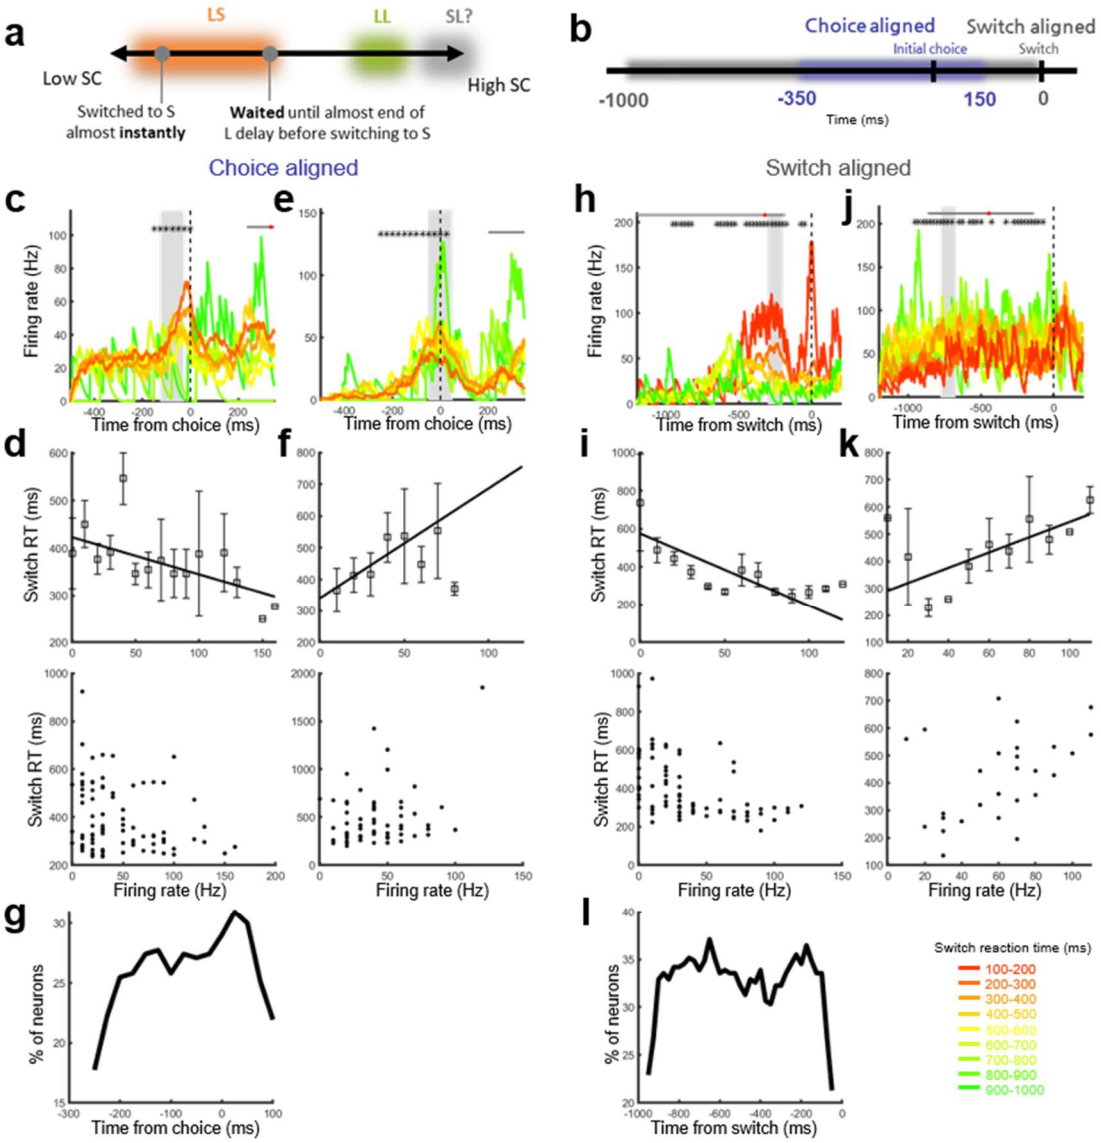

Figure 5 SEF neurons predict finer-grained differences in self-control. **a** Schematic illustration of a continuous gradient of self-control levels, and their potential behavioral consequences in the self-control task. **b** Analyzed time periods and their alignments. **c** Spike density functions by switch reaction time for an example SEF neuron whose activity robustly predicts switch RT and is negatively correlated with switch RT. Gray box indicates the time period for which panel **d** was plotted. Stars indicate time bins in which the neuron's activity predicted switch RT. Black bar indicates timing of switch, red dot indicates median time of switch in trials where this neuron was recorded. **d** Box plot (top) and scatter plot (bottom) of the relationship between switch RT and firing rate of the neuron in **c** in the grayed out period. **e-f** Same as in **c-d**, but for a neuron with a positive correlation with switch RT. **g** Proportion of neurons whose activity robustly predicts switch RT, aligned to choice. **h-l** Same as in **c-g**, but aligned to the time of switch. Black bars in **h**, **j** indicate timing of initial choice, and red dots indicate median time of choice.

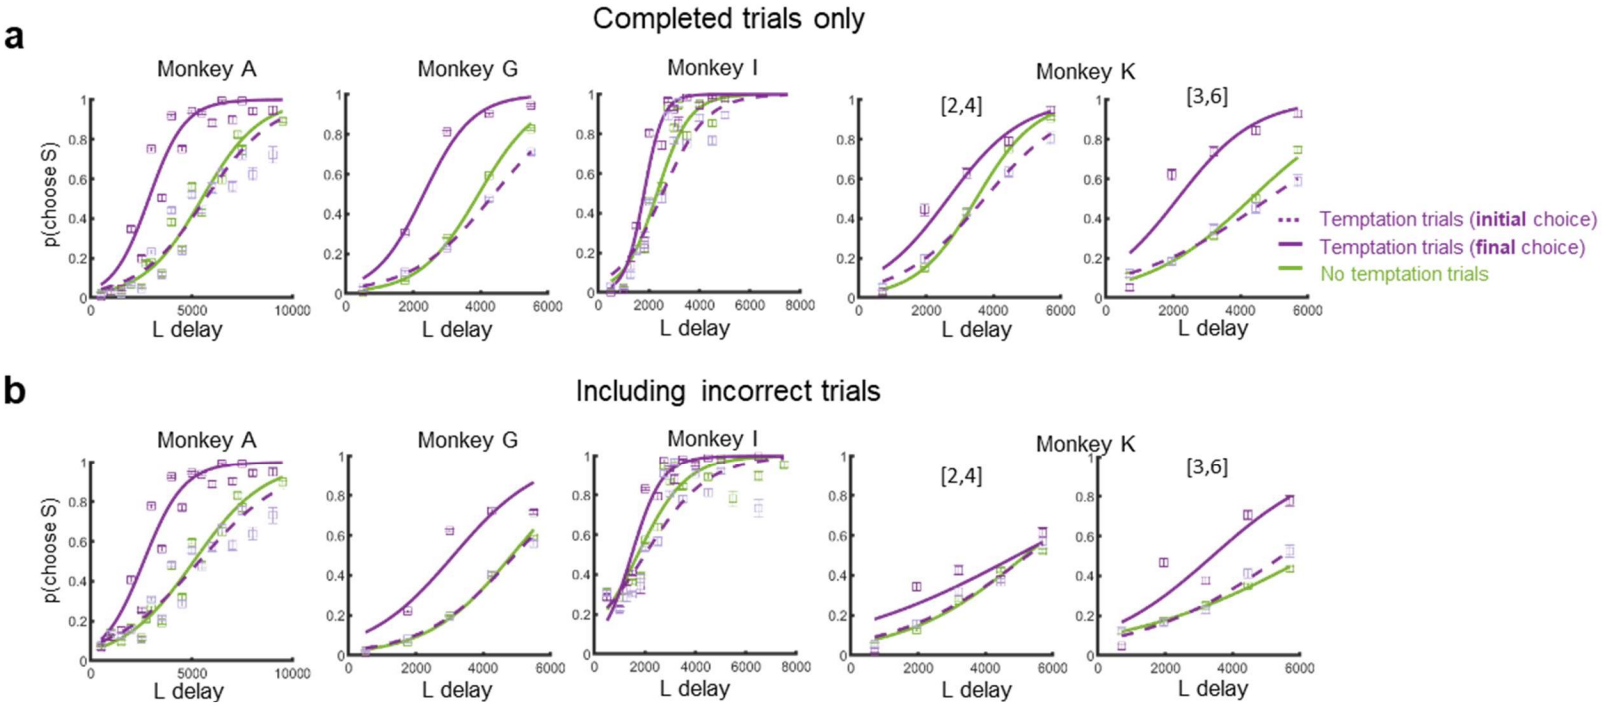

Supplementary figure 1 Choice functions. **a** Choice functions estimated by using completed trials only (i.e., the monkey did not break fixation and received a reward). The initial choice function of no temptation trials has a steeper curve than that of temptation trials, a difference more pronounced at the longer L delays. This apparent systematic difference in delay attitude is likely the result of only using completed trials in analyses. By design, when the monkey makes an L choice in no temptation trials, the monkey must remain fixated until the end of the delay in order to be rewarded. If the monkey breaks fixation, the trial is aborted and is not included in analyses. However, in temptation trials, after the monkey makes an L choice, the monkey can break fixation to settle for the S option, such that the trial is completed and becomes included in analyses. Thus, a higher proportion of initial L-choice trials were included in analysis for temptation trials than for no temptation trials (initial L-choice proportion for no temptation trials: 63.9%, temptation trials: 67.4%; paired t-test, t = 7.7804, p = 3.0378e-12). This also explains why the difference in initial choice functions between no temptation and temptation trials is more pronounced for longer L delays. **b** Choice functions estimated by including incomplete trials (i.e., including error trials where the monkey broke fixation before the end of delay and was not rewarded). Choice functions that include all trials, regardless of whether they were completed, exhibit no such significant difference in initial choice between no temptation and temptation trials (paired t-test, t = 0.6873, p = 0.4932).