Journal of Experimental Psychology:
Human Perception and Performance
2009, Vol. ●●, No. ●, 000–000

0096-1523/09/$12.00 DOI: 10.1037/a0016802

Motivational Influences on Response Inhibition Measures

Lauren A. Leotti and Tor D. Wager
Department of Psychology, Columbia University

Psychological research has placed great emphasis on inhibitory control due to its integral role in normal cognition and clinical disorders. The stop-signal task and associated measure—stop-signal reaction time (SSRT)—provides a well-established paradigm for measuring response inhibition. However, motivational influences on stop-signal performance and SSRT have not been examined. We conceptualize the stop-signal paradigm as a decision-making task involving the trade-off between fast responding and accurate inhibition. In 4 experiments, we demonstrate that performance trade-offs are influenced by inherent motivational biases and explicit strategic control. As a result, SSRT was lower when participants favored correct stopping over fast responding than when the same participants favored fast responding over correct stopping. We present a novel variant of the stop-signal task that uses monetary incentives to manipulate motivated speed-accuracy trade-offs. By sampling performance at multiple-trade-off settings, we obtain a measure of inhibitory ability that is independent of trade-off bias, and thus, more easily interpretable when comparing across participants. We present a working theoretical model to explain the effects of motivational context on response inhibition.

Keywords: inhibitory control, SSRT, speed-accuracy trade-offs

Inhibitory control refers to the ability to suppress a prepotent response, be it a behavioral process, impulse, or inappropriate thought. As a result, inhibitory control can be the difference between stepping back onto a curb and walking straight into traffic; between tasting desert and having a second piece of cheesecake; and between telling your spouse what she wants to hear and putting your foot in your mouth. There has been tremendous interest in inhibitory control, largely due to its apparent deficit in certain clinical disorders, such as attention deficit hyperactivity disorder (ADHD; Barkeley, 1997; Schachar, Tannock, Marriott, & Logan, 1995). Furthermore, inhibitory control warrants considerable attention because it has broad applications to diverse spheres of psychosocial functioning, including decision making (Bechara, Damasio, & Damasio, 2000), social competence (Eisenberg et al., 1996; Kochanska, Murray, & Harlan, 2000; Thorell, Bohlin, & Rydell, 2004), and emotion regulation (Ochsner, Bunge, Gross, & Gabrieli, 2002; Phan et al., 2005; Walcott & Landau, 2004).

Inhibitory control has been operationalized in laboratory paradigms that require inhibition of prepotent motor responses, such as the stop-signal paradigm, which has been shown to differentiate between individuals with inhibitory deficits (e.g., ADHD) and controls (Schachar, Mota, Logan, Tannock, & Klim, 2000) and is correlated with self-reported impulsivity (Logan, Schachar, & Tannock, 1997). The stop-signal paradigm provides a sensitive, quantitative measure of response inhibition, stop-signal reaction time (SSRT), that correlates with other measures of inhibition and executive control (Friedman & Miyake, 2004). In the stop-signal task, participants perform a speeded response task (e.g., indicating whether an arrow points to the left or right). On "go" trials, they are instructed to respond as quickly as possible. On "stop" trials, a stop-signal (e.g., a tone) is presented concurrently with or just after the target, and participants must withhold their response. SSRT (Logan, Cowan, & Davis, 1984) is based on a simple "race" model, which asserts that on a given stop-signal trial, a "go-process" races against a "stop-process." If the stop-process is faster than the go-process, a response will be successfully inhibited. With extensive evidence of reliability and face validity, SSRT has become a standard measure of response inhibition, used to study inhibitory control across stages of normal development (Band, van der Molen, Overtoom, & Verbaten, 2000; Kramer, Humphrey, Larish, Logan, & Strayer, 1994), as well as in the context of clinical populations (Badcock, Michie, Johnson, & Combrinck, 2002; Nigg, 2001; Schachar et al., 2000).

SSRT is assumed to provide reliable estimates of a fixed ability: inhibitory control. However, it is unclear the extent to which SSRT values reflect pure estimates of inhibitory control ability. Simulations of the stop-signal task demonstrate that variability in the primary task reaction time and the correlation between mean response time and true SSRT can both significantly influence the estimation of SSRT (Band et al., 2003; De Jong, Coles, Logan, & Gratton, 1990). Furthermore, several studies have found that SSRT systematically varies within subjects as a function of experimental manipulations of the expected probability or salience of the stop-signals (van den Wildenberg, van der Molen, & Logan, 2002; van der Schoot, Licht, Horsley, & Sergeant, 2005). Given the wide acceptance of SSRT as a measure assessing inhibitory control ability, it is important to determine whether SSRT estimates reflect

---

Lauren A. Leotti and Tor D. Wager, Department of Psychology, Columbia University.

We acknowledge Jennifer Rineer and Deanna Duval for their contributions in the collection of this data. We thank Ed Smith for reviewing early drafts of this manuscript.

Correspondence concerning this article should be addressed to Lauren Leotti, Department of Psychology, Columbia University, 406 Schermerhorn Hall, 1190 Amsterdam Avenue MC 5501, New York, NY 10027. E-mail: lak2116@columbia.edu

cognitive processes that are truly impenetrable to external influence.

The stop-signal paradigm, like other basic reaction time tasks, involves making decisions about whether and when to initiate responses. The paradigm asks participants to accomplish two competing goals: Respond as quickly as possible on a choice task, and withhold or cancel the response if a secondary stop signal is presented. The probabilities of responding and withholding depend on the decision-maker's response bias, or decision criterion, which reflects the subjective relative importance of the two goals. For example, an individual who values accurate inhibition of responses on stop trials over fast responses on go trials may be less likely to prepare motor responses to the go-stimuli than would an individual who places greater value on fast responding. Thus, the task involves an inherent trade-off between these goals, which we refer to as the "stop–go trade-off."

A key question is whether different trade-off biases could produce differences in performance on stop-signal tasks, even in SSRT. Many researchers assume that strategic influence on performance is minimal because participants are informed that speed on go trials and accuracy on stop trials is equally important. Even so, it is assumed that variations in stop-signal performance measures, such as go-response time and stop-signal response rate, should not lead to variations in SSRT. However, this has not, to our knowledge, been experimentally demonstrated. Strategic trade-offs (e.g., speed vs. accuracy) have been critically important in research on psychophysics (Green & Swets, 1966) and memory (Yonelinas, 1994), in which trade-off biases are often controlled with rewards or other psychophysical manipulations. By contrast, studies of inhibitory control, and specifically the stop-signal task, have not experimentally controlled or measured stop–go trade-off biases. As a result, it is unknown whether such trade-offs influence measures of inhibitory ability, such as SSRT. If individual differences in trade-off biases systematically influence SSRT, then differences in SSRT between individuals or even subject groups would not necessarily have the straightforward interpretation that makes the stop-signal task so popular.

In this paper, we present empirical evidence from four experiments that reveal the effects of motivational bias on stop-signal performance and measures of inhibitory ability. In Experiments 1 and 2, we demonstrate that individual differences in performance bias (stop–go trade-off) in a standard adaptive stop-signal task are predictive of individual differences in SSRT. These findings suggest that the SSRT measure depends on context-dependent processes, such as the interpretation of task goals or motivational orientation toward speed or accuracy. They also suggest that the standard task does not constrain behavior well enough to eliminate the bias in SSRT related to stop–go trade-off. We tested this hypothesis experimentally in two additional studies. In Experiment 3, we present a novel variant of the stop-signal task that uses monetary reward and punishment schedules to vary the motivational context favoring speed or accuracy, effectively shifting stop–go trade-off biases and producing systematic variations in SSRT. This experiment shows that experimental manipulations of the importance of go versus stop goals can influence SSRT, demonstrating that the use of the SSRT measure does not preclude their influence and so is not a pure measure of a person-level ability. Experiment 4 replicates this result and explores some of the conditions necessary to elicit performance shifts in the stop-signal task, demonstrating the important roles of both explicit strategic control (instructions about the relative importance of go and stop goals) and performance feedback (rewards and punishments that shape the stop–go trade-off). Collectively, these findings provide experimental evidence that SSRT does not provide a measure of trait inhibitory ability, but rather is influenced by motivational state. In the General Discussion, we present a preliminary theoretical model to explain motivational influences on SSRT.

Experiment 1

A popular approach to constraining performance on the stop-signal task involves the use of an algorithm that targets the same stopping accuracy across all participants. This tracking procedure (Slater-Hammel, 1960) was adopted to ensure reliable estimates of SSRT. Simulations have shown that the SSRT model produces stop latency estimates that are unreliable and biased when stop-signal response rates deviate substantially from 50%. Because SSRT is derived from the response time distribution on go trials, stop latency estimates are most reliable when sampling from the densest part of the response time distributions (50%) and are least reliable when sampling from the tails of the distribution. As a result, extreme trade-off biases (e.g., error aversion resulting in 90% stop accuracy) may compromise the reliability of stop latency estimates from the SSRT model. The tracking procedure has become a popular solution to address this issue.

To achieve 50% stopping accuracy, the onset time of the stop-signal relative to the go stimulus, or the stop signal delay (SSD), is updated throughout the experiment to change the level of difficulty of the task. In theory, the tracking procedure should eliminate potential influences of motivational bias on SSRT by controlling the correct stopping rate across participants. However, the tracking procedure itself may introduce another type of motivational bias into SSRT estimates for some participants. Because successful response inhibition is more probable if the SSD is shorter rather than longer, the tracking procedure lengthens the SSD every time the participant successfully inhibits a response and decreases the SSD every time the participants fails to inhibit a response to a stop-signal. If individuals value the "correct stopping" goal highly, they may continue to wait longer and longer, resulting in longer SSDs. This presents a problem because this feedback cycle can result in SSDs so long that the participant could never stop successfully without waiting for some period of time to see whether the stop signal will appear. In this case, the tracking procedure may actually undermine the face validity of the stop-signal task by changing it from a task of response inhibition, to a task of decision making, in which the participant decides to either respond immediately on half of all trials or to wait for a stop-signal to occur.

Although researchers acknowledge this is a potential problem, as it is common practice to caution participants to avoid such waiting strategies, there have been no reports of this type of behavior and how it influences SSRT. The tracking procedure targets 50% correct stopping rate, but does not constrain response times on the primary choice reaction time (RT) task. In fact, relatively few studies using the stop-signal task, either with or without the adaptive staircase procedure, provide any reinforcement of go goals. We surveyed 110 stop-signal studies and found that participants usually receive no feedback on performance (104/110 studies), nor do they receive any explicit penalties for slowing responses (103/110 studies). A larger proportion of studies actually provide feedback that would inadvertently promote slowdown of responses (e.g., accuracy on the primary choice response task). Moreover, stop goals are automatically reinforced after each stop trial because participants receive immediate internal feedback about their inhibition success. Although participants may be highly compliant, they may nonetheless perform differently on the task depending on inherent motivational orientations that are not constrained by the ambiguity of the task. As a result, the stop–go trade-off bias can vary between subjects even though stop accuracy may be controlled.

In Experiment 1, we investigated whether individual differences in trade-off bias predicted SSRT when a tracking algorithm was used to target 50% accuracy on stop trials. We expected that individual differences in stop–go performance trade-offs would exist despite the use of a tracking procedure, and that differences in performance bias would explain differences in SSRT. Furthermore, we expected the tracking procedure to be unsuccessful at targeting 50% stopping accuracy for a subset of participants, suggesting that SSRT estimates cannot be easily interpreted for many healthy adults.

Method

Participants

Thirty-four individuals (16 women, 18 men) from the Columbia University community participated in Experiment 1 (median age = 22 years). All individuals received monetary compensation for their participation in this study. This research was approved by the Columbia University Institutional Review Board. All participants provided informed consent prior to participation.

Stimuli

All instructions and stimuli were presented on a monitor using Psychtoolbox for Matlab (Brainard, 1997). On go trials, a 1 × 1 inch go-signal (green arrow) was presented centrally. The direction of the arrow (50% left, 50% right) on any given trial was random, but appeared equally often for both go and stop trial types. The go-signal remained on the screen until the participant responded by pressing the *B* or *N* keys on the keyboard with their right index and middle fingers to indicate whether the arrow pointed toward the left or right, respectively.

On stop trials, the go-signal was presented centrally, and after a delay, a visual stop-signal was presented behind the go-signal: a yellow circle with a diameter 1.5 times the width of the arrow. Stop-signal onset time relative to arrow onset varied based on participant performance, as described in the procedure below. The stop-signal remained on the screen for 1,000 ms.

Procedure

Speed block. Participants first completed a speeded RT) task, in which only go trials ($n = 70$) were presented. Participants were instructed to indicate the direction of the arrow as quickly and accurately as possible. During the speed block, an adaptive procedure was used to establish the RT cut-off, the point that divided fast and slow responses. If a response was faster than the RT cut-off, the message “Great!” was displayed for 500 ms, but if the response was too slow (RT > RT cut-off), the message “Too Slow!” was displayed. After 25 trials, the cut-off was updated on every trial to reflect the 70th percentile of the RT distribution for the last 25 trials. This tracking procedure encouraged participants to increase in speed throughout the block and provided a strategy independent baseline measure of an individual’s RT distribution that could be compared with the RT distribution in the subsequent stop-signal task.

Tracking block. Participants completed a mixed stop-signal block with 80 stop and 80 go trials randomly intermixed. A four-staircase procedure, described by Aron & Poldrack (2006), was used to titrate the SSD throughout the experiment. On a given stop-signal trial, the SSD was randomly chosen from one of the four staircases, with starting values of 100, 150, 200, and 250 ms. If inhibition was successful, the value of the current staircase was increased by 50 ms, making inhibition more difficult on a future stop trial. Conversely, if a response was made, the value of the current staircase was decreased by 50 ms, making inhibition more probable on a future stop-signal trial. These staircases are expected to converge around the SSD that results in stop-signal responses 50% of the time. Convergence should occur after about ten stop-signal trials. Examples of staircase movement throughout the block can be found in Figure 1. F1

Participants were instructed to respond as quickly and accurately as possible to the direction of the arrow and to try to withhold a response if the stop-signal occurred. Participants were told to avoid slowing down responses to try to improve stopping accuracy and explicitly instructed not to wait on any trials to see if a stop signal would occur. Go trials that were faster than the RT cut-off were awarded 10 points, and trials that were slower were penalized 10 points. Stop trials on which the participant successfully inhibited were awarded 10 points, and stop trials on which responses were made were penalized 10 points. Participants were instructed to try to earn as many points as possible. On every trial, a 500 ms interval followed stimulus offset and preceded a feedback screen where the amount won or lost was displayed for 1,000 ms. A randomly jittered interval (range: 500 to 1,000 ms) followed each trial.

Analyses

From the tracking block, we calculated stop-signal response rate (RR), or the proportion of stop-signal trials on which the participants failed to inhibit a response. We identified participants whose stop-signal RRs deviated at least 10% from the targeted 50%. These participants (9/34) were classified as “noncompliant” because the tracking procedure was unsuccessful at achieving the targeted 50% RR for these individuals. The remaining participants who achieved RRs within 10% of the targeted 50% were classified as “compliant.” Stop-signal performance measures were compared for compliant and noncompliant participants.

Mean RT was calculated for the speed block, for the first 40 go trials of the tracking block (early), and for the second 40 go trials of the tracking block (late). Incorrect go responses and go-trial RTs less than 50 ms or greater than three times the standard deviation of the block’s distribution were excluded from RT analysis. We also calculated the absolute amount of slowdown for the two halves of the tracking block—referred to as early- and late-slowdown—by subtracting the mean go-RT of the speed block from early-go-RT and from late-go-RT, respectively. We exam-

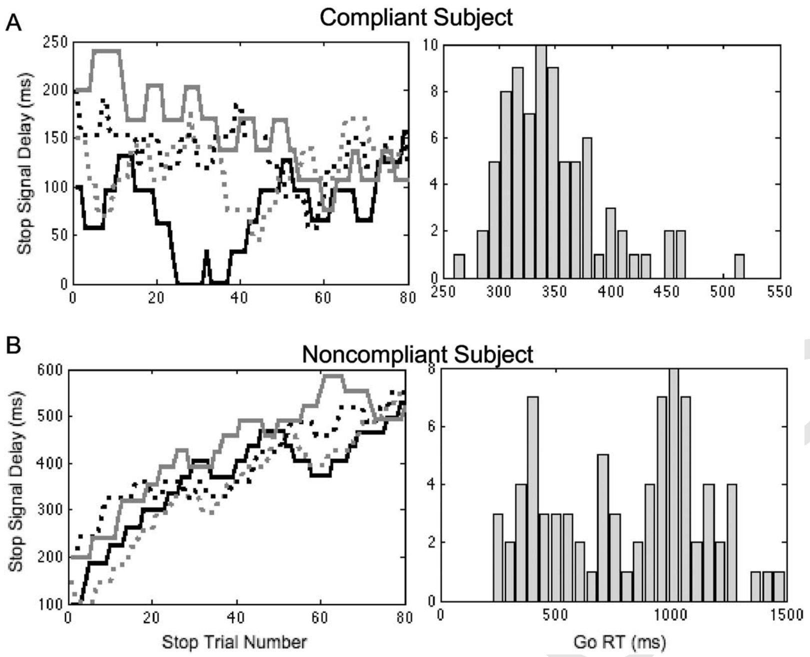

AQ: 5

Figure 1. Interpretations of performance differ for compliant (A) and noncompliant (B) participants on the stop signal delay (SSD) tracking procedure (Experiment 1). On the left, we display the (SSD) values as a function of stop trial number. Each line represents one of four staircases with different starting values. On a given stop signal trial, the SSD is chosen from one of the four staircases. If inhibition is successful, the SSD for that staircase will be increased by 50 ms and if inhibition is unsuccessful and a response is made, the SSD will be decreased by 50 ms. Ideally, the four staircases will converge on the SSD value that results in 50% stopping accuracy. On the right, we display the go reaction times (RT) distributions from the tracking block. (A) The compliant participant consistently responded fast throughout the experiment, resulting in a typical choice RT distribution that is unimodal and slightly positively skewed (right). The four staircases (left) converge and oscillate around SSD = 125 ms. (B) In contrast, the staircases for the noncompliant participant (left) steadily increase throughout the tracking block, suggesting the participant increasingly slowed responses to avoid responding on stop trials. The bimodal distribution of dramatically lengthened RTs (right) confirms this hypothesis and demonstrates a type of noncompliant behavior that can compromise the utility of an SSD tracking procedure.

ined the effects of compliance (compliant vs. noncompliant) and block (speed vs. early vs. late) on go-RT in a 2 (between-) × 3 (within-subjects) analysis of variance (ANOVA). By comparing go-RT for these separate epochs, we can demonstrate the extent that responses increasingly slow over the duration of the tracking block. We expected that late-go-RT would be significantly greater than early-go-RT, demonstrating that slowdown experienced in the tracking block relative to the speed block is more than would be expected simply by adding an additional task component (go + stop). Furthermore, we expected noncompliant participants to demonstrate this effect to a greater extent than compliant participants. For all ANOVAs, Huynh–Feldt corrected *p* values are reported as well as partial eta squared values (η_p^2).

SSRT was calculated for each motivation block according to the method outlined by Logan (1994). If the finishing time of the go-processes is faster than the finishing time of the stop-processes on a given stop trial, then a response should occur. In practice, the relative finishing time of the stop-processes is modeled as the percentile of the go-RT distribution equal to the observed stop-signal RR. For example, if a participant responded on 50% of stop-signal trials, then the stop finishing time would be equal to the median RT on go trials. SSRT, an estimate of stop latency, is equal to the difference between the stop finishing time and the stop-signal onset time. Because the tracking procedure is assumed to target RR = 50%, researchers have adopted the practice of subtracting the central SSD from the median of the go-RT distribution to estimate SSRT. The central SSD, estimated by averaging the last 10 values in each of the four staircases, is assumed to be the delay time that results in correct stopping 50% of the time (Aron & Poldrack, 2006).

SSRT was compared for compliant and noncompliant participants using an independent samples *t* test. Increasing slowdown of go responses over the course of the tracking block can be interpreted as a shift in performance bias favoring accuracy over speed, potentially mediated by explicit strategy or implicit motivational bias. Correlations between SSRT and mean go-RT on the speed block and slowdown measures demonstrate probable motivational biases in SSRT from the tracking procedure.

Results

Across all participants in Experiment 1, the average correct stopping rate was 0.56 (*SD* = 0.093), which was significantly different from the targeted 0.5, $t(33) = 3.67$, $p < .001$. Table 1 **T1**

Table 1
RT Slowdown, RR, and SSRT in Experiment 1

Factor	Speed RT, ms	Tracking RT (ms) Early	Tracking RT (ms) Late	RR	SSRT, ms
Compliant (*N* = 25)	343 (23)	438 (92)	438 (90)	0.49 (.05)	202 (40)
Noncompliant (*N* = 9)	347 (39)	566 (123)	698 (142)	0.31 (.07)	94 (59)

Note. Data are *M* (*SE*). Participants were defined as noncompliant if they responded to stop -signals less than 40% of the time. Targeted stop-signal RR was 50%. RT = reaction time; RR = response rate; SSRT = stop-signal reaction time.

compares performance for compliant and noncompliant participants.

One-way repeated measures ANOVA revealed that mean go-RT significantly slowed from the speed block ($M = 343$ ms, $SD = 31$ ms) to early-go-RT trials ($M = 463$ ms, $SD = 95$ ms), and slowed even further for late-go-RT trials ($M = 488$ ms, $SD = 124$ ms), $F(2,66) = 46.01$, $p < .001$, $\eta_p^2 = 0.53$. An increase in RT from the speed block to the tracking block is expected because the inclusion of stop-signal trials increases the task's complexity. However, the continued increase in RT from early trials to late trials on the tracking block suggests strategic slowdown. Indeed, by comparing RTs between the compliant ($n = 25$) and noncompliant participants ($n = 9$), we observed a significant interaction between compliance and block on RT, $F(2,66) = 31.1$, $p = < .001$, $\eta_p^2 = 0.49$; indicating that participants who deviated from the target 50% RR were also those that slowed down most during the course of the tracking procedure, presumably to wait for the stop signal.

SSRT values were not correlated with go-RT on the speed block ($r = .11$, $p = .53$), or with early slowdown ($r = –.26$, $p = .15$), but were highly correlated with late slowdown ($r = -0.58$, $p < .001$), suggesting that more effective inhibition (shorter SSRT) was not associated with basic processing speed but was associated with strategic slowdown during the tracking procedure. This relationship was largely driven by the noncompliant participants, who demonstrated significantly greater RT slowdown as well as significantly lower SSRT scores than compliant participants, $t(32) = 6.2$, $p < .001$.

Discussion

Although the tracking procedure may be successful in targeting 50% correct stops for many compliant participants, it proved to be unsuccessful for a substantial proportion of participants in Experiment 1. Despite instructions to weigh stop and go trials equally and to avoid strategic slowdown, some individuals demonstrated extreme error aversion and strategic slowdown of responses. Though researchers have not reported this type of behavior, to our knowledge, it has been acknowledged by other researchers as a potential concern, and some researchers have developed enhanced feedback procedures to further constrain strategic influences on the stop-signal task using the tracking procedure (Swylan, 2004). Thus, it is likely that the noncompliance observed in our sample is generalizable to other healthy adult populations.

Examination of individual participants go-RT distributions suggests that some participants engaged in deliberate decision making, resulting in a bimodal go-RT distribution (see Figure 1). As a result, even for participants who were successfully tracked with the paradigm to 50% correct stops, it is unclear whether SSRT actually reflects response inhibition ability, given participants may have been performing a qualitatively different task. Moreover, individual differences in strategic slowdown across all participants predicted individual differences in SSRT, such that stop latency was estimated as faster for individuals who slowed responses more throughout the task. This suggests, that perhaps even minor variations in performance trade-off bias, unconstrained or even exaggerated by the SSD tracking procedure, may systematically bias SSRT estimates of stop latency.

Because we did not collect participant reports of strategy used in this task, we cannot conclude whether the participants identified as noncompliant, intentionally disregarded the instructions to avoid strategic slowdown or whether their performance was influenced by implicit motivational biases toward accuracy. In Experiment 2, we compared explicit reports of strategy to performance to address this question.

Experiment 2

In Experiment 1, we demonstrated limitations to the tracking procedure when used in a stop-signal task with stop-signals occurring on 50% of all trials and performance feedback after each trial. By making stop and go trials equally probable and providing feedback after both trial types in Experiment 1, individual differences in trade-off bias were perhaps more apparent than they might have been had we used another design. Typically, researchers who use the tracking procedure do not provide trial-by-trial speed-related feedback and include a lower proportion of stop trials (e.g., 25%). In Experiment 2, we administered the same tracking procedure described in Experiment 1, but did not provide any performance feedback and lowered the stop-signal probability to 25% of all trials. We expected stop-signal RR and go-RT slowdown to be lower than what we observed in Experiment 1 because this version of the task disproportionately favors the go response. Despite the shift in these outcome measures, we nonetheless expected to observe similar patterns of bias that would be predictive of SSRT. To determine whether noncompliance or extreme trade-off bias was due to explicit strategy or implicit bias, we obtained participant reports of strategic bias. If noncompliant behavior was intentional, we would expect explicit reports of behavior to match stop-signal performance trends in response time.

Method

Participants

Thirty-two healthy adults (18 women, 14 men) ages 18 to 33 years (median = 22) participated in Experiment 2.

Stimuli and Procedure

All stimuli and procedures were identical to those used in Experiment 1, with two exceptions: (a) Feedback was not displayed after every trial to indicate whether go responses were fast enough or too slow, or whether inhibition was successful or not for stop-signal trials; and (b) a total of 360 trials were included in the tracking block, including 180 go trials and 80 stop-signal trials (25%). Because of the extended format of this version, a pause screen was inserted every 80 trials to prevent participant fatigue.

At the end of the block, half of the participants were asked the following: "About what percentage of your total effort on the task was focused on accurately inhibiting responses to the stop-signal (relative to effort focused on responding quickly to the arrow)?" The other half of the participants was asked the question in terms of effort focused on responding quickly relative to accurate inhibition. To respond, participants typed a single number from 0 to 100. Responses made in terms of quick responding were subtracted from 100 to establish the effort for accuracy. Explicit reports of effort for stop-signal accuracy were used for all participants.

Analyses

Because we expected stop-signal RR to be within 10% of the targeted 50%, due to the task's emphasis on speed, we used different criteria for identifying noncompliance than those used in Experiment1. In Experiment 2, we simply compared stop-signal performance and SSRT between those individuals who demonstrated the greatest slowdown across the task (top 25%, $N = 8$) to the remaining participants ($N = 24$). Stop-signal performance measures described in Experiment 1 were all included in the present analyses.

Results

Across all participants in Experiment 2, the average correct stopping rate was 0.53 ($SD = 0.059$), which was significantly different from the targeted 0.5, $t(31) = -2.74$, $p < .01$. Individuals identified as noncompliant demonstrated significantly higher correct stop rates ($M = .59$, $SD = .07$) than did compliant participants who demonstrated less extreme slowdown of responses ($M = .51$, $SD = .04$), $t(31) = 3.3$, $p < .005$.

One-way repeated measures ANOVA revealed that mean go-RT significantly slowed from the speed block to early-go-RT trials ($M = 71.8$ ms, $SD = 33.8$ ms) and slowed even further for late-go-RT trials ($M = 91.8$ ms, $SD = 57.9$ ms), $F(1,31) = 12.85$, $p < .001$, $\eta_p^2 = 0.28$. The amount of continued slowdown from early to late was significantly higher for the noncompliant group ($M = 68.9$ ms, $SD = 24.01$ ms) than the compliant group ($M = 5.1$ ms, $SD = 15.8$ ms), $t(31) = -7.1$, $p < .001$; confirming the basis of our classification of participants as compliant or noncompliant by the extent of this slowdown.

SSRT values were not correlated with go-RT on the speed block ($r = -.08$, p = .65), but were significantly correlated with early slowdown ($r = -.56$, $p = .001$), and even more highly correlated with late slowdown ($r = -.65$, $p < .001$). As in Experiment 1, this relationship was largely driven by the noncompliant participants, who demonstrated significantly greater RT slowdown as well as significantly lower SSRT scores ($M = 126.8$ ms, $SD = 75.6$ ms) than compliant participants ($M = 200.6$ ms, $SD = 42.1$ ms), $t(31) = 2.8$, $p = .02$.

On average, participants in Experiment 2 reported spending 33.5% of their effort on stopping accurately, and 66.5% of their effort on responding quickly to the arrow. Reports of explicit trade-off bias did not significantly differ between noncompliant participants ($M = 42\%$, $SD = 20.6\%$) and compliant participants ($M = 30.6\%$, $SD = 20\%$), $t(31) = -1.4$, $p > .1$; suggesting that differences in behavior were not necessarily driven by differences in explicit strategy, that is, participants were unaware of their bias. Consistent with this finding, explicit reports of trade-off bias did not significantly correlate with extent of RT slowdown during the task ($r = .3$, $p = .1$) or with SSRT ($r = -.17$, $p = .34$). **AQ: 1**

Discussion

Experiment 2 revealed that individual differences in stop–go trade-off biases predicted differences in SSRT estimates, even when the stop-signal probability is relatively low (25%). Thus, the findings presented in Experiment 1 likely are not representative of a special case of motivational bias induced by high stop-signal probability (50%) and thus greater bias toward accurate inhibition. Even under experimental conditions that promote speeded responses, a subset of individuals demonstrate an accuracy bias that undermines the utility of the tracking procedure to target 50% correct stops, and effectively compromises the validity of the SSRT measures based on performance on the task.

In Experiment 2, the low probability of stop-signals should favor fast responding over accurate inhibition. Participants reported favoring speed over accuracy on the task, yet a subset of participants demonstrated strategic slowdown that would suggest accuracy on stop trials was more important than speed on go trials. These findings indicate that performance trade-offs may be driven by implicit motivational biases rather than by explicit strategic control. This is important because researchers have only acknowledged the potential role of consciously executed strategy on stop-signal performance. Although the tracking procedure may attempt to constrain the influence of explicit strategy, findings from Experiment 2 suggest it does not constrain implicit motivational biases, which not only impact stop-signal task performance, but also influence SSRT estimates.

The findings from the SSD procedure do not imply a criticism of the tracking procedure as a whole, but merely illustrate that the procedure does not completely eliminate the influence of motivational bias for all individuals. Many groups have used it successfully, and a key factor may be the selection and training of participants. The tracking procedure has become popular because it helps eliminate potential bias in SSRT estimates that can occur when performance deviates substantially from 50% stop trial accuracy. Although the tracking procedure may be able to match stopping accuracy across most participants, individual differences in motivational bias still exist, and can influence SSRT estimates, as we observed in Experiments 1 and 2. Thus, by only controlling the stop-signal RR, the tracking procedure solves one source of bias in SSRT (estimating from the center of the distribution), but it also introduces another form of bias into SSRT estimates. We make this point with extreme cases of noncompliance, yet it is unknown how even minor variations in motivational bias could impact estimates of SSRT. It is important to point out the existence of this pattern of noncompliant behavior in a formal empirical study because, to our knowledge, no studies using the tracking procedure have reported any sort of quality assurance check of individual participants' data to confirm that the tracking procedure was even successful at achieving 50% RR for each participant. Thus, we cannot be entirely confident that SSRT is measuring the same construct—inhibitory ability—for every participant.

SSRT is generally accepted as a highly reliable and valid measure of response inhibition ability. However, the findings suggest that inhibitory ability, as measured by SSRT, depend on motivational factors. As a result, measures of inhibitory ability that do not account for stop–go trade-off bias may not reliably reflect a trait inhibitory ability. Although studies have demonstrated that SSRT has high test–retest reliability (Kindlon, Mezzacappa, & Earls, 1995; Tannock, Schachar, Carr, Chajczyk, & Logan, 1989), SSRT may be reliably measuring something other than inhibitory ability. Individual differences in interpretation of instructions and task goals, as well as noncompliant behavior observed on the tracking procedure for some participants in Experiments 1 and 2, suggest that people have inherent motivational biases, and it is probable that people approach the task with the same inherent biases in a reliable manner. As a result, experimentally manipulating and measuring variables like motivational bias and strategic control over performance may substantially improve our ability to produce reliable and valid measures of inhibitory ability.

Experiment 3

The previous studies demonstrated that stop-signal performance and measures of inhibitory control (e.g., SSRT) are influenced by motivational bias, reflected in the individual differences in the trade-off between fast responding and accurate inhibition. In the present study, we conceptualize the stop-signal paradigm as a decision-making task involving strategic allocation of resources and addressed two questions. First, does experimental manipulation of stop–go trade-off influence SSRT? A positive finding would provide further evidence that "cognitively penetrable" processes matter in the measurement of SSRT and are conflated with inhibitory ability if not measured or controlled for. Second, could experimental manipulation of SSRT provide a solution to measuring, and thus controlling for stop–go trade-off related bias?

The measurement of SSRT across multiple, experimentally manipulated stop–go trade-off settings was inspired by signal detection theory (Green & Swets, 1966), which has been used in a large number of studies demonstrating the separable influences of decision-making processes and perceptual/mnemonic sensitivity on perception and memory tasks. Signal detection theory (Green & Swets, 1966) suggests that trade-off bias (measured by a bias parameter) should be independent of the measured ability (measured by a sensitivity parameter). In Experiment 3, we adopted a model similar to the receiver operating characteristic (ROC) curves of signal detection theory to calculate response curves for varying strategic go speed and stop accuracy trade-offs. By examining stop-signal performance under multiple-decision criteria, we aimed to calculate an inhibition sensitivity measure that explains performance accounting for motivational bias and strategic control.

In this experiment, we tested the hypothesis that shifts in motivational context—the subjective relative importance of responding quickly on go trials vs. accurate inhibition on stop trials—would produce shifts in response inhibition measures (e.g., SSRT). We manipulated the motivational context for three trial blocks by differentially varying rewards and penalties for each trial type (stop vs. go), resulting in the systematic trade-off between accuracy on stop trials and speed on go trials. We examined the effect of motivational context on stop-signal performance measures and SSRT in two versions of the stop-signal task, which differed only by the proportion of stop-signal trials (25% vs. 50%).

To examine the trade-off between speed and accuracy, ideally (for optimal power) we would use equal numbers of stop and go trials (50%), which would allow trade-off bias to be unconstrained by trial frequencies. However, most stop-signal studies report using fewer stop-signal trials (e.g. 25%). Thus, it is important to determine if motivational context can influence stop-signal performance measures when stop trials are relatively infrequent. We hypothesized that motivational context would significantly influence stop-signal performance measures regardless of stop-signal probability. However we expect the magnitude of this effect would be smaller when the stop-signal probability was relatively low because in the 25% stop-signal condition, there is an existing bias toward speed due to the disproportionate number of go trials.

Method

Participants

Twenty-one individuals (15 women, 7 men) from the Columbia University community participated in Experiment 3 (median age = 23 year). Participants were randomly assigned to one of two versions of the stop-signal task. One version ($n = 11$) had 50% stop trials in each block and the other version ($n = 10$) had 25% stop trials in each block. Instructions and procedures were identical for both task versions.

Stimuli

All stimuli and trial timing were identical to those used in Experiment 1.

Procedure

Speed block. Participants first completed a speed block of 70 go trials, with RT cutoff titrated as in Experiment 1. Subsequent task parameters were derived from each individual's speed block RT distribution, allowing for perceived task difficulty to be controlled across participants. The RT cutoff for subsequent blocks was set to 1.5 times the median of the speed block RT distribution. In Experiment 3, SSD was fixed rather than titrated as in Experiments 1 and 2: the stop-signal was presented 200 ms before the individual's RT cut-off on all stop trials.

Mixed blocks. Participants then completed three stop-signal blocks. For the 50% stop-signal task, each mixed block had 60 go trials and 60 stop trials. For the 25% stop-signal task, each mixed block had 90 go trials and 30 stop trials. Participants were instructed to respond as quickly as possible on all trials, but to try to withhold responses if a stop-signal occurred. They were explicitly instructed to avoid slowing down responses on go trials to improve stopping accuracy, as go and stop trials were equally important. On each trial, feedback was provided, as described in the tracking block of Experiment 1. On go trials, responses faster than the RT cutoff were rewarded with monetary gains and responses slower than the RT cutoff were penalized with monetary losses. On stop trials, participants received monetary gains for successfully inhibiting a response and were penalized with monetary losses for failing to do so. If participants responded before the go-signal appeared, they were penalized 100 points to discourage premature responses and to promote attentive performance. Participants were instructed to earn as many points as possible, for points were translated into bonus reward money. The point total was displayed at the end of each block. On average, participants earned \$4.00 bonus money (range = \$0.00 to \$9.00).

Monetary rewards and punishments were titrated adaptively during task performance, depending on the difference between an individual's current trade-off bias and the targeted trade-off for each of three motivational contexts: go bias, no bias, and stop bias. Block order was counterbalanced across participants. On each trial, the stop–go trade-off (SGT) bias was defined as the average of the success rate on the last ten stop trials (correct stops) and the failure rate on the last 10 go trials (RT > RT cutoff). SGT values closer to zero reflect a stronger bias for responding quickly on go trials, and values closer to one reflect a stronger bias for stopping correctly on stop trials. Each motivational context block targeted a different trade-off bias (go bias = 0.3, no bias = 0.5, stop bias = 0.7). In practice, if a participant initially demonstrated a trade-off bias toward stopping accuracy in the go bias block, then rewards and penalties for stop trials would be decreased in magnitude and rewards and penalties for go trials would increase in magnitude.[1] The titration began with equal rewards (+50) and penalties (–50) for both trial types, and was bounded at a maximum of 100 points and a minimum of 0 points for rewards, and a maximum of 0 points and a minimum of –100 points for penalties. Rewards and punishments initially varied as the participants adapted to the payoffs, but increasingly stabilized over time.

Participants were explicitly trained to strategically trade off between speed and stop accuracy in response to changes in rewards and punishments. For example, participants were told, "If you notice that you are receiving larger rewards and penalties for your performance on go trials than on stop trials, you should focus on making speeded responses on go trials to earn the most money possible." Three practice blocks of 20 trials each proceeded the three mixed blocks. Participants were not informed of the number of trials in a block, or of the probability of stop trials in a block.

Analyses

One-way repeated measures ANOVA models were used to examine the influence of motivational context (3 levels: stop-bias, no-bias, go-bias) on four outcome variables: SGT, stop-signal RR, RT slow down relative to the speed block, and SSRT. Separate analyses were conducted for each task version (25% vs. 50% stop trials), and combined data were also analyzed with 2 (between: task version) × 3 (within: motivational context) mixed-effects ANOVAs.

To determine the amount of "strategic" slowdown for each participant, we calculated the differences between the mean of the RT distribution for each motivation block and the mean of the RT distribution for the speed block. Incorrect go responses and go-trial RTs less than 50 ms or greater than three times the standard deviation of the block's distribution were excluded from RT analysis.

For each block, a participant's average SGT was defined as the average of the stop trial accuracy rate and the go trial failure rate across the entire block. Thus, a given SGT value reflects a particular trade-off between time outs (go-RT > RT cutoff) and correct stops, shown by the diagonal lines in Figure 2 for three SGT values (0.3, 0.5, and 0.7). SGT values vary between 0 and 1, with lower values reflecting a stronger speed bias, and higher values reflecting a stronger accuracy bias. An SGT of 0.5 indicates no bias toward speed or accuracy. As Figure 2 illustrates, holding SGT constant, better inhibitory efficacy would lead to higher correct stop rates and lower time-out rates. Thus, the diagonal lines in Figure 2 reflect performance at a constant SGT but for varying levels of inhibitory efficacy. As a result, SGT is a measure comparable to the bias measure used in signal detection theory, and the each SGT value approximates a position along a trade-off curve (similar to an ROC curve) defined by the participant's SSRT. However, the curves in Figure 2 differ critically from typical ROC curves because here, response curves can be both concave (good inhibitor) and convex (poor inhibitor) due to the nonlinear trade-off between correct stop rate and time-out rate, which is the proportion of go trials that were penalized for being too slow.

For each of the three blocks, correct stopping rate (y) was plotted against the time-out rate (x). From the data, we estimated the participant's speed-accuracy trade-off curve using a general linear model. Area under the curve (AUC) provides a measure of the efficiency of inhibition. Higher values of AUC suggest that less strategic slowdown is necessary to increase stopping accuracy.

Results

Overall accuracy on the primary response task (left/right discrimination) was high ($M = 94\%$, $SD = 5.0\%$). Mean response time on the speed block was 264.6 ms ($SD = 33.3$ ms). Independent-samples t tests revealed that block order did not influence SGT for each motivational context (all $ps > .1$). In the 3 (motivational context) × 2 (task version) analysis, there was a significant main effect of motivational context on SGT, $F(2,38) = 70.99$, $p < .001$, $\eta_p^2 = 0.79$; RR, $F(2,38) = 55.91$, $p < .001$, $\eta_p^2 = 0.75$; go-RT slowdown, $F(2,38) = 18.81$, $p < .001$, $\eta_p^2 = 0.50$); and SSRT, $F(2,38) = 4.7$, $p = .024$, $\eta_p^2 = 0.20$. As rewards and penalties increasingly favored correct stopping over fast responding (go bias block > no bias block > stop bias block), SGT values increased, RR decreased, RTs increased, and critically, SSRT decreased (see Figure 3). Main effects of motivational context

---

[1] Amount of change in rewards depended on (a) trial number: initial reward jump size was 25 and exponentially decayed to a minimum jump size of 5; and (b) difference between true trade-off bias and targeted trade-off: jump size was multiplied by this difference to determine shift in rewards and penalties.

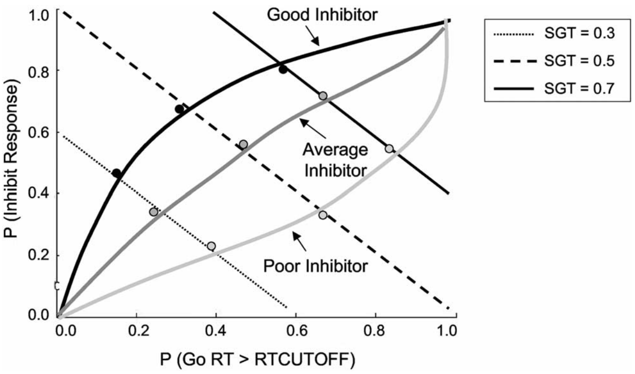

Figure 2. Stop-signal task performance was sampled at multiple-stop–go trade-off (SGT) settings to obtain an inhibitory control measure that is orthogonal to SGT bias. In Experiment 3, rewards and penalties for performance were titrated to target each of three SGT settings (diagonal lines) for the go bias, no bias, and stop bias blocks, respectively. For example, if a participant adopts a strategy that favors fast responding over accurate inhibition, performance would lie somewhere along the diagonal where SGT = 0.3. A measure of inhibitory ability can be estimated by estimating the speed-accuracy trade-off curve that best explains performance across the three SGT points. SGT should be independent of the curve because the same SGT setting is possible anywhere along its diagonal, as demonstrated with three sample participants of varying inhibitory ability: a relatively good inhibitor (concave curve), an average inhibitor (middle gray curve), and a relatively poor inhibitor (convex curve).

were also significant on all outcome variables for each task version
T2 independently (see Table 2).

There were significant interactions between task version and motivational context on SGT, $F(2,38) = 10.4$, $p < .001$, $\eta_p^2 = 0.35$; RR, $F(2,38) = 8.13$, $p = .001$, $\eta_p^2 = 0.3$; and go-RT slowdown, $F(2,38) = 8.42$, $p = .002$, $\eta_p^2 = 0.31$. As predicted, larger changes in SGT, RR, and go-RT slowdown were observed across motivational contexts for the 50% task version than for the 25% task version. However, the interaction between task version and motivational context on SSRT was not significant, $F(2,38) = 1.8$, $p = .19$, $\eta_p^2 = 0.09$; suggesting that motivational effects on SSRT were roughly equivalent across 25% and 50% stop-trial versions.

Compared to the 50% stop-signal task, the 25% stop-signal task produced higher overall RRs, $F(1,19) = 6.58$, $p = .010$, $\eta_p^2 = 0.26$; less go-RT slowdown, $F(1,19) = 19.4$, $p < .001$, $\eta_p^2 = 0.51$; and longer SSRTs, $F(1,19) = 13.4$, $p = .002$, $\eta_p^2 = 0.41$; across all blocks. The task versions did not differ on overall SGT, $F(1,19) = 0.036$, $p = .85$, $\eta_p^2 = 0.002$.

To successfully estimate a stop–go trade-off curve, there must be sufficient variability in the modeled data. Individuals who did not substantially shift performance across the different motivational contexts showed little to no variability in the timeout rate and/or the correct stopping rate. As a result, curves that best fit the data of these participants may not accurately portray inhibitory efficiency, as measured by AUC, because performance was not sampled at varying trade-off points. To ensure that the AUC metric has the same interpretation across participants, the curves should be fit to data occurring at similar trade-off points (e.g. 0.3, 0.5, and
F4 0.7) for all participants (see Figure 4A) Participants who demonstrated a shift in SGT bias less than 0.20 units between the go bias and stop bias blocks were excluded from AUC analysis. For the remaining 15 participants, the mean AUC was 0.45 (*SD* = 0.18), with values ranging from 0.16 to 0.72 (Figure 4B). Because participants in the 25% stop-signal condition demonstrated significantly less shift in SGT across blocks, reliable estimates of the stop–go trade-off curve could be made for relatively few participants (*N* = 3) in this condition. As a result, AUC could not be compared across subjects in the 25% and 50% stop-signal tasks.

Discussion

The results of Experiment 3 demonstrate that stop signal task performance and measures of response inhibition are influenced by motivational context. Participants responded more accurately on stop trials (lower RR) and responded more slowly overall, as rewards and punishments increasingly favored stop accuracy over speed. The significant change in SGT across blocks suggests the adaptive titration of rewards and penalties was successful in producing the desired performance shifts. Critically, SSRT was greater when a participant demonstrated a bias toward speed and was lower when a participant demonstrated a bias toward stop accuracy. Based on these results, we might conclude that inhibitory control is worse when favoring a speed bias and better when favoring an accuracy bias. Alternately, inhibitory control ability may not actually be changes with motivational shifts, but rather,

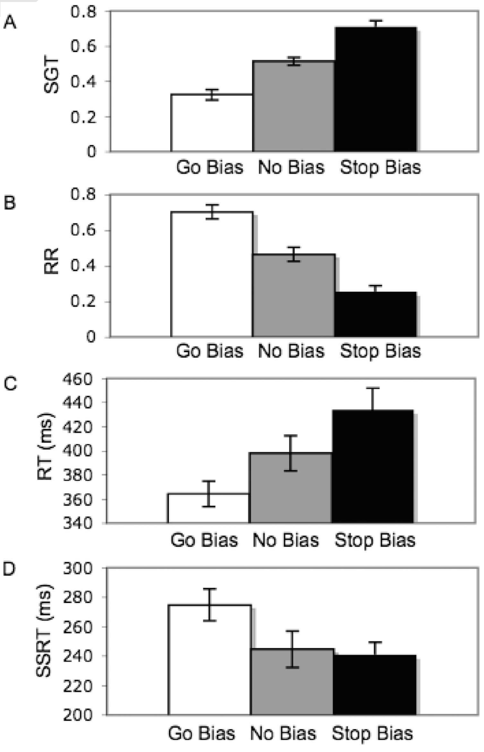

Figure 3. Variations in rewards and punishments result in strategic changes in stop signal performance in Experiment 3. The figures display, respectively, that as payoffs increasingly favor stop accuracy over speed: (A) SGT values increase; (B) stop-signal RR decreases; (C) go response times increase; and (D) SSRT estimates decrease. SGT = stop–go trade-off; RR = response rate; SSRT = stop-signal reaction time.

Table 2
Variation in SGT, RR, RTs, and SSRT Across Motivational Contexts in Experiment 3

Measure	Go bias	No bias	Stop bias	*F*	*p* <
SGT					
25%	0.42 (0.04)	0.53 (0.04)	0.63 (0.05)	16.6	.001
50%	0.29 (0.04)	0.53 (0.03)	0.79 (0.04)	62.4	.001
RR					
25%	0.49 (.05)	0.35 (0.06)	0.26 (0.05)	9.28	.005
50%	0.79 (.05)	0.50 (0.05)	0.24 (0.04)	59.9	.001
RT Δ, ms					
25%	–15 (13)	–3 (20)	8 (25)	3.75	.05
50%	47 (12)	105 (18)	171 (23)	20.8	.001
SSRT, ms					
25%	294 (21)	268 (23)	264 (14)	16.2	.005
50%	257 (19)	214 (20)	209 (13)	19.3	.005

Note. Data are *M* (*SE*). Statistics are from separate repeated measures analyses of variances for each task version (25% vs. 50% stop-signal trials). SGT = stop-go trade-off; RR = response rate; RT = reaction time; RT Δ = relative to speed block RT; SSRT = stop-signal reaction time.

our ability to measure stop latency does not account for motivational context and associated performance shifts.

We also found that the shifts in performance induced by changes in motivational context were less extreme when stop-signal probabilities were lower (i.e., in the 25% version). However, there was no significant difference between task versions on the change in SSRT across motivational contexts. We had hypothesized that stop-signal performance measures, including SSRT, would show less change across motivational contexts in the 25% stop-signal task than in the 50% stop-signal task because there is less competition for resources, and thus less SGT trade-off, when the stop-signal probability is relatively low. Furthermore, we expected motivational context to have a smaller impact on SSRT in the 25% task relative to the 50% task because response inhibition should be more difficult in the 25% task version, and studies of attentional control have found that extraneous dimensions exert less impact on performance when the task is sufficiently difficult (Lavie, 2005). The significantly higher SSRT scores for the 25% task relative to the 50% task suggest the 25% task was indeed more difficult. Nonetheless, the influence of motivational context on SSRT did not vary across the two versions of the task, suggesting that motivational context is a relevant predictor of stop-signal performance and SSRT regardless of task difficulty.

The stop-signal task variant outlined in Experiment 3 has the potential to provide estimates of inhibitory ability, as well as measures of individuals' strategic bias and reward and punishment sensitivity (not examined here). This may serve as an important tool for characterizing clinical disorders, such as ADHD, that have been associated not only with impaired response inhibition (Schachar et al., 2000), but also with motivational deficits (Slusarek, Velling, Bunk, & Eggers, 2001) and altered processing of rewards (Luman, Oosterlaan, & Sergeant, 2005). If SSRT varies with SGT setting, as observed in Experiment 3, then estimates of stop latency confirming inhibitory control deficits in clinical populations may reflect different motivational or strategic concerns rather than differences in inhibitory ability. Experiments should account for these strategic and motivational biases in some way, either by simply measuring bias (e.g., SGT measure), attempting to control it (e.g., via feedback and reinforcement), or both. One advantage of the procedures we employ here is that we control strategies and performance trade-offs using quantitative manipulations of objective payoffs, rather than relying solely on instructions to participants. Researchers wishing to implement the procedure should manipulate motivational context in a way that, ideally, targets approximate RR values of 30%, 50%, and 70% for each participant. This could be done by manipulating motivational context, or in principle by using adaptive tracking procedures targeting multiple-performance accuracies (e.g., by varying SOA), though other methods were not tested here.

By sampling performance at multiple-SGT points, it may be possible to establish a measure of inhibition efficiency that accounts for stop-signal performance in various motivational contexts. The area under the speed-accuracy trade-off curve (AUC) is one such measure. Higher AUC values reflect more efficient inhibitory ability, that is the participant requires less slowdown of responses to achieve greater stopping accuracy. An inefficient inhibitor (lower AUC) would require greater response slowdown to achieve the same level of stopping accuracy. To use a metric like the AUC, it is necessary to target similar trade-off biases across participants. As Figure 4A illustrates, modeling invariant data can result in trade-off curves that fit the data well, but have unknown predictive validity, and which cannot easily be interpreted. Most of the participants in the 50% stop-signal task achieved the targeted SGT biases, allowing for successful estimation of the speed-accuracy trade-off curve and AUC. However, most of the participants in the 25% stop-signal task did not reach the targeted trade-off points, suggesting that the reward titration used was not appropriate for both groups. Enhancement of this procedure would likely improve the ability to target various trade-off points in a 25% stop-signal task so that summary measures like AUC could be estimated reliably. Using the GLM to find the AUC is only one approach to measuring inhibition under varying motivational contexts. Additional research is necessary to determine if other statistical models would provide a better explanation of stop-signal performance. AQ: 2

Experiment 4

Based on the results of Experiment 3, we can conclude that motivational context, defined in terms of monetary gains and losses, can shape performance bias, which influences SSRT. However, we cannot determine whether such shifts in performance were the result of explicit strategy, or whether performance was simply shaped by reinforcing feedback in a way that may or may not be cognitively accessible to participants. In Experiment 4, we addressed this question by comparing a group of participants who received instructions to strategically shift behavior in response to reinforcement (as in Experiment 3) to a group of participants who only received instructions to strategically shift behavior without receiving any trial-by-trial feedback or reinforcement. Because the groups differ only by performance feedback, any differences observed between the groups can be attributed to properties of reinforcement.

If explicit strategic control of behavior mediates performance shifts, we would expect to see comparable changes in stop-signal task performance and SSRT as a function of variations in motivational context. If variations in stop-signal performance and SSRT

A Curve estimation depends on shift in stop-go tradeoffs (SGT)

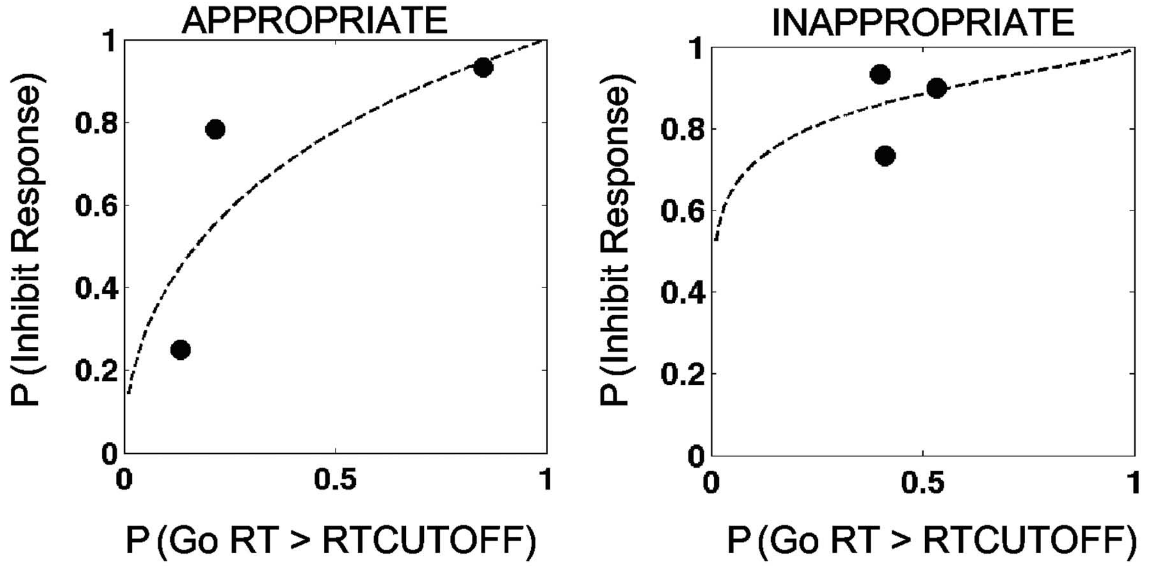

B Tradeoff curves for subjects with adequate SGT (n=15)

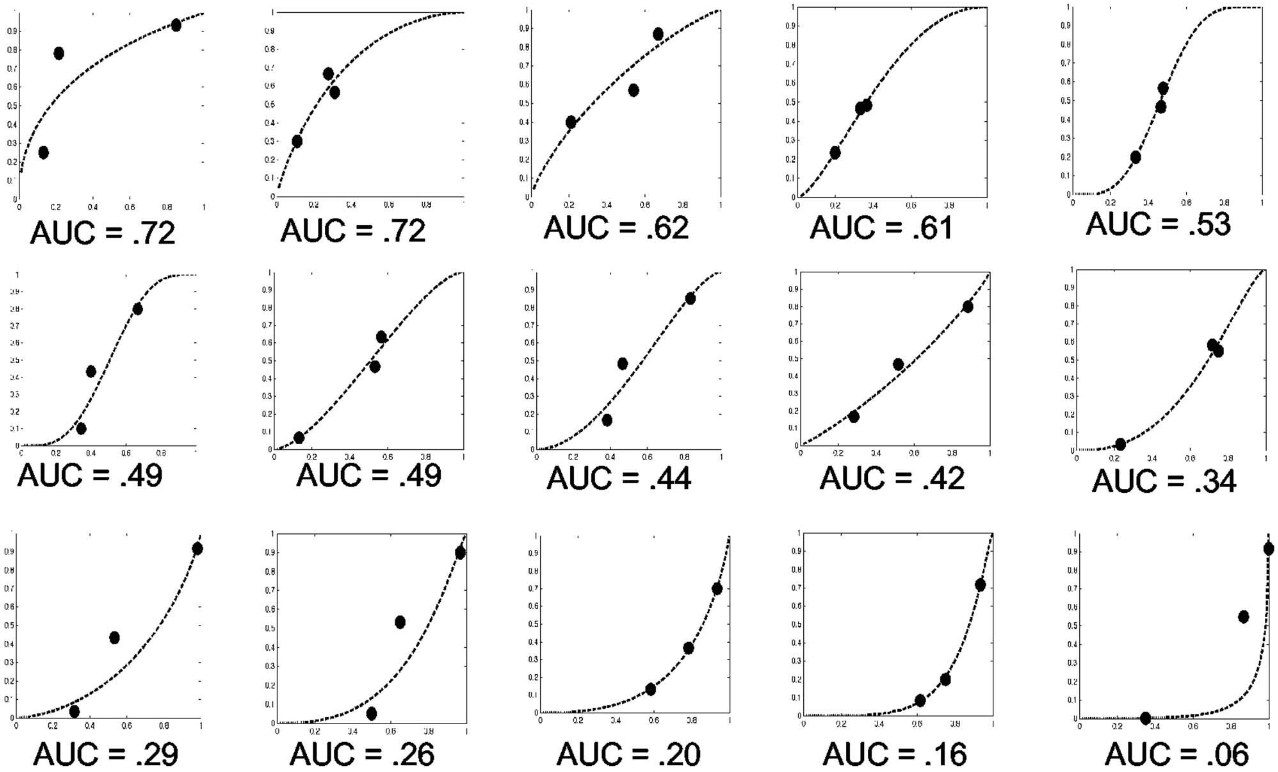

AQ: 6

Figure 4. (A) Example SGT curves based on data from two Experiment 3 participants. Successful estimation of the speed-accuracy trade-off curve depends on appropriate variability in the SGT settings sampled. The participant on the left shifted SGT adequately (AUC = 0.72), whereas the participant on the right did not shift SGT settings at all (AUC = 0.85). As a result, the curve that best explains the data on the right does not have very good predictive validity at other SGT settings, and so accuracy of the curve and of the resulting AUC measure are questionable. (B) Trade-off curves for the 15 participants with appropriate shifts in SGT across motivational contexts. SGT = stop–go trade-off; AUC = area under the curve.

were driven by reinforcement learning in a nonstrategic way, rather than by explicit strategic control, we would expect to observe gradual shifts toward the targeted trade-off bias. We hypothesized that stop-signal performance and SSRT would vary systematically with motivational context for both groups, demonstrating that explicit strategic control can influence stop-signal performance and SSRT.

Method

Participants

Thirty-one individuals (12 men, 19 women) ages 18 to 35 years (median = 23) participated in Experiment 4. Participants were randomly assigned to the reinforcement group ($n = 16$) or to the strategy-only group ($n = 15$). Participants were informed they could earn bonus money for good performance, but were not given the criteria for earning bonus money. All participants were paid $12 for their participation.

Stimuli

All stimuli are identical to those described in Experiment 1. All trials lasted 1,200 ms. For the reinforcement group, feedback was displayed for 500 ms after each trial. No feedback was provided for the strategy-only group. Between trials, the interval was randomly selected from a list of values between 300 to 500 ms in 50 ms intervals. During this interval, a fixation cross was displayed centrally. On go trials, arrows were presented until a response was made or 1,200 ms had elapsed. The fixation cross was displayed for any time remaining (1,200—go-RT). On stop trials, stop signals were presented after the participant specific delay (based on values derived from speed block). Both the arrow and the stop signal then remained on screen until 1,200 ms had elapsed from the trial start.

Procedure

Speed block. Participants first completed a speed block of 70 trials. RT cutoff and stop signal onset times were determined from the speed block as in Experiment 3.

Mixed blocks. Participants were given a practice block of 100 trials (30% stop), on which they were instructed to respond to the arrow as quickly and accurately as possible and to try to inhibit their response if the stop-signal appeared. They were instructed to avoid delaying responses to improve stopping accuracy and that speeded responses and correctly stopping were equally important. Next, participants completed two strategy blocks, presented in counterbalanced order. Each block had 70 go trials and 30 stop trials randomly intermixed. Before each block, the phrase “Speed Next” or “Accuracy Next” was presented for 3,000 ms to indicate which type of performance should be favored in the upcoming trial block. For the go bias block, all participants were instructed to focus 70% of their effort on making fast responses and 30% of their effort on inhibiting responses when the stop-signal occurred. For the stop bias block, participants were instructed to focus 30% of their effort on making fast responses and 70% of their effort on correctly inhibiting responses to the stop-signal.

Participants in the reinforcement condition were informed they would receive feedback after every trial to indicate whether they were successful. Instructions were virtually identical to those received by participants in Experiment 3. Rather than the reward titration procedure used in Experiment 3, we used fixed rewards and penalties for each block in Experiment 4. On the go bias block, responses faster than the RT cut-off received 20 points, and slow responses lost 20 points; successful inhibition resulted in 10 points, and stop-signal responses were penalized 10 points. On the stop bias block, fast responses won 10 points and slow responses lost 10 points; successful inhibition resulted in 30 points and stop-signal responses were penalized 30 points. Participants were told to try to earn as many points as possible.

Analyses

Using 2 (between) × 2 (within) ANOVAs, we examined the effects of feedback type (reinforcement vs. strategy only) and motivational context (speed bias vs. accuracy bias) on SGT, stop-signal RR, and go-RT slowdown.

With respect to SSRT, we did not target specific SGT biases as in Experiment 3, and thus we expected that individuals would differ in both the magnitude of shift in SGT and (as a result), shift in SSRT induced by the feedback/strategy manipulation. Thus, we included a strategic shift measure as a covariate in the analysis of feedback and motivational context on SSRT. Strategic shift was operationalized as the difference between mean go-RTs (accuracy – speed) divided by difference in stop accuracy (1 – RR) across blocks. Thus, strategic shift reflected the overall change in SGT between blocks, as well as the efficiency of trading off between speed and accuracy: lower shift scores reflect less strategic slowdown per unit increase in stopping accuracy.

To examine the time course of shifts in performance trade-off bias, we calculated the RR for every 10 stop-signal trials in a motivation block. Within the speed block, and separately within the accuracy block, we examined the RR trajectory across the two participants groups in a 3 (within—time: early, middle, late) × 2 between—reinforcement vs. strategy only) ANOVA. If strategy shifts are achieved through reinforcement learning, stop-signal RR should gradually approach the targeted trade-off for the reinforcement group. Otherwise, the targeted trade-off should be achieved in the first 10 trials of the block (early), suggesting the trade-off bias was induced by adopting an explicit strategy.

Results

Performance

Accuracy on the left/right discrimination task was high overall ($M = 98.1\%$, $SD = 0.018$). Mean response time on the speed block was 324.6 ms ($SD = 37.9$). Neither accuracy nor speed block RT differed by feedback type (all ps $> .5$). SGT for go bias and stop bias blocks did not depend on which block was presented first (all ps $> .1$).

The results of Experiment 4 are summarized in Table 3. As in Experiment 3, we observed significant main effects of motivational context on SGT, $F(1,29) = 123.7$, $p < .001$, $\eta_p^2 = 0.81$; stop-signal RR, $F(1,29) = 140.3$, $p < .001$, $\eta_p^2 = 0.83$; and go-RT slowdown, $F(1,29) = 64.8$, $p < .001$, $\eta_p^2 = 0.69$. Between-subject groups there was a significant main effect of feedback on SGT, $F(1,29) = 24.0$, $p < .001$, $\eta_p^2 = 0.453$; RR, $F(1,29) = 6.3$, $p =$ T3

Table 3
Variation in SGT, RR, RTs, and SSRT Across Motivational Contexts in Experiment 4

Context	SGT	RR	RT Δ, ms	SSRT, ms
Go bias block				
Reinforcement	.33 (.03)	.75 (.04)	15 (10)	220 (8)
Strategy-only	.50 (.04)	.65 (.04)	65 (10)	206 (9)
Stop bias block				
Reinforcement	.65 (.03)	.37 (.04)	82 (15)	211 (8)
Strategy-only	.84 (.03)	.23 (.04)	153 (16)	192 (8)

Note. Data are *M* (*SE*). For SSRT, means reported are adjusted for model including strategy shift as a covariate. SGT = stop-go trade-off; RR = response rate; RT = reaction time; RT Δ = relative to speed block RT; SSRT = stop-signal reaction time.

.018, $\eta_p^2 = 0.18$; and go-RT slowdown, $F(1,29) = 15.2$, $p = .001$, $\eta_p^2 = 0.34$. Relative to the reinforcement group, the strategy-only group demonstrated higher average SGT biases, lower average stop-signal RR, and greater strategic slowdown across both motivational context blocks. There were no significant interactions between feedback and motivational context (all $ps > .1$).

SSRT

Consistent with our expectations, strategic shift varied considerably across subjects. After removing an extreme outlier with a AQ: 3 shift score of 1,389 (IQR = 112 to 287), we found that strategic shift was highly correlated with shift in SSRT across blocks ($r = –.68$, $p < .001$). Figure 2 illustrates that participants with the lowest shift scores had the largest difference in SSRT (go SSRT > stop SSRT), and participants with the highest shift scores had the most negative change in SSRT (go SSRT < stop SSRT). Strategic shift was not significantly different between feedback groups, $t(28) = -0.48$, $p = .64$. In an analysis of covariance (ANCOVA) that controlled for strategic shift, the main effect of motivational context on SSRT was significant: SSRT was higher for the go bias block relative to the stop bias block, $F(1, 27) = 4.2$, $p < .05$, $\eta_p^2 = 0.14$.

Immediate Versus Gradual Shifts in Strategy

T4 Table 4 displays the RRs as a function of time (every 10 stop trials) for each of the motivation blocks. In the go bias block, there was a significant main effect of group, $F(1,29) = 4.18$, $p = .05$, $\eta_p^2 = 0.13$; such that RR were higher overall for the strategy-only group. The main effect of time on RR was marginally significant, $F(2,58) = 2.79$, $p = .07$, $\eta_p^2 = 0.08$; and there was a significant interaction between group and time on RR, $F(2,58) = 5.43$, $p = .007$, $\eta_p^2 = 0.16$. Throughout the speed block, RR was consistent for the reinforcement group but increased over time for the strategy only group. In the stop bias block, there was a significant main effect of group, $F(1,29) = 7.55$, $p = .01$, $\eta_p^2 = 0.21$; with RR being higher overall for the strategy-only group. There was also a significant main effect of time on stop-signal RR, $F(2,58) = 6.21$, $p = .004$, $\eta_p^2 = 0.18$. In both participant groups, RR increased over the course of the block. However, the amount of increase was not significantly different between groups, $F(2,58) = 0.562$, $p = .573$, $\eta_p^2 = 0.02$.

Discussion

Experiment 4 replicated the results of Experiment 3 with respect to the influence of motivational context on stop-signal performance. Moreover, those participants who only received instructions to shift strategy demonstrated similar shifts in performance bias as those who were given trial-by-trial feedback and reinforcement for performance. The lack of significant interactions between motivational context and feedback indicate that the observed changes in stop-signal performance in both experiments is likely due to conscious shifts in strategy rather than more automatic processes that could have occurred in the course of reinforcement learning.

However, it is possible that both explicit strategy and reinforcement learning could produce the same pattern of results independently. In the present study, we did not include a group of participants who received trial-based reinforcement but no strategy instructions. The feedback group also received strategy instructions and only differed in the receipt of reinforcement feedback after every trial. However, we implemented the strategic control training used in Experiment 3, and adapted in Experiment 4, because participants in earlier pilot studies who received only reinforcement and no explicit instruction in strategy shifting showed little to no change in SGT across motivational contexts. Although reinforcement may contribute somewhat to performance shifts, it seems that explicit strategic control is necessary to achieve shifts in performance across multiple-SGT settings.

The results of the time-course analysis of RR changes over the course of each motivational context block are consistent with the hypothesis that explicit strategy drives performance shifts. If reinforcement learning significantly contributed to performance shifts, we would expect the reinforcement group to demonstrate gradual shift toward the targeted stop-signal RR. More interesting, participants in the reinforcement group consistently responded to

Table 4
Stop-Signal RR Changes Within Go Bias and Stop Bias Blocks in Experiment 4

	Go bias			Stop bias		
Factor	Early	Middle	Late	Early	Middle	Late
Reinforcement	.27 (.20)	.26 (.23)	.22 (.20)	.56 (.20)	.63 (.22)	.71 (.23)
Strategy-only	.25 (.16)	.42 (.20)	.44 (.22)	.73 (.14)	.79 (.16)	.82 (.13)

Note. Data are *M* (*SD*). Stop-signal response rates were averaged over the first 10 stop trials (early), the second 10 stop trials (middle), and the last 10 stop trials (late) for each motivation block. RR = response rate.

the stop-signal at the same rate across the go bias block, although participants in the strategy-only group steadily decreased in response rate over the course of the block. This suggests that participants in both groups initially adopted the appropriate strategy, but without the benefit of performance feedback, participants in the strategy-only group increasingly strayed away from the target over the course of the block. Both groups demonstrated a gradual decrease in RR over the course of the stop bias block. Again, the shift in RR for the strategy-only group can be explained by lack of constraints on performance. On the other hand, the gradual change observed in the reinforcement group likely reflects the gradual shift toward a strategy favoring stop accuracy, which is at odds with the prepotent response to the go stimulus, because the go response is reinforced more frequently than performance on stop trials.

Our interpretations of performance shift within each motivation block are consistent with the observation of main effects of feedback group on overall SGT, stop-signal RR, and go-RT slowdown across motivation blocks. Greater SGT values, lower stop-signal RR, and greater go-RT slowdown observed in the strategy-only group suggest greater inherent bias favoring accuracy over speed in our population. The observed differences in overall strategy between groups may have implications for the design of future experiments. Without the benefit of feedback following each trial, participants in the strategy-only group ended up being slower overall, which increases likelihood of stopping correctly on stop trials. This demonstrates the importance of using some kind of performance feedback on the stop-signal task to adequately constrain participants' behavior, which can influence performance bias and our estimation of SSRT.

Another implication for future studies is that it is important to control the strategy setting adopted by participants (e.g., the relative importance of stop- and go-trial performance as measured by SGT). Because we used fixed rewards and penalties (reinforcement group) or no feedback at all (strategy-only group) in Experiment 4, rather than the adaptive titration procedure used in Experiment 3, we observed much variability between participants in the amount of strategic shift between the two motivational contexts. It is important to measure this variability and account for it when both comparing SSRT values across participants (or patient groups) and measuring strategy-induced SSRT changes within-participants. Another, approach is to experimentally control stop–go trade-off using adaptive payoffs, as in Experiment 3, or other manipulations that impact stopping accuracy and go-RT slowdown. This approach is recommended, if it can be employed, because it ensures that all participants adopt similar strategy settings and obviates the need to include SGT changes as a covariate (though the SGT-SSRT correlation can be used to check whether participants' strategy settings were equated).

F5 More interesting, Figure 5 illustrates that those individuals with the largest strategic shift scores actually demonstrated higher SSRT scores on the stop bias block than the go bias block. This was inconsistent with our hypotheses as well as with the results from Experiment 3 that showed that SSRT was higher for the go bias block relative to the stop bias block. As the shift scores reflect the ratio of the change in go-RT slowdown to the change in stopping accuracy, higher strategic shift scores indicate that participants slowed down considerably, but did not substantially increase inhibition accuracy. This pattern may suggest a nonlinear trade-off between speed and accuracy, or it could be the result of poor decision making on a limited number of trials for participants who may have found the strategy instructions confusing or difficult, though no measures of perceived difficulty were obtained. Either way, the pattern of performance associated with higher strategic shift scores brings to light another way that strategic bias can influence SSRT estimates. Because the stop-signal onset time relative to the go-signal onset time was fixed across motivational contexts, a dramatic slowdown in go-RT on the stop bias block resulted in a lengthened SSRT. Certain participants may show this type of pattern in the absence of any explicit strategy instructions, and thus our findings have implications for the interpretation of SSRT in more commonly used variants of the stop-signal task.

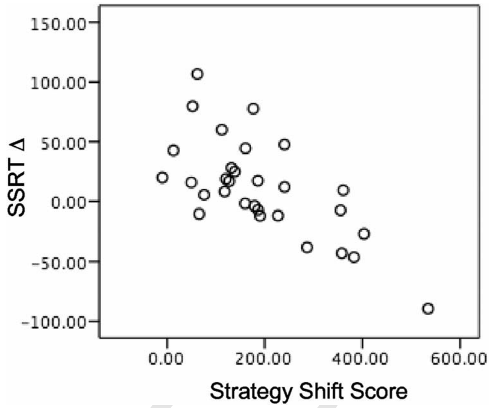

Figure 5. Change in SSRT across motivational contexts is inversely related to shift in strategic trade-off in Experiment 4. On the *y*-axis, positive SSRT change scores reflect that SSRT speed block > SSRT accuracy block, whereas negative change scores reflect SSRT speed < SSRT accuracy. Strategic shift is a measure of the slowdown in go-RT divided by the increase in stop accuracy when comparing performance on the stop bias block relative to the go bias block. Thus, strategic shift reflects the overall change in SGT between blocks, as well as the efficiency of trading off between speed and accuracy: Higher strategic shift scores were observed for participants who slowed considerably when shifting strategies but did not improve stopping accuracy. As a result, these participants demonstrated a negative change in SSRT, reflecting a larger SSRT for the stop bias block relative to the go bias block, whereas participants with lower strategic shift scores demonstrated more positive change in SSRT, reflecting a larger SSRT for the go bias block relative to the stop bias block. SSRT = stop-signal reaction time; SGT = stop–go trade-off.

General Discussion

The stop-signal task and associated SSRT measure (Logan et al., 1984) provide a simple and reliable probe of response inhibition ability. In this paper, we address the central role of motivational influences on stop-signal performance and measurements of inhibitory ability. Although strategic bias has been acknowledged as a potential nuisance factor since the stop-signal paradigm's inception (Logan & Cowan, 1984), the overwhelming majority of research using the stop-signal task does not consider motivational influences when interpreting the results of SSRT differences across individuals or groups, in part because there has been no empirical demonstration of motivational effects on SSRT.

In the present study, however, we demonstrate that performance on the stop-signal task—and in particular SSRT, which is widely assumed to be immune to strategic influences—is indeed susceptible to changes in motivational context. Experiments 1 and 2 showed that measured stop–go trade-off correlates with SSRT in commonly used adaptive stop-signal task variants. Experiments 3 and 4 showed that participants' stop–go trade-off was influenced by experimentally manipulating the payoffs associated with fast responses or with accurate inhibition. As a result, SSRT estimates of stop-latency systematically varied, suggesting that the SSRT measure, as typically implemented, is not insensitive to motivational and strategic influence. The change in SSRT observed in Experiment 3 is comparable to the difference in SSRT reported between individuals with ADHD and healthy controls in several studies (Aron, Dowson, Sahakian, & Robbins, 2003; Epstein, Johnson, Varia, & Conners, 2001; Manassis, Tannock, & Barbosa, 2000). If motivational bias or strategy shifts can affect SSRT estimates within individuals, then inherent motivational biases may be driving much of the difference observed in SSRT estimates between individuals. Further, Experiment 3 showed that experimentally manipulating stop–go trade-off provides a way to measure it and de-confound it from inhibitory ability.

Motivational context may produce systematic variations in SSRT for two major reasons: (a) Inhibitory ability is independent of motivational context, but our measures of inhibitory ability (e.g., SSRT) depend on processes influenced by motivational context. Violation of the major assumptions of the SSRT measurement model (Logan & Cowan, 1984), in particular the notion that the stop process has a fixed duration, may lead to biased estimates of stop latency (Band, van der Molen, & Logan, 2003; De Jong et al., 1990). One way to address this issue would be to obtain summary measures of inhibitory ability, which account for variations in performance across multiple contexts, such as the AUC measure used in Experiment 3. Alternately, (b) inhibitory ability may not be a fixed-trait ability, but may be mutable, so that stopping correctly becomes more or less difficult for an individual depending on the primary motivational focus toward speed or accuracy. A plausible mechanism might be that putative go and stop pathways in the brain (Frank, Seeberger, & O'Reilly, 2004) can be differentially potentiated by mental preparation, and that shifts in motivation to go fast or stop correctly may change the relative potentiation of the pathways. Future studies designed to test these alternative hypotheses would make significant contributions to our understanding of inhibitory control and how it is modulated by motivational context.

A Theoretical Account of Motivational Context Effects in Stop-Signal Performance

Below, we describe a theoretical information-processing account that captures the experimental effects on SSRT, go-RT, and SGT that we have observed in these experiments. Motivational influences on stop-signal performance can be understood in terms of a simple quasi-Bayesian decision framework. In this framework, one fundamental decision needs to be made in the stop-signal task: Which kind of trial is this trial—stop or go? We denote the posterior probability that this trial is a go trial as P(R), the probability that a motor response is required. Like the race model, we posit two different kinds of information-accumulation processes: P(go) is the evidence, expressed as a likelihood, that the go signal is present, and P(stop) is the likelihood that the stop signal is present. These two signals develop over time, influenced by both the sensory evidence (presentation of go and stop signals) and prior probabilities. Go stimuli update P(go), and stop signals update P(stop), and these evolve over time (*t*). The stronger the evidence for the go signal, the more likely that a response is required, P(R) increases; but it must be weighted against the subjective probability that a stop trial will occur at some later time, which reduces P(R). As in information accumulator models (Ratcliff, 1985, 1988; Ratcliff & Smith, 2004; Smith & Vickers, 1988; Usher & McClelland, 2001), if P(R) reaches a decision threshold, a response is made. A simple calculation of P(R) across time is as follows:

$$P(R)_t = \int_{i=1}^{t} \mathrm{P}(\mathrm{go})_i - \int_{i=1}^{t} \mathrm{P}(\mathrm{stop})_i - \int_{i=t}^{T} \mathrm{P}(\mathrm{Stop})_i,$$

where go represents the probability a go signal has occurred, stop represents the probability a stop signal has occurred, hedging represents the predicted probability of a future stop signal before the end of the trial (*T*). Thus, even if the go signal has been presented and strong evidence has accumulated that the trial will be a go trial, the evidence that the stop signal is still likely to occur may be strong as well, and P(R) may not reach threshold. As time elapses and the subjective likely interval for stop signal presentation passes, hedging decreases, increasing P(R), and enabling a response. Hedging in turn depends on the prior probability of a stop signal occurring on the trial and the likely temporal intervals. With a choice-RT go-task (e.g., left or right arrow), this model would be supplemented with a second, temporally overlapping function that determines which of the possible motor responses is appropriate.

Normatively, when the go stimulus is presented and a trial begins, then P(go) increases. However, it cannot achieve its maximum value (and should usually remain initially below threshold) because the hedging term is nonzero, equal in value to the predicted stop-signal base rate (e.g., 0.25 for 25% stop trials). If a stop-signal is presented, then evidence accumulates that decreases P(R) further and maintains it at a low level (e.g., if P(go) = 1 and P(stop) = 1, P(R) = 0), and no response is made. If a stop signal is not presented, then P(go) accumulates normally, but the hedging term remains high until the likely time interval for the stop-signal presentation (based on previous experience) has elapsed. When enough time has elapsed, the hedging term drops and P(R) is increased to threshold, driven by P(go).

A distinctive and important feature of this model is that the posterior probability of the trial being a stop versus go trial, P(R), is not only influenced by the stimuli presented during the trial (as in the race model), but also by the subjective probability that a stop signal will occur at some time in the future before the end of the trial, which is learned across trials and depends on both the estimated stopping base rate Pr(S) and a function of learned subjective probabilities over time. Across trials, the stop-signal base rate Pr(S) may be updated.

This feature is likely to be important for explaining both why RTs are substantially slower in mixed blocks of stop and go trials than they are in pure blocks of go-only trials (as our experiments demonstrate), and several other effects as well—including why some participants continually slow down on adaptive "tracking" tasks, and why SSRT grows shorter with increasing accuracy bias and longer with speed bias. For example, in go-only blocks, the hedging term is essentially zero throughout the trial (because Pr(S) is close to zero), enabling rapid RTs. With mixed stop and go trials, some time must elapse before the likelihood of a future stop signal (the hedging term) decreases enough that P(R) exceeds the response threshold; thus, RT is lengthened. On tracking tasks, changes in the stop-signal delay make the subjective distribution of likely stop-signal times less precise. If the distribution is sufficiently broad and the response threshold sufficiently high, then the participant will wait for evidence that the stop-signal is coming, and thus slow down. However, because the stop-signal delay increases as participants slow down, then the internal model of likely stop-signal times also becomes more diffuse, and the participant must wait yet longer to be sure that the stop-signal will not occur and P(R) can reach threshold.

Motivational context may affect one or more of several parameters in this model that result in both performance trade-offs and differences in SSRT. Motivational factors may influence (a) the response threshold for P(R) to drive action; (b) the rate at which stop and go information are accumulated on stimulus presentation; or (c) the subjective probability that a stop trial will occur in the future (Pr(S)). All of these might be expected to influence RT means and variabilities, RR, SGT, and SSRT in the directions found in our experiments. For example, actual probabilities may be subjectively weighted for several reasons, including (a) individual differences in the interpretation of task goals and what it means to be successful on the stop-signal task, explaining differences in performance and in SSRT in participants from Experiments 1 and 2; (b) variations in explicit reward and penalties for performance on stop (and go) trials, as in Experiment 3, as well as sensitivity to those rewards and penalties (Yechiam et al., 2006); and (c) informational value imposed by instructions to adopt a specific strategic trade-off, as in Experiment 4. Figure 6 outlines the influences of both stop-signal frequencies and reward payoffs on subjective stop-signal probabilities, and their resulting influences on performance outcome measures.

F6

We present this model as a preliminary theoretical explanation of motivational influences on response inhibition in the stop-signal task. Though formal testing of the model is beyond the scope of the present paper, we believe that this framework is an appropriate starting point for explaining the observed phenomena. The proposed model is consistent both with signal detection theory and with error-driven learning models, and many researchers have applied similar decision models to explain response times, decisions, and speed-accuracy trade-offs in cognitive tasks (Busemeyer & Townsend, 1993; Ratcliff, 1985, 1988; Ratcliff & Smith, 2004; Smith & Vickers, 1988; Usher & McClelland, 2001). Across the four experiments presented here, we demonstrate various sources of motivational and strategic bias that influence stop-signal performance and measures of inhibitory ability. Ultimately, each of the motivational effects we observed, as well as some effects commonly observed in other research (e.g. frequency and history effects), can be effectively explained by the proposed conceptual framework.

It is important to address the influence of motivation on inhibitory control, since these inhibitory control processes are thought to play principal roles in other basic cognitive processes such as attention (Neill, 2007; Rafal & Henik, 1994) and memory (Anderson & Green, 2001; Hasher & Zacks, 1988), as well as in complex regulatory behaviors related to emotion processing and social interaction (Kunda & Spencer, 2003).

Many researchers have questioned the reliability and validity of response inhibition measures as indexes of inhibitory control, largely because zero-order correlations between different measures of inhibitory control tend to be very low (Friedman & Miyake, 2004; Kramer et al.. 1994; Shilling, Chetwynd, & Rabbitt, 2002). However, small correlations do not necessarily imply that the measures are unrelated. Friedman and Miyake suggested that low zero-order correlation could be explained by the impurity of the inhibition measures. We believe that the development of task variants that reflect more transparent measures of cognitive abilities, uncontaminated by strategic and motivational effects, may provide a clearer picture of the structure of cognitive control processes. In this paper, we have taken steps towards developing the stop-signal paradigm in that direction.

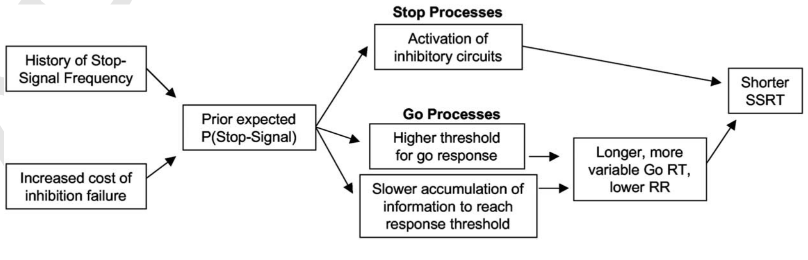

Figure 6. Theoretical model for explaining motivational influences on response inhibition. Motivational context influences the expected probability the stop-signal will occur, either directly by changing the frequency of stop trials, or indirectly, by changing the subjective weighting of probabilities due to internally generated value of stopping or externally imposed rewards and penalties for performance. Increase in P(stop-signal) may alter inhibitory control either directly by affecting stop-processes and engaging inhibitory circuits, or indirectly, by changing the decision threshold (or latency to reach the decision threshold) for go responses. In either case, we would expect to observe lengthened go responses and shortened SSRT values. SSRT = stop-signal reaction time; RR = response rate; RT = reaction times.

Here, we argued that slight differences in task instructions, stimuli, and experimental parameters can substantially alter measures of inhibitory control in the stop-signal task. However, we also showed that using explicit payoffs and adaptive procedures that target the same performance strategies in all participants provide a way to remove these sources of bias. The stop-signal is an excellent model paradigm for this effort because of its relative simplicity, and because of its face validity as a measure of a particular type of response inhibition. Similar efforts could be made with other RT-based tasks. We hope that the present study will encourage researchers to take these considerations into account when designing future studies, ultimately contributing to a unified goal of developing an improved taxonomy of inhibitory control.

References

Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature, 410,* 366–369.

Aron, A. R., Dowson, J. H., Sahakian, B. J., & Robbins, T. W. (2003). Methylphendiate improves response inhibition in adults with attention-deficit/hyperactivity disorder. *Biological Psychiatry, 54,* 1465–1468.

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

Badcock, J. C., Michie, P. T., Johnson, L., & Combrinck, J. (2002). Acts of control in schizophrenia: Dissociating the components of inhibition. *Psychological Medicine, 32,* 287–297.

Band, G. P. H., van der Molen, M. W., & Logan, G. D. (2003). Horse-race model simulations of the stop-signal procedure. *Acta Psychologica (Amsterdam), 112,* 105–142.

Band, G. P. H., van der Molen, M. W., Overtoom, C. C., & Verbaten, M. N. (2000). The ability to activate and inhibit speeded responses: Separate developmental trends. *Journal of Experimental Child Psychology, 75,* 263–290.

Barkeley, R. A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. *Psychological Bulletin, 121,* 65–94.

Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making, and the orbitofrontal cortex. *Cerebral Cortex, 10,* 295–307.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459.

De Jong, R., Coles, M. G. H., Logan, G. D., & 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,* 164–182.

Eisenberg, N., Fabes, R. A., Karbon, M., Murphy, B. C., Wosinski, M., Polazzi, L., et al. (1996). The relations of children's dispositional prosocial behavior to emotionality, regulation, and social functioning. *Child Development, 67,* 974–992.

Epstein, J. N., Johnson, D. E., Varia, I. M., & Conners, C. K. (2001). Neuropsychological assessment of response inhibition in adults with ADHD. *Journal of Clinical and Experimental Neuropsychology, 23,* 362–371.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science, 306*(5703), 1940–1943.

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

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and new view. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 193–225). New York: Academic.

Kindlon, D., Mezzacappa, E., & Earls, F. (1995). Psychometric properties of impulsivity measures: Temporal stability, validity and factor structure. *The Journal of Child and Adolescent Psychology, 36,* 645–661.

Kochanska, G., Murray, K. T., & Harlan, E. T. (2000). Effortful control in early childhood: Continuity and change, antecedents, and implications for social development. *Developmental Psychology, 36,* 220–232.

Kramer, A. F., Humphrey, D. G., Larish, J. F., Logan, G. D., & Strayer, D. L. (1994). Aging and inhibition: Beyond a unitary view of inhibitory processing in attention. *Psychology and Aging, 9,* 491–512.

Kunda, Z., & Spencer, S. J. (2003). When do stereotypes come to mind and when do they color judgment? A goal-based theoretical framework for stereotype activation and application. *Psychological Bulletin, 129,* 522–544.

Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Science, 9*(2), 75–82.

Logan, G. D. (1994). On the ability to inhibit thought and action. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language*: San Diego, CA: Academic. AQ: 4

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91,* 295–327.

Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit simple and choice reaction time responses: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 276–291.

Logan, G. D., Schachar, R. J., & Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science, 8,* 60–64.

Luman, M., Oosterlaan, J., & Sergeant, J. A. (2005). The impact of reinforcement contingencies on AD/HD: A review and theoretical appraisal. *Clinical Psychology Review, 25*(2), 183–213.

Manassis, K., Tannock, R., & Barbosa, J. (2000). Dichotic listening and response inhibition in children with comorbid anxiety disorders and ADHD. *Journal of the American Academy of Child and Adolescent Psychiatry, 39,* 1152–1159.

Neill, W. T. (2007). Mechanisms of transfer-inappropriate processing. In D. S. Gorfein & C. M. MacLeod (Eds.), *Inhibition in cognition* (pp. 63–78). Washington, DC: American Psychological Association.

Nigg, J. T. (2001). Is ADHD a disinhibitory disorder? *Psychological Bulletin, 127,* 571–598.

Ochsner, K. N., Bunge, S. A., Gross, J. J., & Gabrieli, J. D. E. (2002). Rethinking feelings: An fMRI study of the cognitive regulation of emotion. *Journal of Cognitive Neuroscience, 14,* 1215–1229.

Phan, K. L., Fitzgerald, D. A., Nathan, P. J., Moore, G. J., Uhde, T. W., & Tancer, M. E. (2005). Neural substrates for voluntary suppression of negative affect: A functional magnetic resonance imaging study. *Biological Psychiatry, 57,* 210–219.

Rafal, R., & Henik, A. (1994). The neurology of inhibition: Integrating controlled and automatic processes. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention memory and language* (pp. 1–51). San Diego, CA: Academic.

Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review, 92,* 212–225.

Ratcliff, R. (1988). Continuous versus discrete information processing modeling accumulation of partial information. *Psychological Review, 95,* 238–255.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111,* 333–367.

Schachar, R., Mota, V. L., Logan, G. D., Tannock, R., & Klim, P. (2000).

Confirmation of an inhibitory control deficit in attention-deficit/hyperactivity disorder. *Journal of Abnormal Child Psychology, 28,* 227–235.

Schachar, R., Tannock, R., Marriott, M., & Logan, G. (1995). Deficient inhibitory control in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology, 23,* 411–437.

Shilling, V. M., Chetwynd, A., & Rabbitt, P. M. A. (2002). Individual inconsistency across measures of inhibition in older adults. *Neuropsychologia, 40,* 605–619.

Slater-Hammel, A. T. (1960). Reliability, accuracy, and refractoriness of a transit reaction. *The Research Quarterly, 31,* 217–228.

Slusarek, M., Velling, S., Bunk, D., & Eggers, C. (2001). Motivational effects on inhibitory control in children with ADHD. *Journal of the American Academy of Child & Adolescent Psychiatry, 40,* 355–63.

Smith, P. L., & Vickers, D. (1988). The accumulator model of two-choice discrimination. *Journal of Mathematical Psychology, 32,* 135–168.

Swylan, R. P. (2004). The control of deliberate waiting strategies in a stop-signal task. *Brazilian Journal of Medical and Biological Research, 37,* 853–862.

Tannock, R., Schachar, R. J., Carr, R. P., Chajczyk, D., & Logan, G. D. (1989). Effects of methylphenidate on inhibitory control in hyperactive children. *Journal of Abnormal Child Psychology, 17,* 473–491.

Thorell, L. B., Bohlin, G., & Rydell, A. M. (2004). Two types of inhibitory control: Predictive relations to social functioning. *International Journal of Behavioral Development, 28,* 193–203.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108,* 550–592.

van den Wildenberg, W. P. M., van der Molen, M. W., & Logan, G. D. (2002). Reduced response readiness delays stop signal inhibition. *Acta Psychologica (Amsterdam), 111,* 155–169.

van der Schoot, M., Licht, R., Horsley, T. M., & Sergeant, J. A. (2005). Effects of stop signal modality, stop signal intensity and tracking method on inhibitory performance as determined by use of the stop signal paradigm. *Scandinavian Journal of Psychology, 46,* 331–341.

Walcott, C. M., & Landau, S. (2004). The relation between disinhibition and emotion regulation in boys with attention deficit hyperactivity disorder. *Journal of Clinical Child and Adolescent Psychology, 33,* 772–782.

Yechiam, E., Goodnight, J., Bates, J. E., Busemeyer, J. R., Dodge, K. A., Pettit, G. S., et al. (2006). A formal cognitive model of the go/no-go discrimination task: Evaluation and implications. *Psychological Assessment, 18,* 239–249.

Yonelinas, A. P. (1994). Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory & Cognition, 20,* 1341–1354.

Received May 19, 2008
Revision received March 25, 2009
Accepted March 31, 2009 ■

AUTHOR QUERIES

AUTHOR PLEASE ANSWER ALL QUERIES

AQ1: Author: Please provide values to the same decimal points. APA recommends two decimal places. Please be consistent throughout.

AQ2: Author: Please define this acronym at its first usage here.

AQ3: Author: Please define this acronym at its first usage here.

AQ4: Author: Logan (1994): Please provide the page range for this chapter.

AQ5: Author: If figure can be edited subject should be changed to participant in all cases.

AQ6: Author: If figure can be edited tradeoff should be hyphenated to trade-off; *n* should be italed.