Response biases in simple decision making: Faster decision making, faster response execution, or both?

Jeffrey J. Starns[1,2] · Qiuli Ma[1]

Abstract When people are biased to use one response more often than an alternative response in a decision task, they also make the preferred response more quickly. Sequential sampling models can accommodate this difference in response time (RT) by changing the relative amount of evidence that must accumulate to decide in favor of one versus the other response, but nondecision processes might also play a role, such as the amount of time between selecting and executing a response. We investigated the influence of decision and nondecision processes in two experiments. In Experiments 1a and 1b, arrows appeared on the screen, and participants were asked to move a joystick in the direction of the arrow or make a keypress as quickly as possible. Results showed that motor execution times were faster for expected directions than unexpected directions. In Experiments 2a and 2b, participants decided whether a high or low number of asterisks was displayed on the screen. Decision times were faster for the stimulus class that was more likely to appear, and this effect was larger when participants could anticipate both the likely stimulus class and the motor response needed to identify it than when they knew the likely stimulus class but the associated motor response changed probabilistically from trial to trial. These results show that both decision and nondecision factors contribute to bias effects on RT.

Imagine that you are at a social gathering, and you see someone quickly approaching for a conversation. In seconds, you must decide whether or not you have met this person before. Your decision will obviously be informed by the characteristics of the person, but you might also consider contextual information, such as whether the guests are mostly people that you have met before (e.g., it is your high-school reunion) or mostly people that you have not met before (e.g., it is your partner's high-school reunion). The influence of this contextual information is often referred to as "bias" because it affects decisions independently of the characteristics of a target stimulus. For example, you might decide that an approaching person who is moderately familiar is someone that you met before if you are at your own reunion, but make the opposite decision if you are at your partner's reunion.

Psychologists have extensively investigated the cognitive processes involved in making simple decisions such as the one described above. The most successful models have been sequential sampling models (e.g., Ratcliff & Smith, 2004), which assume that decisions are made by accumulating information from a stimulus over time. The information available in any small time window (a "sample") is assumed to be very noisy due to external variation in the stimulus and/or intrinsic variation in the processing system (e.g., neural noise). To deal with this uncertainty, the models collect new samples until the amount of evidence in favor of one decision meets a criterion. Decision biases can be introduced in the models by changing

✉ Jeffrey J. Starns
jstarns@psych.umass.edu

1 University of Massachusetts Amherst, Amherst, MA, USA

2 Department of Psychological and Brain Sciences, University of Massachusetts, 135 Hicks Way, Amherst, MA 01003-9271, USA

the amount of information needed to make one response versus the alternative response.[1] Other crucial processing components of these models are the quality of evidence, the caution of the decision maker, and the time required for "nondecision" processes such as hitting a response key after the decision is made.

Sequential sampling models have been very successful in accommodating the effects of response bias on both response proportions and response time (RT) distributions (e.g., Criss, 2010; Ratcliff & McKoon, 2008; Starns, Ratcliff, & McKoon, 2012; Starns, Ratcliff, & White, 2012; Voss, Rothermund, & Voss, 2004). When participants are biased to use a response more frequently than an alternative response, they are also quicker to make the preferred response for both correct and error responses. This speedup is a natural product of the way that bias is represented in the models: If less information must accumulate to trigger a biased response, then the accumulation process leading to a biased response should be faster.

Although bias effects on RT can be successfully modeled by changing only bias parameters, nondecision time might also contribute to faster RTs for biased responses (Voss, Voss, & Klauer, 2010). That is, people might be faster when executing expected movements than unexpected movements, so the motor execution time following a decision could be faster for the response that is consistent with the prevailing bias. Using parameter recovery simulations, Voss et al. (2010) showed that introducing differences in nondecision time across responses can lead to severe distortions in parameter estimates in the diffusion model, one of the most popular sequential sampling models. These problems with parameter estimation mean that researchers are at risk of making incorrect conclusions about the underlying cognitive processes involved in biased decision making. Thus, empirically determining whether or not bias affects nondecision processes is an important goal for the community of researchers who use sequential sampling models.

We explored the extent to which bias effects on RT are produced by changes in decision time versus nondecision time by attempting to selectively influence these components. Our first experiment was designed to explore the motor requirements of biased responding, and we were especially interested in whether RTs were faster for expected than for unexpected movements. The task itself was quite simple: An arrow appeared on the screen, indicating the direction to move the joystick (or which key to press in Experiment 1b), and participants were asked to complete the movement as quickly as possible when they saw the arrow. The arrows were offset

from the center of the screen in the same direction that they pointed, so it was very easy to discriminate the different arrow stimuli. We did this to minimize the need to accumulate evidence for an uncertain decision outcome, allowing us to focus on the motor execution process.

To explore the effects of anticipating a movement on response time, we compared a range of conditions from blocks with only one possible movement to blocks with four possible movements. If response preparation has an effect, then participants should make movements more slowly when multiple directions are possible. We also explored anticipation effects with a standard bias manipulation. That is, some multiple-direction blocks introduced biases in the frequency with which each direction appeared to determine if motor execution times were faster for the more frequent direction. If bias manipulations affect motor execution times, then we should see a bias effect on RT in the arrow task even though it essentially eliminates the need to accumulate evidence for an uncertain stimulus classification. Finally, we directly compared the joystick response modality with more typical keypress responding for the two-direction conditions.

Although the arrow task should have minimal decision requirements, we cannot completely rule out the possibility that decision processes produced any potential bias effects in this task. To address this limitation, Experiment 2 explored the effect of response bias on a standard decision-making task. We used a numerosity task in which a variable number of asterisks appeared on the screen and participants had to quickly decide if a high or low number was presented (Ratcliff, Thapar, & McKoon, 2001). High and low stimuli were equally likely to appear in neutral blocks, and in biased blocks each trial was preceded by a hint about which stimulus class was likely to appear on the trial.

We explored the relative contributions of decision and non-decision time by comparing different "mapping" conditions within Experiment 2. In the constant-mapping condition, the same movement was used to respond "high" or "low" on every trial. Thus, biased blocks allowed participants to anticipate both the likely stimulus class and the likely response, and both decision and motor preparation factors could contribute to the bias effect on RT. In the variable-mapping condition, trials switched unpredictably between a horizontal or vertical response orientation such that there were two responses for "high" (e.g., up or to the left) and two responses for "low" (e.g., down or to the right). Whether the participant should respond with the horizontal or vertical orientation was cued on each trial immediately after the asterisk stimulus was presented. Thus, participants could anticipate the likely stimulus class for biased blocks but could not anticipate the movement required for the response, preserving the role of decision factors but minimizing the effects of motor preparation. A sizeable effect in the variable-mapping condition would provide evidence for the role of decision processes because nondecision

---

[1] Sequential sampling models incorporate two mechanisms for introducing response biases: changing the amount of information that must accumulate for a response and changing the standards for which response is supported by a given evidence sample (e.g., Starns, Ratcliff, & White, 2012). We focus on the first because it is the type of bias primarily affected by stimulus proportion manipulations.

factors should be minimized. By the same logic, a larger effect in the constant-mapping condition than in the variable-mapping condition would provide evidence for the role of nondecision processes.

Our predictions for the mapping manipulation rest on the assumption that knowing which one of two responses to anticipate speeds motor execution times in the constant-mapping condition, but knowing which two of four responses to anticipate has little or no effect in the variable-mapping condition. Another goal of Experiment 1 was to test this assumption directly by comparing two types of blocks in the arrow task: (1) blocks in which one direction had a 75% chance of appearing and a second direction had a 25% chance, and (2) blocks in which two directions each had a 37.5% chance of appearing and another two directions each had a 12.5% chance of appearing. The first type of block was analogous to the constant-mapping condition (two possible responses), and the second type was analogous to the variable-mapping condition (four possible responses). If the logic underlying our predictions holds, then the RT difference between probable and improbable directions should be larger on the two-direction blocks than on the four-direction blocks.

Experiments 1a and 1b

Method

Participants University of Massachusetts Amherst undergraduates participated to earn extra credit in their psychology courses. We ran 47 participants for Experiment 1a and 22 for Experiment 1b.

Procedure

Each trial began with either one, two, or four thin (1 pixel wide) arrows on the screen. At some point, a thick (4 pixels wide) arrow appeared covering (one of) the thin arrow(s). Each arrow began at the center of the screen and extended left, right, up, or down. The participant was instructed to respond with the direction of the thick arrow as quickly as possible when it appeared. The thick arrow appeared after a delay of 300 ms plus a value sampled from an exponential distribution with a mean of 400 ms. Delays were truncated at 3 s.

In Experiment 1a people responded with a joystick, and there were five block types differing in the number of possible arrow stimuli and whether the presentation frequencies were biased (B) or unbiased (U). For single-arrow (simple-reaction time) blocks, the arrow pointed the same direction on every trial, and both the thin and thick arrows were displayed with white lines on a black screen. For two-arrow unbiased (2U) blocks, the trial began with thin white arrows pointing both left and right, and the thick arrow appeared in one of these

directions with equal probability. For two-arrow biased (2B) blocks, the thick arrow appeared in one direction 75% of the time and in the other only 25% of the time. The more likely direction was left for half of these blocks and right for half. The arrows pointing in the likely and unlikely directions were displayed in green and red, respectively, and participants were informed of this. For four-arrow unbiased (4U) blocks, four thin white arrows appeared at the beginning of the trial pointing left, right, up, and down, and the thick arrow appeared in one of these four directions with equal probability. For four-arrow biased (4B) blocks, one vertical and one horizontal direction each had a 37.5% chance of displaying the thick arrow (e.g., left and up), and the other two each had a 12.5% chance (e.g., right and down). Arrows for the two more-likely directions were displayed in green and the others were red. Participants completed five practice blocks (one for each block type) followed by 14 critical blocks of 64 trials each in a random order (two one-arrow, two 2U, two 4U, four 2B, and four 4B blocks).

Experiment 1b had only two-arrow biased blocks, and participants responded with either the joystick or by pressing "z" or "/" for "left" and "right." Participants completed two practice blocks followed by 12 critical blocks with 64 trials each. Half of the blocks were randomly assigned to each response modality.

Results and discussion

Raw data for all experiments are available on the Open Science Framework (https://osf.io/ws3cj/). We analyzed results using Bayesian t tests with a standard or folded unit-normal prior on effect size for nondirectional and directional tests, respectively. Bayes factors were often very small, so we report results in terms of the posterior probability of either the alternative (alt.) or null hypothesis (whichever was favored) assuming equal prior probabilities.

As in previous research (e.g., Fay, 1936), simple reaction times to a single stimulus were quite fast: a typical participant had a median of 279 ms for the one-arrow condition. RTs slowed when there were more possible responses (351 and 388 ms for two and four arrows unbiased, respectively), which is consistent with the idea that motor execution times are slower when it is more difficult to anticipate the required response. However, a typical RT median was below 400 ms even in the four-arrow blocks, and the low RTs across all conditions suggest that the arrow task is very easy and has minimal decision requirements, as intended. Further supporting this claim, errors occurred on less than 2% of trials for all blocks types.

Figure 1b shows the effect of response bias on RT medians for correct responses. The two-arrow-biased blocks are labeled as constant mapping and the four-arrow-biased blocks are labeled as variable-mapping because these conditions have
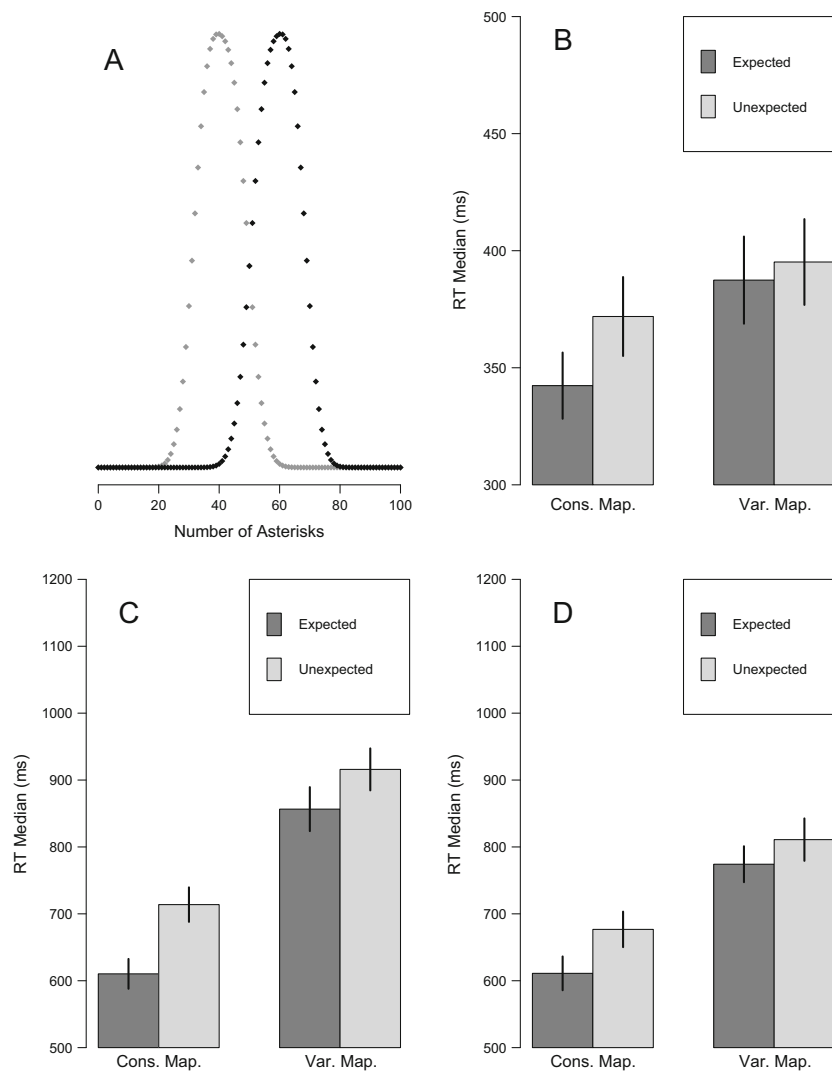
Fig. 1 **a** Distributions for the number of asterisks displayed for "low" (*gray*) and "high" (*black*) stimuli in Experiment 2. **b** Average correct RT medians for the expected and unexpected stimulus in the constant-mapping (Cons. Map.) and variable-mapping (Var. Map.) conditions of Experiment 1a. **c** Same for Experiment 2a. **d** Same for Experiment 2b. *Error bars* on all plots are 95% confidence intervals

similar response requirements to the corresponding conditions in Experiment 2. For Experiment 1a, participants responded more quickly when the arrow pointed in the expected direction than in the unexpected direction, $t(46) = 7.29$, p(alt.) > .99. The bias effect was larger with constant mapping (30 ms) than with variable mapping (8 ms), $t(46) = 5.85$, p(alt.) > .99. There was clear evidence for an effect even with variable mapping, $t(46) = 3.36$, p(alt.) = .96. For Experiment 1b (which only used constant mapping), people were again faster for expected than unexpected responses both when they responded with the joystick (297 vs. 324 ms), $t(21) = 5.28$, p(alt.) > .99, and when they responded with the keyboard (286 vs. 308 ms), $t(21) = 5.81$, p(alt.) > .99. RTs were very similar across the two response modalities, and a test comparing the size of the bias effect across modalities found suggestive evidence for the null hypothesis, $t(21) = 1.24$, p(null) = .70.

The most important implication of these results is that people make expected responses more quickly than unexpected responses even in a task with minimal decision requirements. Experiment 1b showed that this pattern was not unique to the joystick response modality. In subsequent experiments we use the joystick to facilitate four-option responding in the variable-mapping conditions, but Experiment 1b suggests that the conclusions can be generalized to keypress responding.

Experiments 2a and 2b

Method

Participants University of Massachusetts Amherst undergraduates participated to earn extra credit in their psychology

courses. For Experiment 2a, we ran 62 participants with constant mapping and 57 with variable mapping. Participants with accuracy below .55 in the unbiased condition were excluded from analyses, resulting in the loss of one constant-mapping participant and two variable-mapping participants. For Experiment 2b, we ran 33 participants and excluded one from analyses for near-chance accuracy.

Procedure We used a numerosity task in which a variable number of asterisks appeared on the screen and participants were instructed to quickly decide whether the stimulus came from a high set or a low set. To construct the stimulus display for each trial, each position on a 10 × 10 grid was randomly assigned to display either an asterisk or a space with a probability parameter p, with $p > .5$ for high stimuli and $p < .5$ for low. To introduce additional variability beyond the binomial distributions produced by this process, the probability of displaying an asterisk in each position switched between .36 or .45 for low stimuli and .55 and .64 for high stimuli, producing the distributions seen in Fig. 1a. As in previous applications of this task (e.g., Ratcliff et al., 2001), the distributions overlapped such that some errors were unavoidable.

For Experiment 2a, each participant completed two practice blocks of 70 trials each and six critical blocks of 96 trials each. Half of the blocks were "hint" blocks and half were "no hint" blocks. On hint blocks, trials began with a message reading "Probably H" or "Probably L" to indicate that the display for that trial had a 75% chance of coming from the high or the low set, respectively. Trials with the two cues were randomly intermixed and occurred with equal probability. Trials in no hint blocks began with a message reading "No Hint," and participants were informed that there was an equal chance of high and low. Participants pressed the right "bumper" on the game controller to initiate the stimulus display. The asterisk display remained on the screen for 250 ms and was immediately followed by a screen showing the response mapping (the short duration of the stimulus display ensured that participants would not have to delay responding to wait for the response mapping). For constant-mapping participants, the response display showed an "H" and "L" side-by-side in the center of the screen. When they were ready to respond, they moved the joystick in the direction of the "H" or "L" to respond "high" or "low." The mapping of responses to directions was selected randomly for each participant but was constant across the session within a participant. Variable-mapping participants saw the horizontal response mapping described above on 50% on the trials and a vertical mapping on the other 50%. The vertical mapping showed an "H" and "L" one on top of the other, and participants had to move the joystick up or down to select a response. Trials with a vertical and horizontal response mapping were randomly intermixed.

Experiment 2b used only the hint condition and alternated between blocks of constant and variable mapping. The mapping condition was signaled to participants before each block with the messages "left-right only" (constant mapping) or "direction will switch" (variable mapping). Participants completed one 45-item practice block with variable mapping followed by twelve 96-item critical blocks. Sessions ended either when all blocks were complete or after 42 min. The critical blocks were evenly divided between constant and variable mapping with a new random order for each participant.

Results

Participants used the correct response orientation on variable mapping trials an average of 97% of the time for Experiment 2a and 95% for Experiment 2b. We included trials with orientation errors in all analyses and coded them as correct if participants indicated the correct stimulus class even though they used the wrong orientation. The reported results do not change in any meaningful way if these trials are excluded. Participants responded correctly on 80% of the trials on unbiased blocks in Experiment 2a. With the distributions that were used (see Fig. 1a) the maximum possible accuracy is 92%. The fact that participants were well below this value means that many errors were driven by imprecise perception of the number of asterisks displayed. On the hint blocks, accuracy was 87% for the expected stimulus category and 66% for the unexpected category in Experiment 2a, with corresponding values of 86% and 67% for Experiment 2b. This pattern demonstrates that participants used the hint information to inform their choices. Accuracy levels were very similar with constant and variable mapping across all trial types in both Experiment 2a (.80 and .79 for no hint; .87 and .87 for expected stimuli; .67 and .64 for unexpected stimuli) and 2b (.85 and .86 for expected stimuli; .67 and .67 for unexpected stimuli).

Figure 1c shows correct RT medians for Experiment 2a. RTs were longer with variable than with constant mapping, $t(114) = 11.87$, $p(\text{alt.}) > .99$. As in the arrow task, RTs were shorter for the expected direction than the unexpected direction for both the constant-mapping condition, $t(60) = 10.95$, $p(\text{alt.}) > .99$, and the variable-mapping condition, $t(54) = 8.36$, $p(\text{alt.}) > .99$. The bias effect was larger with constant mapping (103 ms) than with variable mapping (59 ms), $t(114) = 3.67$, $p(\text{alt.}) > .99$. Figure 1c shows data for Experiment 2b. Bias effects were smaller overall than in Experiment 2a, but the difference between experiments is unlikely to be systematic given that they used identical bias manipulations. Experiment 2b replicated Experiment 2a by showing faster RTs for expected responses in both the constant-mapping condition, $t(31) = 8.39$, $p(\text{alt.}) > .99$, and the variable-mapping condition, $t(31) = 8.80$, $p(\text{alt.}) > .99$, with a larger bias effect for constant mapping (66 ms) than for variable mapping (37 ms), $t(31) = 4.20$, $p(\text{alt.}) > .99$.

If the mapping manipulation affects nondecision times, then it should affect not just the median of the RT distribution

but also the "leading edge." We evaluated this by replicating the analyses above on the .1 quantile of the RT distribution (the time at which 10% of responses have already been made). All of the patterns reported for median RTs were also observed for the .1 quantiles. Most critically, the leading edge was lower for expected than for unexpected responses in both the constant-mapping condition [81 ms bias effect for Experiment 2a, $t(60) = 10.65$, $p(\text{alt.}) > .99$; 59 ms bias effect for Experiment 2b, $t(31) = 6.26$, $p(\text{alt.}) > .99$] and the variable-mapping condition [59 ms bias effect for Experiment 2a, $t(54) = 9.87$, $p(\text{alt.}) > .99$; 33 ms bias effect for Experiment 2b, $t(31) = 4.79$, $p(\text{alt.}) > .99$]. The bias effect on leading edge was larger with constant than with variable mapping, with suggestive support for a mapping effect in Experiment 2a, $t(114) = 2.21$, $p(\text{alt.}) = .65$, and strong support in Experiment 2b, $t(31) = 3.12$, $p(\text{alt.}) = .96$.

We claim that the bias effect on RT is larger with constant mapping because variable mapping attenuates nondecision time differences for expected versus unexpected movements. An alternative explanation is that our variable-mapping condition simply disrupted the bias manipulation. We tested this possibility by evaluating whether the bias effect on accuracy was also larger with constant than with variable mapping. If variable mapping simply disrupts the bias effect, then it should attenuate the effect in accuracy as well as RT. In contrast, nondecision time uniquely affects RT. The accuracy bias effect in Experiment 2a was very similar in size for the constant-mapping (.21) and variable-mapping (.23) conditions, and a directional t test supported the null, $t(114) = -0.79$, $p(\text{null}) = .90$. The same held for Experiment 2b, which had a .18 bias effect in accuracy for both mapping conditions, $t(31) = -0.13$, $p(\text{null}) = .86$. Thus, our results support the claim that nondecision time plays a role in bias effects by showing that the mapping effect is unique to RT, with no effect on accuracy.

Diffusion modeling

In this section, we more carefully consider the mechanism that produces our mapping effect from the standpoint of a particular sequential sampling model, the diffusion model (Ratcliff & McKoon, 2008). The diffusion model assumes that evidence accumulation proceeds toward an upper or lower boundary, with an average drift rate representing the strength of evidence gleaned from the stimulus. Accumulation continues until one of the boundaries is reached, and then the response associated with that boundary is made. Biases can be introduced by moving the starting point of evidence accumulation closer to one boundary than the other or by shifting drift rates up or down across all conditions (i.e., changing the "drift criterion"; Starns, Ratcliff, & White, 2012). Our goals for this section are (1) to make sure that parameters with our joystick responding are similar to previous fits of the asterisk

task and (2) to compare the empirical results with predictions produced by changing different model parameters. We do not use the model to try to estimate different nondecision times across responses or compare models with and without different nondecision times, because the model is not able to do this effectively with a typical number of observations (Voss et al., 2010). The full modeling details are described in the Diffusion Modeling section available in the Supplementary Material and the Open Science Framework (https://osf.io/q486r/).

To compare with previous fits, we fit the constant-mapping data using standard modeling procedures. The results were well in line with previous asterisk-task fits, as detailed in the Diffusion Modeling supplement. This suggests that our joystick response procedure produces results comparable to a keypress procedure. To explore the mechanisms driving our mapping effects, we took the best fitting parameters for each participant, adjusted one parameter to match the mapping effect in median correct RT, and compared the mapping effect on accuracy with these parameters to the observed accuracy effect.[2] We considered scenarios in which the mapping effect was produced because the variable-mapping procedure disrupts the bias manipulation (in terms of either stating point or drift-rate biases) or because it attenuates the nondecision time advantage for expected responses. Accommodating the mapping effect on RT in terms of starting-point bias produces a .07 mapping effect on accuracy, and drift-rate bias has a .25 mapping effect. The observed mapping effect on accuracy was -0.003 (95% CI [-0.041, 0.036]), much more in line with the null effect predicted by the nondecision time mechanism than the effects predicted by a starting-point or drift-bias mechanism.

General discussion

Our results support the claim that both motor and decision factors contribute to RT bias effects in simple decision tasks. Supporting a role for decision factors, we found a clear effect of bias on RT even in the variable-mapping conditions of Experiment 2—conditions that should minimize nondecision factors like motor preparation. The relative size of the bias effects in Experiment 1 versus Experiment 2 also support a role for decision processes. Across both mapping conditions, bias effects were larger in the task that involved making decisions from uncertain evidence (Experiment 2) than the task that used easily discriminable arrow stimuli (Experiment 1). Supporting a role for nondecision factors, participants in Experiment 1a responded more quickly when they knew which movement would be required (one-arrow blocks) than when they could not anticipate the movement (two- and four-arrow blocks). For two-arrow blocks with unequal stimulus

---

[2] We used data from Experiment 2b because the within-subject design allowed us to define a mapping effect for each participant.

frequencies, participants were faster to make the expected movement than the unexpected movement in Experiments 1a and 1b. These effects provide evidence that motor execution times are faster when the required response can be anticipated. Alternative explanations are possible, however, as we cannot be sure that the task completely eliminates the need to accumulate evidence. Stronger support for nondecision factors is provided by the fact that bias effects in the asterisk task were larger with constant mapping than with variable mapping. Moreover, the fact that the mapping effect was observed in RT but not accuracy shows that it is best explained by nondecision time as opposed to bias parameters.

Researchers usually ignore potential nondecision effects when modeling bias manipulations (e.g., Criss, 2010; Ratcliff & McKoon, 2008; Starns, Ratcliff, & McKoon, 2012; Voss et al., 2004), and Voss et al. (2010) demonstrated that this practice distorts parameter estimates if nondecision times truly differ across responses. Our results show that a standard bias manipulation does affect nondecision (motor execution) time, so theorists will need to explicitly model nondecision effects to make accurate conclusions about the cognitive processes that produce response biases. Voss et al. proposed a model with free nondecision time parameters for each response instead of a single parameter across all responses (the standard practice). This solution could be beneficial for experiments with a very large number of conditions and many observations, but introducing free nondecision parameters greatly increases model flexibility and results in poor parameter recovery in typical research scenarios (Voss et al., 2010). Verdonck and Tuerlinckx (2016) developed a method for estimating diffusion model parameters without specifying the distribution of nondecision times, and this technique might provide a more accurate characterization of bias effects on decision processes than fitting the standard diffusion model. Both of these techniques rely on the assumption that the diffusion process accurately characterizes the decision times in order to estimate nondecision time.

Cognitive psychologists are faced with the difficult task of identifying and measuring underlying processes from observable behavior. Decision theorists have progressed admirably in terms of identifying processing components from fits to accuracy and RT data, with many reported successes in detecting changes in components such as response caution and information quality (e.g., Ratcliff & McKoon, 2008; Voss et al., 2004). Despite many successes, the models still handle nondecision processes in a rudimentary fashion. Our results show that thinking more carefully about the interplay of decision and nondecision processes will be an important advance.

References

Criss, A. H. (2010). Differentiation and response bias in episodic memory: Evidence from reaction time distributions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 36,* 484–499.

Fay, P. J. (1936). The effect of cigarette smoking on simple and choice reaction time to colored lights. *Journal of Experimental Psychology, 19,* 592–603.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20,* 873–922.

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

Ratcliff, R., Thapar, A., & McKoon, G. (2001). The effects of aging on reaction time in a signal detection task. *Psychology and Aging, 16,* 3323–3341.

Starns, J. J., Ratcliff, R., & McKoon, G. (2012). Evaluating the unequal-variance and dual-process explanations of zROC slopes with response time data and the diffusion model. *Cognitive Psychology, 64,* 1–34.

Starns, J. J., Ratcliff, R., & White, C. N. (2012). Diffusion model drift rates can be influenced by decision processes: An analysis of the strength-based mirror effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 1137–1151.

Verdonck, S., & Tuerlinckx, F. (2016). Factoring out nondecision time in choice time data: Theory and implications. *Psychological Review, 123,* 208–218.

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition, 32,* 1206–1220.

Voss, A., Voss, J., & Klauer, K. C. (2010). Separating response-execution bias from decision bias: Arguments for an additional parameter in Ratcliff's diffusion model. *British Journal of Mathematical and Statistical Psychology, 63,* 539–555.