UvA-DARE (Digital Academic Repository)

Modeling distracted performance

Hawkins, G.E.; Mittner, M.; Forstmann, B.U.; Heathcote, A.

DOI

[10.1016/j.cogpsych.2019.05.002](https://doi.org/10.1016/j.cogpsych.2019.05.002)

Publication date

2019

Document Version

Final published version

Published in

Cognitive Psychology

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Hawkins, G. E., Mittner, M., Forstmann, B. U., & Heathcote, A. (2019). Modeling distracted performance. *Cognitive Psychology*, 112, 48-80.
<https://doi.org/10.1016/j.cogpsych.2019.05.002>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Modeling distracted performance^{☆,☆☆}Guy E. Hawkins^{a,*}, Matthias Mittner^b, Birte U. Forstmann^c, Andrew Heathcote^d^a School of Psychology, University of Newcastle, Australia^b Department of Psychology, University of Tromsø, Norway^c Integrative Model-Based Cognitive Neuroscience Unit, University of Amsterdam, the Netherlands^d School of Medicine – Department of Psychology, University of Tasmania, Australia

ARTICLE INFO

Keywords:

Mind wandering
Task-unrelated thought
Sustained attention
Decision making
Evidence accumulation
Cognitive model

ABSTRACT

The sustained attention to response task (SART) has been the primary method of studying the phenomenon of mind wandering. We develop and experimentally test the first integrated cognitive process model that quantitatively explains all stationary features of behavioral performance in the SART. The model assumes that performance is generated by a competitive race between a stimulus-related decision process and a stimulus-unrelated rhythmic response process. We propose that the stimulus-unrelated process entrains to timing regularities in the task environment, and is unconditionally triggered as a habit or ‘insurance policy’ to protect against the deleterious effects of mind wandering on ongoing task performance. For two SART experiments the model provided a quantitatively precise account of a range of previously reported trends in choice, response time and self-reported mind wandering data. It also accounted for three previously unidentified features of response time distributions that place critical constraints on cognitive models of performance in situations when people might engage in task-unrelated thoughts. Furthermore, the parameters of the rhythmic race model were meaningfully associated with participants’ self-reported distraction, even though the model was never informed by these data. In a validation test, we disrupted the latent rhythmic component with a manipulation of inter-trial-interval variability, and showed that the architecture of the model provided insight into its counter-intuitive effect. We conclude that performance in the presence of mind wandering can be conceived as a competitive latent decision vs. rhythmic response process. We discuss how the rhythmic race model is not restricted to the study of distraction or mind wandering; it is applicable to any domain requiring repetitive responding where evidence accumulation is assumed to be an underlying principle of behavior.

1. Introduction

“Mind wandering” – losing track of time, place, or current task goals – is very common in everyday life, occupying up to 50% of

* Funding: This work was supported by: Australian Research Council (ARC) Discovery Early Career Researcher Award (Hawkins, DE170100177); ARC Discovery Project (Hawkins, DP180103613); European Research Council Council Starting Grant (Forstmann); Netherlands Organisation for Scientific Research Vidi Grant (Forstmann); ARC Discovery Project (Heathcote, DP160101891); Visiting Professor Grant, University of Amsterdam (Heathcote). The funding sources had no role in the study design; in the collection, analysis and interpretation of data; in the writing of the report; and in the decision to submit the article for publication.

☆☆ Declarations of interest: none.

* Corresponding author address: School of Psychology, University of Newcastle, University Drive, Callaghan NSW 2308, Australia.

E-mail address: guy.hawkins@newcastle.edu.au (G.E. Hawkins).

<https://doi.org/10.1016/j.cogpsych.2019.05.002>

Accepted 10 May 2019

Available online 23 May 2019

0010-0285/ © 2019 Elsevier Inc. All rights reserved.

our waking hours (Killingsworth & Gilbert, 2010). Phenomena related to mind wandering have been referred to as task-unrelated thought, stimulus-independent thought, distraction or attentional lapses, among others, and can have positive effects (e.g., freeing the mind to think creatively, or prospectively thinking about the future; Smallwood & Schooler, 2015). However, task-unrelated thoughts and distraction are most commonly studied in terms of their negative impact, particularly on ongoing task performance. That is, when thoughts drift from the task at hand, how is performance on that task affected? The most common findings are that mind wandering leads to greater error rates, more variable responding and response lapses, and deficits in reading comprehension and working memory (for review, see Mooneyham & Schooler, 2013). It is particularly important to understand the underlying causes of the negative outcomes of mind wandering in light of the contexts where mind wandering is most likely to occur: situations requiring vigilance, such as monitoring tasks, or sustained attention, in tasks such as driving, reading and comprehension, and executive control (Smallwood & Schooler, 2015).

Various theoretical proposals attempt to explain task performance in the presence of mind wandering. For example, one key hypothesis states that executive resources are called upon to complete goal-directed tasks, and this finite pool of resources is depleted when the mind wanders: internally-directed thoughts consume mental energy, and as such there is a diminished pool of resources available to commit to the ongoing external task, which reduces task performance (Smallwood & Schooler, 2006; Teasdale et al., 1995). This “executive resources” account has been keenly debated against an alternative “executive failure” account (McVay & Kane, 2010, 2012; Smallwood, 2010), which proposes that executive control over thoughts and behavior is achieved proactively – such as actively maintaining a task set with the current goal in mind. If proactive control fails, the current goal escapes the task set, subsequently leading to distraction and goal neglect (Duncan, Emslie, Williams, Johnson, & Freer, 1996). Yet another proposal is that people alternate between a task-focused state of perceptual coupling – where attentional processes are directed toward sensory input – and a task-disengaged state of perceptual decoupling – where attention is diverted from sensory inputs toward inner thoughts (for reviews, see Schooler et al., 2011; Smallwood & Schooler, 2015).

Such theorizing has led to considerable insights into the potential generators and effects of mind wandering. Nevertheless, we argue that further theoretical progress is likely to be limited as long as the field continues to operate with such verbally specified or qualitative theories. Qualitative theories cannot, by definition, generate precise quantitative predictions for observed behavior, which can be necessary to discriminate between qualitatively-described theories (Lewandowsky & Farrell, 2011). We propose this can only be remedied by the development and testing of integrated cognitive process models of performance in the presence of mind wandering. As we have previously argued (Hawkins, Mittner, Boekel, Heathcote, & Forstmann, 2015), cognitive process models of task performance will be instructive in enabling more decisive arbitration among competing theories of mind wandering.

1.1. Quantitative implementations of qualitative theories of mind wandering

To conceptually illustrate arbitration through modeling, Hawkins, Mittner, Forstmann, and Heathcote (2017) outlined how two of the qualitative theoretical positions in the mind wandering literature could be mapped to quantitative cognitive process model implementations. They proposed a *discrete state representation* where, at any moment in time, task performance is generated by one of two discrete latent states – on-task or off-task – each of which is specified with a set of potentially different model parameters. When specified in this manner, as the respondent switches from the on-task to the off-task state, or vice versa, the set of model parameters that generates performance also changes, which allows the model to generate different behavioral predictions as a function of mind wandering. This provides a potentially deeper link between observed behavior and mind wandering theorizing, because changes in different model parameters over the latent states will lead to different predictions for observed performance across those states, including differential causes for more variable responding, or higher error rates. This approach, therefore, provides a quantitative approximation to the perceptual decoupling theory, and was successfully tested from a neural perspective in Mittner et al. (2014). In their work, the perceptually decoupled (off-task) state reduced the efficiency with which people acquired information about target stimuli, and slightly decreased cautiousness in responding, relative to the perceptually coupled (on-task) state.

In a similar vein, executive resources theories of mind wandering could be quantitatively specified in terms of Hawkins et al.’s (2017) *continuous dimension representation*. From this perspective, mind wandering is represented as a latent continuum that varies from completely on-task through to completely off-task. At any moment in time, a respondent will fall somewhere along the task-focus continuum, where the position provides a natural yet quantitative mapping from the latent (mind wandering) state to the data-generating parameters of the cognitive model. Although this mapping of resource theories to a continuous latent dimension does not exhaustively account for all features of the qualitative proposals, it does appear to capture the key motivating idea: ‘resources’ are a fluid concept that varies across time, such that the resources allocated to a goal-directed task can dynamically increase (greater attention) or decrease (less attention), on the assumption that the quantity of resources allocated to a task influences performance. The latent continuum could even be mapped to external covariates; for example, to neural signatures thought to track components of mind wandering, such as activation in the default mode network (Christoff, Gordon, Smallwood, Smith, & Schooler, 2009; Mason et al., 2007; Weissman, Roberts, Visscher, & Woldorff, 2006). In this manner, the activation of a neural signal thought to reflect the task-focused state of the respondent is structurally mapped to trial-by-trial changes in behavior (for similar applications in other domains, see Cavanagh et al., 2011; Frank et al., 2015; Nunez, Srinivasan, & Vandekerckhove, 2015; Nunez, Vandekerckhove, & Srinivasan, 2017).

These examples highlight two ways that qualitatively described theories might be implemented in well-developed quantitative modeling architectures. Although we do not believe there is a conceptually exhaustive, one-to-one mapping from the qualitative theories (such as perceptual coupling) to quantitative models (e.g., a discrete state representation), we do argue that using the qualitative theories to aid development of a set of quantitative cognitive models that generate different predictions will lead to

further theoretical clarity between the alternative proposals. Such discrimination between theories is a key goal in the scientific study of mind wandering research. However, a necessary first step toward achieving this goal is to develop a quantitatively precise cognitive model of task performance in the presence of mind wandering. We emphasize that this step is *not* a model of mind wandering per se. Rather it is a model of observed performance while people may, or may not, be focused on the external task at hand. We argue that this is a prerequisite to fulfilling the broader goal as it will allow one to quantitatively discriminate between potential latent architectures underlying mind wandering; that is, to tease apart predictions of, for example, the discrete state and continuous dimension representations. In the absence of a quantitatively precise model of task performance, it will be conceptually and practically challenging to reliably discriminate between the latent generators of mind wandering.

Some steps have been taken toward developing quantitative theories of mind wandering and their predictions for behavioral performance. There has also been similar progress in related contexts such as the study of fatigued performance (Gunzelmann, Gross, Gluck, & Dinges, 2009; Walsh, Gunzelmann, & Van Dongen, 2017). As we outline below, although these developments have provided insights, they have tended to emphasize theorizing about the latent mechanisms of mind wandering over a quantitatively precise explanation of behavioral data. As a consequence, all existing theories fall short in terms of their ability to completely account for observed behavior. We address this issue here.

Almost all previous quantitative theories of mind wandering have studied performance in the sustained attention to response task (SART), as it is the most widely used experimental paradigm in the study of mind wandering. Given the predominance of the paradigm, we also focus on it, first providing an overview of the SART and common experimental findings in the paradigm, followed by an overview of previously proposed quantitative theories of mind wandering in the SART. We note, however, that the quantitative model developed here generalizes beyond this paradigm, a point we return to in the General Discussion.

1.2. Sustained attention during mind wandering

The sustained attention to response task (SART; e.g., Robertson, Manly, Andrade, Baddeley, & Yiend, 1997; Smallwood et al., 2004; Smallwood & Schooler, 2006; Smilek, Carriere, & Cheyne, 2010) is a go-nogo task with a small proportion of (“target”) trials that require one to withhold a response. The standard format is that each decision trial involves presentation of a single digit, randomly selected from the digits 1–9. If the digit is the target number, such as 3, the participant is to withhold their response; if any other digit is shown the participant is to press a response button (the same button irrespective of digit). Minor variants of the SART use different stimuli (e.g., letters, character strings, colors, etc.), but generate the same key trends in data that are described below. In all cases, the target stimulus is relatively rare, which induces a pattern of habitual, very fast responding and a relatively high proportion of commission errors (i.e., ‘go’ response to the target stimulus) and a low proportion of omission errors (i.e., ‘nogo’ response to non-target stimuli).

As Robertson et al. (1997) originally proposed, the aim of the SART is to habituate the respondent to the repetitive, non-arousing qualities of the stimulus, leading to automatic patterns of responding and hence frequent responding to target (nogo) stimuli. Faster mean response times tend to be observed preceding errors of commission relative to correct responses to the target stimulus (i.e., withholding response; Manly, Robertson, Galloway, & Hawkins, 1999; Robertson et al., 1997). This result is also observed across participants: people with faster mean response times across the entirety of the task tend to make more commission errors (Manly, Davison, Heutink, Galloway, & Robertson, 2000). These findings suggest the SART induces a speed-accuracy tradeoff in performance (for similar argument, see Dang, Figueroa, & Helton, 2018), and that the best account of SART performance might, therefore, be a framework that naturally accounts for speed-accuracy tradeoffs.

Although the SART was originally developed to study the effect of brain injury on sustained attention, it has been widely adopted in the study of mind wandering due to its simple structure and non-arousing nature that allows the respondent to take part in various cognitive activities during task completion, including thoughts unrelated to the task, if they so wish. To index mind wandering, participants are occasionally interrupted with *thought probes* that ask them to introspectively judge whether they were focused on-task or off-task in the preceding trial or series of trials (e.g., Giambra, 1995; for review, see Smallwood & Schooler, 2006).

Through the use of thought probes, much investigation in the SART and related tasks has found performance to be strongly affected during phases of self-reported mind wandering. For example, prior to reporting off-task versus on-task thoughts, people tend to produce more variable response times and are more likely to incorrectly respond when they should have withheld their response (a greater number of commission errors; e.g., Bastian & Sackur, 2013; Cheyne, Solman, Carriere, & Smilek, 2009; Leszczynski et al., 2017; McVay & Kane, 2009, 2012; Mrazek, Smallwood, & Schooler, 2012; Stawarczyk, Majerus, Maj, Van der Linden, & D’Argembeau, 2011); similar effects have also been observed in conceptually related sustained attention tasks (e.g., Esterman, Noonan, Rosenberg, & Degutis, 2013; Seli, Cheyne, & Smilek, 2013). Mean response times are sometimes faster prior to off-task relative to on-task thoughts (e.g., McVay & Kane, 2009, 2012), and sometimes slower (e.g., Stawarczyk et al., 2011).

1.3. Cognitive models of sustained attention and mind wandering

1.3.1. Cognitive models of aggregate-level data

A few cognitive process models have been proposed to account for the trends observed in group-level data in the SART (i.e., performance averaged across participants). For example, Peebles and Bothell (2004) considered performance in the SART as a competition between two strategies, from the perspective of the Adaptive Control of Thought–Rational (ACT-R) framework (Anderson & Lebiere, 1998). The first strategy involves rapid detection of the stimulus, where a response is given based on the presence of the stimulus but not its identity. The second strategy is slower and deliberative, based on the detection and subsequent

identification of the stimulus and its associated correct response. The strategy used on each trial is determined by the current utilities of both strategies, which in turn depend on the history of correct and erroneous responses that followed the use of each strategy, and the (time) costs associated with performing each strategy. With exposure to the task, the model learns greater utility (i.e., expected usefulness) for the simpler stimulus-detection strategy, which is correct most of the time given the low rate of targets; in this sense the model learns to minimize a joint function of expected response time and error rate. The model captures a number of trends in the choice data (cf. Fig. 3, Peebles & Bothell, 2004), including the expected number of errors of commission (responding to targets) and omission (failing to respond to non-targets), yet fails to predict any meaningful differences in response times (cf. Fig. 3, Peebles & Bothell, 2004). As outlined above, and demonstrated in detail in Experiment 1 below, response time effects are a crucial component of performance in the presence of mind wandering.

Building on Peebles and Bothell's (2004) two-strategy competition model of the SART, Van Vugt, Taatgen, Sackur, and Bastian (2015) proposed that mind wandering is composed of two sub-models, each with a particular goal: a task-attending sub-model with the goal of focusing on the task, and a distracted sub-model with the goal of pursuing internal thoughts. An intriguing feature of this approach is that mind wandering is explicitly represented as a latent process; a "distracted" state, similar to the perceptually decoupled state, or Hawkins et al.'s (2017) discrete state representation. As in Peebles and Bothell (2004), the goal that is pursued at any moment in time is a function of each sub-model's activation; that is, the goal of attending to the task, or factors unrelated to the task. Critically, activation of the "attend" goal is not constant throughout a task. When its activation drops the distraction goal takes over, and the distracted sub-model continually retrieves items from declarative memory, which form the content of mind wandering, until the system encounters an item in declarative memory that reminds it to attend to the task, which subsequently increases activation of the "attend" goal. In this sense, Van Vugt et al.'s (2015) model can be considered an exemplar computational implementation of the executive failure hypothesis of mind wandering (e.g., McVay & Kane, 2009), where mind wandering arises due to a failure to proactively maintain a goal to attend to the current task. Although conceptually intriguing, and able to capture the overall frequency of self-reported mind wandering reasonably well, Van Vugt et al.'s (2015) approach does not provide a strong account of the data, even at the group level; the model incorrectly predicts much lower error rates and variability in response times than observed.

In a similar vein, Hiatt and Trafton (2015) also built on the work of Peebles and Bothell (2004), proposing that mind wandering arises when functions of executive control fail, but in a different manner to Van Vugt et al. (2015). In particular, they proposed that mind wandering begins when there is a natural break in task-oriented thought, which can even occur when someone is working toward their goal but not currently actively reasoning about that goal. In this way the model can continue to respond to the task while mind wandering, and is therefore consistent with existing qualitative theories in the literature, such as the perceptual decoupling hypothesis (Schooler et al., 2011), and recent neural theories consistent with perceptual decoupling and task performance in the presence of off-task thoughts (Mittner, Hawkins, Boekel, & Forstmann, 2016). Nevertheless, although Hiatt and Trafton's (2015) model accounts for the frequency of mind wandering, it misses features in the error rate data, and makes no predictions for response times.

A common feature across the few existing quantitative models of mind wandering is an emphasis on providing an account of the cognitive processes that control the frequency of mind wandering during task performance rather than providing a precise account of task performance itself. When the models have been evaluated against task performance, it has only been at the aggregate level (i.e., data collapsed across participants). It is well known that trends at the group level can obscure important individual differences (e.g., Estes, 1956; Estes & Maddox, 2005), and that more detailed insight into underlying cognitive processes driving task performance can be obtained by evaluating cognitive models at the level of individual participants (e.g., Lee & Webb, 2005) or below (e.g., at the single-trial level).

1.3.2. Cognitive models of individual-participant data

Where the emphasis is on providing an account of individuals' performance in the presence of mind wandering, models from the ACT-R tradition are arguably challenging to apply, and so attention has shifted to the class of sequential sampling models, also known as *evidence accumulation* models, which have been broadly applied to speeded decision making in the psychology and neuroscience literatures (e.g., Brown & Heathcote, 2008; Busemeyer & Townsend, 1993; Ratcliff, 1978; Usher & McClelland, 2001; Van Zandt, Colonius, & Proctor, 2000; Vickers, 1979; for reviews, see Forstmann, Ratcliff, & Wagenmakers, 2016; Gold & Shadlen, 2007; Ratcliff & Smith, 2004). Evidence accumulation models assume that simple decisions – such as responding to non-targets and withholding a response to targets – are made through a process of gradually accumulating sensory information to a threshold.

We are aware of two studies to date that have used evidence accumulation models to analyze task performance during mind wandering. In their SART data, McVay and Kane (2012) observed that the accumulation of sensory information was more variable across decisions for people who were more prone to mind wandering. This result, which can only be considered across rather than within participants, provides an interesting conceptual link to the greater response time variability observed in SART data that is typically associated with off-task thoughts. Second, although Mittner et al. (2014) focused on a different experimental paradigm, a stop-signal task, they found that when participants were off-task relative to on-task, the rate at which they acquired information for the correct response, and for the stop signal, was lower, and they accumulated less evidence before making a response. Mittner et al.'s (2014) approach enables understanding of the latent processes underlying task performance, that is, the cognitive processes affected during on-task and off-task states, within a participant.

A key factor for these two existing studies is that they employed off-the-shelf evidence accumulation models: McVay and Kane (2012) used the Linear Ballistic Accumulator (LBA; Brown & Heathcote, 2008) and Mittner et al. (2014) used a stochastic accumulator-based implementation of the independent race model (Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014). This turns out

to be a problematic assumption when they are applied to the SART task. As we demonstrate in Experiment 1, performance in a standard SART experiment gives rise to response time distributions that do not conform to a number of the trends typically observed in rapid decision-making studies. It follows that evidence accumulation models designed to account for response time data typically observed in decision-making studies may fail to account for response times observed in the SART, a possibility we verify later. This conclusion holds even when considering evidence accumulation models modified for performance in go-nogo tasks (Gomez, Ratcliff, & Perea, 2007; Ratcliff, Huang-Pollock, & McKoon, 2018). To our knowledge, no existing cognitive model can account for the combination of response time patterns that we report.

1.4. Mind wandering as a rhythmic response process

Here, we propose the first integrated cognitive process model of task performance in the presence of mind wandering. Unlike others (e.g., Van Vugt et al., 2015), we do not aim to identify the content of task unrelated thoughts, nor do we propose a model of mind wandering. Rather, our aim is to develop and validate a quantitatively precise account of task performance in the presence of mind wandering, which is achievable without making any assumptions about the content of mind wandering. We argue that such a quantitatively precise account of individual performance is a prerequisite to understanding mind wandering more broadly. In our opinion, once we understand the relationship between the latent constructs of a cognitive model and observed performance, we can begin to hypothesize, develop and test quantitative mechanisms of the interplay between mind wandering and those latent model constructs (cf. Hawkins et al., 2017), which in turn will allow for the development of deeper, more complete quantitative theories addressing issues such as the content of mind wandering.

Our proposed cognitive model places a key emphasis on ‘autopilot’ behavior, where people maintain performance on an external task, typically with acceptable levels of accuracy, even when mental focus drifts from the external task. We propose that autopilot behavior leads to rhythmic response patterns – responding that is unconditionally triggered by the onset of external events, independent of stimulus identity – which turn out to be critical in SART performance. Below, where we describe the model in detail, we motivate the hypothesis that such rhythmic processing is central to performance in the SART. The newly developed model provides an excellent account of the observed choice and response time data, and theoretically meaningful interpretation of latent constructs related to performance in the presence of mind wandering.

We first report a standard SART experiment and demonstrate heretofore unexplored response time trends in the SART data that are incompatible with existing cognitive process models typically used to understand response time data. Next, we propose a simple modification to the evidence accumulation model framework that emphasizes autopilot behavior, which allows it to provide a precise account of the full range of choice and response time data in the SART. Finally, we provide a validation and novel test of the model in a new data set.

2. Experiment 1

To begin, we report a standard SART experiment to highlight the set of benchmark response time phenomena that a complete cognitive process model of task performance must explain.

3. Method

3.1. Participants

Twenty undergraduate psychology students (8 female, 12 male) from the University of Tromsø participated in the experiment for course credit. The age-range was from 20 to 32 years ($M = 23.85$, $SD = 2.9$). Due to a technical issue, data from one participant were not recorded, leaving data from 19 participants for analysis.

3.2. Materials

3.2.1. SART

Stimuli were presented with a custom program written in the simulation software Psychopy (Peirce, 2009) on a MacBook Pro (Retina display, 15-in.). Participants completed the experiment on a laptop computer in a testing laboratory. Each trial consisted of a single digit, randomly sampled from the digits 1–9, presented in black font in a large point size in the center of a display with a light gray background. If the presented digit was 3, the participant was instructed to withhold their response. If the presented digit was 1–2 or 4–9, the participant was instructed to respond by pressing the space bar. Participants completed 640 go trials and 80 nogo trials for a total of 720 SART trials.

A trial began with a centered fixation cross that was displayed for .25 s. The fixation cross was then replaced with the digit stimulus for the trial, also displayed for .25 s, which was followed by a blank display. The participant could provide a response from the time the digit was presented. The blank display was shown for .9 s. Once the ISI period elapsed, the next trial commenced with the display of the fixation cross.

3.2.2. Thought probes

Participants were occasionally presented with *thought probes* throughout the SART. Thought probes are a self-report measure

commonly used in the mind wandering literature to assess when participants were focusing on the external task (i.e., the SART) or focusing elsewhere (e.g., an internal stream of thoughts; Smallwood & Schooler, 2006). Thought probes consisted of the question “Where was your attention during the previous trial?” and were presented with a 4-point Likert scale with labels “on-task” (leftmost position on scale) and “off-task” (rightmost position on scale). Participants moved a slider along the 4-point scale with left and right arrow keys and pressed a button to confirm their response. The initial location of the slider along the four possible positions on the scale was randomized on each probe trial. For simplicity throughout, we interpret Likert scale responses as a value from 1 (on-task) through to 4 (off-task).

Participants responded to 20 thought probes throughout the experiment. Probe trials were pseudo-randomly dispersed amongst SART trials such that there was a minimum of 30 and maximum of 45 SART trials between successive probe trials. Within the 30–45 trial window, the trial on which a thought probe would appear was uniformly sampled.

3.3. Procedure

Participants provided informed consent prior to participation. They were informed that they would complete a simple, computerized decision-making task and were given brief instructions for the experimental task. They were informed prior to beginning the experiment that they would be periodically asked whether they were paying attention to the task. Participants were instructed that it was more important for them to respond truthfully than to simply report that they were paying attention (even if they were not). Participation lasted for approximately half an hour.

This data set was originally collected with the intention of analyzing facial expressions, so in addition to recording behavioral performance participants were instructed that their face would be recorded by a video camera. Participants were seated roughly 1 m from the screen. Stimuli were adjusted such that they did not exceed a length of 5.24 cm which corresponds to 3 degrees of visual angle. To ensure adequate and stable lighting conditions, a 11 W fluorescent tube office lamp with a flat 15 cm × 28 cm lamp shade-reflector was installed 60 cm above the laptop and facing the participant at a distance of about 1.3 m. A blank sheet of white A4 paper was taped onto the lamp, fully covering the shade reflector, to reduce brightness. After the behavioral task, participants completed the Mindful Attention Awareness Scale (MAAS; Brown & Ryan, 2003) and were asked about their experience with the task. Data from these latter measures (video recording, questionnaire responses) are not further analyzed here.

4. Results

We first report an analysis that demonstrates our experiment reproduces the key trends observed in typical SART experiments. We then report three patterns in SART response time data that are challenging to explain from the perspective of standard evidence accumulation models.

4.1. Data analysis

We focused on 4 outcome variables that are prominent in the SART/mind-wandering literature, each of which were calculated across the trials leading up to a thought probe: mean response time (RT), coefficient of variation in response times (RTCV), proportion of omission errors, and proportion of commission errors. We examined these 4 outcome variables across a number of window sizes, where the window denotes the number of trials prior to a probe response that was included in the calculation of the outcome variable. For example, a window of size 5 would mean the 5 SART trials ending with the probe trial, separately for each probe trial. Our motivation for testing different window sizes was to ensure the robustness of effects across minor variations in trial numbers.

One challenge in analyzing these outcome variables as a function of another dependent variable – that is, probe responses – is that it leads to an uneven number of observations in each cell of the design (possible probe responses) across participants. To avoid these issues, we developed a custom hierarchical Bayesian analysis that appropriately accounted for the uneven cell numbers across participants. Conceptually, our analysis approach provides a group-level estimate for each outcome variable at each position on the probe scale; that is, for the trials that precede each type of probe response (on-task through to off-task), we estimate the group mean RT, group mean RTCV, and mean proportion of omission and commission errors. We report the results of this analysis in the main text and refer the reader to [Appendix A](#) for details of the analysis.

4.1.1. Mean RT

Posterior distributions of the group-level mean RT for selected window sizes are shown in the upper row of [Fig. 1](#). It is clear that mean RT was faster in the trials preceding self-reported off-task thoughts relative to on-task thoughts, and this effect did not appear to differ for smaller versus larger window sizes (i.e., there was no interaction with window size).

Our statistical inference on the group-level posterior distributions compared the outcome measure prior to off-task responses (4 on the probe scale) versus on-task responses (1 on the probe scale). Although [Fig. 1](#) demonstrates a smooth gradation between the two end points of the probe scale for each of the 4 outcome measures, we restricted our analysis in this manner because our primary interest is how the end points of the scale differ (on-task vs. off-task). We repeated the comparison across all window sizes, calculating for each the median and the 95% highest density interval (HDI; Kruschke, 2011) – the smallest interval that contains 95% of the density of a distribution. If the 95% HDI calculated on the difference in the posterior distribution between on-task and off-task responses did not contain 0, we concluded there was a difference in the outcome measure.

When considering the 3 responses immediately preceding a thought probe, participants were approximately 85 ms faster before

reporting they were off-task on the subsequent probe trial compared to on-task, 95%HDI [−133 ms, −38 ms]. This relationship held with a similar magnitude across all examined window sizes (i.e., 95%HDI excluded 0 for $w = 4, 5, \dots, 10$). This finding is consistent with McVay and Kane (2009, 2012) and Smallwood, McSpadden, and Schooler (2007).

4.1.2. RT Coefficient Of Variation (RTCV)

As shown in the second row of Fig. 1, responses were more variable (larger RTCV) in the SART trials preceding reports of off-task thoughts relative to on-task thoughts (95% HDI across all window sizes excluded 0). This replicates a common finding in the mind wandering literature (Bastian & Sackur, 2013; Cheyne et al., 2009; Leszczynski et al., 2017; McVay & Kane, 2009, 2012; Mrazek et al., 2012; Stawarczyk et al., 2011). We confirmed that increasing RTCV across probe responses was not entirely driven by corresponding decreases in mean RT, by repeating the analysis on the RT standard deviation (data not shown).

4.1.3. Proportion of omission and commission errors

The third and fourth rows of Fig. 1 show the probability of an omission and commission error, respectively, across window sizes. Although the probability of an omission error tended to increase prior to self-reported off-task thoughts, the effect was small, reflecting the overall scarcity of omission errors. This is further emphasized by the fact that there was only a statistically reliable yet very small effect for the largest 2 window sizes we tested ($w = 9, 10$, 95%HDI of the difference distribution [.001,.020] for both w). This result is broadly consistent with the literature, which tends to observe that participants rarely fail to respond to non-targets.

Fig. 1. Posterior distributions of the group-level effects from the Bayesian data analysis of the 4 outcome measures (shown in rows) in Experiment 1. The estimated outcome measures (y-axes) are shown as a function of each probe response (x-axes). The window size, the number of trials preceding each thought probe response over which the outcome measure was calculated, is shown with the shade of the ‘violin’ plots. The violin plots represent the posterior distribution of the estimated parameters, where each ‘violin’ combines a boxplot and a kernel density estimate. The boxplot component is indicated with the white circular symbol (median), the interquartile range (heavy vertical line), and 1.5x interquartile range as an indicator of the range of the posterior distribution (thin vertical line). The ‘violin’-like shape of each distribution is obtained through a smoothed density estimate of the posterior distribution, rotated vertically, and plotted on both sides of the box plot to create a symmetric figure. The width of the violin is proportional to the number of samples that fall in that part of the posterior distribution.

Fig. 2. Individual participant response time distributions from Experiment 1.

In contrast to omission errors, there was a large increase in the probability of a commission error prior to off-task reports ($\sim .89$) relative to on-task reports ($\sim .24$). This replicates the most common SART result in the mind wandering literature (for reviews, see Mooneyham & Schooler, 2013; Smallwood & Schooler, 2015). The 95% HDIs excluded 0 across all examined window sizes.

4.2. SART response time distributions

Now that we have demonstrated the data from Experiment 1 reproduces trends typically observed in the mind wandering literature, we outline three trends in SART RT distributions that present problems for standard evidence accumulation models. We note that evidence accumulation models have previously been modified to account for go-nogo tasks such as the SART (Gomez et al., 2007; Ratcliff et al., 2018). However, even with these modifications, these models are unable to account for all patterns present in SART data.

4.2.1. Very fast responses

Fig. 2 shows distributions of individual participant RTs from Experiment 1. Every participant has at least some responses that are faster than what is generally considered possible for a regular decision process (e.g., faster than .15 s; Luce, 1986). Some participants gave few of these very fast responses (e.g., participants 1 and 12), so for these participants it might be permissible to proceed with the standard practice of excluding those fast trials from further analysis. For other participants, however, these very fast responses represented more than a trivial number of trials (e.g., participants 3 and 11), with some participants producing distributions that were skewed toward many very fast responses (e.g., participants 2 and 13). One approach would be to exclude trials with very fast RTs from all participants according to some cutoff (e.g., McVay & Kane, 2012), and exclude the two participants that produced particularly unusual distributions from further analysis. However, a more satisfying and complete account of the data would account for all trials from all participants. We believe that this is a desirable approach as it is possible that the very fast responses might be important for developing and testing theoretical accounts of performance in the presence of mind wandering.

4.2.2. Shallow leading edge of the RT distribution

Even if all of the very fast responses were excluded from analysis, Fig. 2 shows there is a second issue related to fast RTs: a shallow slope on the leading edge of the distribution. It appears as if the SART leads to a pattern of responses that slowly ramp up to the peak of the body of responses – a pattern produced by every participant. In most cases this led to an almost Gaussian shaped RT

distribution, which is rarely observed unless external factors have manipulated the decision environment (e.g., response deadlines; Evans & Hawkins, 2019; Hawkins, Forstmann, Wagenmakers, Ratcliff, & Brown, 2015). In contrast, standard RT data typically display a steep leading edge, characterized by a sharp, sudden onset to the RT distribution where many responses occur (i.e., the heavy positive skew of most RT distributions; Luce, 1986). Since this is the typically observed pattern in RT data, it is the pattern that standard models of RT data predict.

To our knowledge, the only means by which a conventional evidence accumulation model can predict this pattern of data would be to assume a large amount of ‘noise’ (variance) in the distribution of non-decision times, which leads to a shallower leading edge. Even with this modification, our initial explorations found that the standard models were unable to predict leading edges as extreme as observed in the data from Experiment 1. Even if they did, though, this would likely not be a psychologically satisfying approach to modeling data, as it would assume that almost all of the variability in RTs was due to variability in encoding and response production processes. Although this may be possible a priori, it would lead to the conclusion that the SART is a qualitatively different type of task even though it seems plausible that it shares cognitive structure with other go-nogo tasks that do not demonstrate the RT patterns described here (e.g., Gomez et al., 2007; Ratcliff et al., 2018).

4.2.3. The distribution of target RTs (errors) is very fast relative to non-target RTs (corrects)

Although previous investigations of the SART have not considered in great detail the relative speed of correct and error responses, analyses of such data have proven very insightful in understanding cognitive processing in other decision contexts (e.g., Ratcliff & Smith, 2004; Ratcliff, Smith, Brown, & McKoon, 2016). Fig. 3 shows defective cumulative distribution functions (CDFs) for RTs in target (solid lines) and non-target trials (dashed lines), separately for each participant. When splitting observed responses into corrects (go responses on non-target trials) and errors (go responses on target trials), Fig. 3 makes it clear that almost all participants produced a pattern of data that is consistent with a censored distribution: error RTs were almost-identically distributed to correct RTs, up until a certain time. After this time, which differed across participants, error RTs were no longer observed. This led to an effect where target RTs (errors) were considerably faster than non-target RTs (corrects), on average. Conventional evidence accumulation models can predict the latter result – faster mean RTs for errors than corrects – but they do so in a manner that assumes very fast responses are more likely to be erroneous. This was not the trend observed in data. A given fast RT was equally likely to occur on non-target trials (when it was correct) as on target trials (when it was incorrect). It follows that, given a response was fast, it was about 8

Fig. 3. Defective cumulative distribution functions (CDFs) of individual participant non-target (dashed lines) and target (solid lines) response times from Experiment 1. The y-axes represent the cumulative probability of a response and the x-axes represent response time. Panels show individual participant data, where the lower right of each panel displays the proportion of go (i.e., observed) responses in non-target (go) and target (nogo) trials for each participant; $p(\text{go}|\text{go})$ and $p(\text{go}|\text{nogo})$, respectively.

times more likely to be correct (i.e., a non-target trial) than an error (i.e., a target trial, owing to differences in the base rate of target and non-target trials).

To our knowledge, no existing cognitive model can parsimoniously and simultaneously account for the three patterns in SART RT data within a conventional evidence accumulation model framework, or more generally any cognitive modeling framework. We begin the next section with a brief review of the utility of evidence accumulation models with a focus on their potential for application to the SART and related paradigms. We then provide an overview of previous efforts to explain atypical or ‘contaminant’ responses in the context of evidence accumulation models, followed by our primary innovation: the rhythmic race model.

5. Evidence accumulation models of performance during mind wandering

Evidence accumulation models assume that speeded decisions are made through a process of sequentially sampling information from a stimulus. This information is accrued in one or more evidence counters that track support for the available response options. Once the evidence in one of the counters has reached a pre-determined quantity – the response threshold – a decision is triggered for the threshold-crossing option. This general evidence accumulation framework has been extremely powerful in understanding decision phenomena across a wide range of contexts from simple perceptual decisions through to complex discrete choices (e.g., Hawkins et al., 2014), neural data (e.g., Forstmann et al., 2008), primate decision making (e.g., Gold & Shadlen, 2007), and even clinical populations (e.g., Heathcote, Suraev, Curley, Gong, & Love, 2015), alcohol consumption (e.g., Van Ravenzwaaij, Dutilh, & Wagenmakers, 2012), and sleep deprivation (e.g., Walsh et al., 2017). For reviews of the history and current status of evidence accumulation models, we refer the reader to Forstmann et al. (2016) and Ratcliff et al. (2016).

Not only does our theoretical approach using evidence accumulation models capitalize on the framework’s strong history of accounting for individual participant data, it also naturally accounts for an analysis problem that has been identified in previous SART research: the speed-accuracy tradeoff. The speed-accuracy tradeoff occurs when respondents elect to make faster responses at the expense of a lower probability of a correct choice, or vice versa – making more accurate but slower decisions (for early work, see Pachella, 1974; Reed, 1974; Swensson, 1972; Wickelgren, 1977; for recent review, see Heitz, 2014). In the SART, the speed-accuracy tradeoff manifests as an across-participant negative association between mean RT and errors of commission (Dang et al., 2018; Manly et al., 2000), which has been shown to confound other analyses of interest (Seli, Cheyne, & Smilek, 2012; Seli, Jonker, Cheyne, & Smilek, 2013; Seli, Jonker, Solman, Cheyne, & Smilek, 2013). Some researchers have gone so far as to propose ad-hoc regression-based approaches to circumvent the influence of the speed-accuracy tradeoff on analyses of various SART variables (e.g., Seli et al., 2013).

We argue that such ad-hoc analysis techniques are unnecessary given that the architecture of evidence accumulation models provides an elegant tool to naturally predict the observed covariation between choices and response times (i.e., the speed-accuracy tradeoff). Specifically, one can strategically alter the amount of evidence accrued prior to making a decision, by varying the height of the response threshold; collecting more evidence leads to slower responses with a higher probability of a correct response. This balance between response speed and accuracy has been alluded to in the context of previous models of the SART. For example, Peebles and Bothell (2004) suggested that performance might not be based on strategy selection as hypothesized in ACT-R-based models (e.g., Hiatt & Trafton, 2015; Van Vugt et al., 2015), but rather from balancing the competing demands of responding rapidly and minimizing error; decision optimality of precisely this form has been studied from the perspective of evidence accumulation models for decades (e.g., Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Wald & Wolfowitz, 1948). In this sense, the use of evidence accumulation models to understand performance in the SART, and more generally in the presence of mind wandering, is a natural fit that bypasses the need to develop ad-hoc statistical analysis techniques.

5.1. Contaminant mixture distributions in evidence accumulation models

A potential shortcoming in applying evidence accumulation models, particularly in the presence of mind wandering, is the models’ strong assumption that every response is generated from a process of evidence accumulation. That is, the models assume that participants never generate a response from another latent process, such as an off-task state. This assumption stands in contrast to some of the key theoretical proposals and trends observed in data in the mind wandering literature, such as increased errors of commission and more variable response times during periods of off-task thought. Even when mind wandering is not the explicit focus of study, the assumption still conflicts with most researchers’ intuitions and personal experience that participants’ attention is not always directed toward the decision tasks they have been assigned, which typically involve completing hundreds of trials in hour-long sessions.

Responses generated outside the evidence accumulation process of interest have been referred to as ‘contaminants’, or noise in the cognitive system, and there have been attempts to quantitatively account for them. For example, Ratcliff and Tuerlinckx (2002) assumed a contaminant mixture distribution in the likelihood function of the diffusion decision model (DDM), specified as a uniform distribution between the minimum and maximum observed RTs. This approach assumes there is a small probability of contamination, such that a subset of trials was generated from the uniform contaminant distribution rather than the evidence accumulation process of the DDM. Vandekerckhove, Tuerlinckx, and Lee (2008) implemented the contaminant mixture process in a Bayesian latent mixture modeling framework and extended it to assume that performance on any particular trial could be generated from the DDM, a uniform contaminant distribution, or a delayed-startup DDM (i.e., longer non-decision time), demonstrating the flexibility of Bayesian analysis frameworks to allow for complex mixtures of generating distributions.

Although the addition of a contaminant mixture distribution stabilizes parameter estimates (Ratcliff & Tuerlinckx, 2002), it

provides no insight into *why* some responses were not generated by the process of evidence accumulation. Furthermore, since it assumes that an observed choice-RT pair is generated from the evidence accumulation process *or* the contaminant distribution, it must assume that there is a proportion of trials in which the participant never initiated the evidence accumulation process, regardless of the duration of the observed RT. This begs the question of why this occurs.

5.2. Racing contaminants: the Rhythmic Race Model (RRM)

Here, we take a conceptually different approach and assume that a decision arises from a race between the evidence accumulation process and a contaminant process. The key difference in our proposal is that the contaminant process *races against* the evidence accumulation process; this is an integrated process model of contamination rather than the previous either-or assumption. Our approach provides an explanation for failures of the evidence accumulation process *during* the course of a decision: a contaminant response is produced when the contaminant process terminates before the evidence accumulation process. This means our approach does not need to assume a pre-determined failure probability prior to trial onset, as in contaminant mixture distributions, and means the model is constrained to make strong predictions about when contaminant responses occur relative to the timing of non-contaminated responses. With this simple assumption incorporated in an evidence accumulation model, we account for a range of non-evidence-accumulation driven responses, including very fast responses, failures of the decision process, and simply responding in a rhythm matched to the timing of the external task.

We refer to our model as the rhythmic race model (RRM). The RRM consists of a race between two latent processes: a *stimulus-related decision process* and a *stimulus-unrelated rhythmic response process*. The *stimulus-related decision process* is defined by an evidence accumulation model of the same form as they are typically studied in the literature. The specific evidence accumulation model one chooses to incorporate in the RRM is peripheral to our thesis. For example, to model performance in the SART one could use an independent race model such as the Linear Ballistic Accumulator (LBA) model, in which case there is an accumulator that collects evidence in favor of a ‘go’ response and a second, independent accumulator that collects evidence in favor of a ‘nogo’ response, or a random walk or diffusion model such as the DDM, in which case there is a single decision variable that simultaneously collects evidence for and against the ‘go’ and ‘nogo’ responses. We do not focus on this choice of evidence accumulation model because it does not influence our substantive results or conclusions; any evidence accumulation model that provides an adequate account of typical choice and RT distributions is sufficient. We incorporated the LBA here for computational simplicity, however we emphasize that the stimulus-related decision process could also have been a DDM.

The *stimulus-unrelated rhythmic response process* contains a single runner that represents a ‘rhythmic’ response.¹ We could have chosen different terminology for this latent process, such as an unconditionally-triggered response process. However, we believe that the terminology and concepts associated with rhythmic responding have precedence in the mind wandering literature. For example, performance errors in the SART – such as mistakenly responding to a to-be-inhibited item, or increased response variability – have been linked to mindless or habitual responding (McVay & Kane, 2009; Robertson et al., 1997; Smallwood et al., 2004).

We refer to the latent stimulus-unrelated process as rhythmic as we assume it is triggered by the highly regular sequence of external events – the ‘rhythm’ of the environment – that a participant experiences during completion of many trials in the experimental psychology laboratory, and in real life. In these tasks, events often occur in frequent, discrete units, each of which requires some form of active information processing, and where timing typically becomes an important and predictable component of performance; for example, aiming to complete X jobs – experimental trials, scanning items at a supermarket, or laying bricks to build a wall – before time Y – end of experimental session, or shift at work. During such repetitive tasks, it is common for people to start responding in a habitual manner with minimal or no active information processing or awareness of the current task requirements; colloquially, we refer to this phenomenon as ‘going on autopilot’ (for review, see Gouraud, Delorme, & Berberian, 2017). In this sense, we propose that frequently occurring events in the environment, such as the onset of a stimulus in an experimental task, start to unconditionally trigger a response, such as pressing a button; this is a form of instrumental learning.

The autopilot behavior encoded in the stimulus-unrelated rhythmic response process can be seen as a habit-like response strategy, or an ‘insurance policy’ against the negative effects of mind wandering. We argue that this process is rhythmic in the sense that it is unconditionally triggered by the onset of an external stimulus – such as the fixation cross in our paradigm, which provides a cue to stimulus onset – which has highly predictable timing in the SART when the inter-stimulus interval is constant across trials. We argue that this highly predictable timing very quickly comes to act as a metronome to which participants match or *entrain* their responses; this is the sense in which we use the term rhythmic. The rhythmic process we describe can be thought of as an effective habit-like response strategy or as an insurance policy because it protects the decision maker against some of the potentially deleterious effects of mind wandering on performance: unconditionally triggering a response following a rhythmically-timed external cue (like the fixation cross) will provide a correct response on 8 of every 9 SART trials, on average, without *any* attention to stimulus identity.

Our proposal has links to the interval timing literature, where the canonical task asks participants to repeatedly reproduce experimenter-specified intervals with button presses. We believe that the similarity between repeatedly reproducing externally-specified intervals and our latent process that generates a response at a stereotypical interval following onset of a cue such as a fixation cross provides an interesting avenue to pursue links between literatures. It also provides an avenue to consider potential

¹ We refer to the rhythmic response process as containing a *runner* rather than an *accumulator* or *decision variable*. This is because a rhythmic response does not accumulate information. Rather it is a distribution of finishing times that ‘runs’ against, and may terminate prior to, the stimulus-related decision process.

implementations of the rhythmic response process in the brain, since neurally-based theories of interval timing as an accumulation process have received substantial attention (Simen, Balci, deSouza, Cohen, & Holmes, 2011; Simen, Vlasov, & Papadakis, 2016).

We wish to be clear that we are not proposing the rhythmic response process is mind wandering in some sense. Rather, the insurance policy provided by the rhythmic response process allows a participant to continue performing a task even in the absence of full attention to that task, such as when they are mind wandering. That is, we assume that regularly occurring events in the external environment – like the onset of a fixation cross – unconditionally trigger a response process. With this simple assumption, the rhythmic response process of the RRM produces 89% accuracy on average, yet will also produce atypical patterns in the predicted time to produce each response. We propose that it is these contaminant responses, generated from a rhythmic response process, that give rise to the three patterns in SART RT data that are challenging for conventional evidence accumulation models to explain. We note, however, that the RRM is not restricted to use in the SART, or the study of mind wandering. Regardless of whether mind wandering is the explicit focus of study, the RRM is a cognitive process model of performance in the presence of contaminant responses, so it can be utilized in any context where evidence accumulation models are applied. In this way, the RRM allows the researcher to avoid removing any ‘outliers’ from the observed data, thus providing a more complete understanding of the latent cognitive processes under investigation.

5.3. RRM architecture

The RRM assumes a race between a stimulus-related decision process and a stimulus-unrelated rhythmic responding process with independently distributed finishing times. Throughout the main text we describe and apply the RRM with a stimulus-related decision process specified as a LBA. We provide parallel text specifying the RRM architecture in terms of the DDM in Appendix C. A go response can be generated in two ways:

1. When the go accumulator reaches threshold before the nogo accumulator *and* the rhythmic runner is slower than the time taken for the go accumulator to reach threshold.
2. The rhythmic runner is faster than the time taken for the go and nogo accumulators to reach threshold.

To specify the race architecture we denote the probability density function (PDF) that the *go* and *nogo* accumulators have reached threshold at time t as $f_{go}(t)$ and $f_{nogo}(t)$, respectively, with corresponding cumulative distribution functions (CDF) $F_{go}(t)$ and $F_{nogo}(t)$. Similarly, we denote the PDF and CDF of the rhythmic runner producing a response at time t as $f_{rhythmic}(t)$ and $F_{rhythmic}(t)$, respectively (the form of this distribution is specified below). The PDF of a go response at time t is therefore

$$PDF_{go}(t) = f_{go}(t) \times (1 - F_{nogo}(t)) \times (1 - F_{rhythmic}(t)) + f_{rhythmic}(t) \times (1 - F_{go}(t)) \times (1 - F_{nogo}(t)).$$

A nogo “response” (i.e., withheld response) only occurs if the nogo accumulator reaches threshold before the go accumulator *and* the rhythmic runner is slower than the time taken for the nogo accumulator to reach threshold. Since there is no observed response when the nogo accumulator crosses threshold (i.e., this accumulator corresponds to withholding a response), we integrate over the predicted distribution of response times for a nogo response. The probability of a nogo response is:

$$p(nogo) = \int_{t=0}^{\infty} f_{nogo}(t) \times (1 - F_{go}(t)) \times (1 - F_{rhythmic}(t)) dt.$$

5.4. RRM specification and parameterization

The RRM as implemented here has the parameters of a regular LBA model and additional parameters for the stimulus-unrelated rhythmic response process. We now outline the parameters of the full model followed by the particular constraints we used in our application, first focusing on the parameters of the stimulus-related decision process followed by the stimulus-unrelated rhythmic response process. It is straightforward to substitute another evidence accumulation model and its parameters in place of the LBA, such as the DDM as specified in Appendix C, so we do not discuss these methods further.

5.4.1. Parameters of the stimulus-related decision process

The RRM as implemented here assumes the PDF and CDF of the LBA model reported in Brown and Heathcote (2008) with the drift rate distribution truncated to positive values (Heathcote & Love, 2012). The model contains the following parameters for the two LBA accumulators. There is a mean drift rate, restricted to non-negative values, for the go accumulator to go (non-target) stimuli, $d_{go|go}$ (i.e., the stimulus is one of the numbers 1–2 or 4–9; the go response is correct), and nogo (target) stimuli, $d_{go|nogo}$ (i.e., the stimulus is the number 3; the go response is incorrect), with Gaussian-distributed across-trial variability in drift rate with standard deviation s_{go} . Similarly, there is a drift rate for the nogo accumulator to nogo stimuli, $d_{nogo|nogo}$ (the nogo response is correct) and to go stimuli, $d_{nogo|go}$ (the nogo response is incorrect), with across-trial variability s_{nogo} . The non-drift parameters include the maximum value of the start point distribution, where the start-point of the evidence accumulation process is uniformly distributed between 0 and A independently across LBA accumulators and trials; response thresholds for the go and nogo accumulators, b_{go} and b_{nogo} , and time outside the decision process associated with encoding the stimulus and producing a response, known as non-decision time, t_0 (t_0 is only identifiable when RT is observed, i.e., from go responses).

We imposed sensible constraints on the decision-related parameters of the RRM so as to reduce model freedom. Specifically, there

were almost no trials where the participant failed to give a go response to go stimuli, since go trials made up almost 90% of all trials, so we constrained the drift rate for the corresponding accumulator to 0 ($d_{nogo|go} = 0$). Since there are no observed RTs associated with the nogo response, across-trial variability in drift rate for the nogo accumulator is poorly constrained by data, since this parameter is constrained by the relative speed of correct versus error responses, so we assumed that $s_{go} = s_{nogo}$ in all applications of the model. Similarly, the estimate of an accumulator's response threshold is influenced by the relative frequency of a response *and* its speed. Accordingly, since the time associated with the nogo accumulator is unobserved, we cannot reliably estimate separate thresholds for the go and nogo accumulators, so we assumed $b_{go} = b_{nogo}$. We also set this common value at $b = 1$ as the scaling parameter of the LBA model (Donkin, Brown, & Heathcote, 2009).

5.4.2. Parameters of the stimulus-unrelated rhythmic response process

We assume that the rhythmic runner of the RRM follows a Weibull distribution, which has a long history in accounting for the shape of RT distributions in cognitive psychology (e.g., Logan, 1988; Van Zandt, 2000). Specifically, if a random sample from the Weibull distribution is faster (i.e., has a lower value) than the time taken for the go or nogo accumulators to reach threshold, a rhythmic response is given (i.e., button press, the same response as if the go accumulator had crossed threshold first). We note that alternative descriptive distributions of RTs including the ex-Gaussian, shifted Wald or Gamma are also likely to produce an approximately equivalent descriptive account of the data as the Weibull, since this collection of descriptive distributions tends to demonstrate relatively high levels of mimicry in their ability to capture empirical RT distributions. We chose the Weibull distribution over these alternatives due to a clear mapping from its two core parameters to the latent construct of rhythmic responding.

We assume that the rhythmic runner is triggered by onset of the fixation cross (250 ms prior to stimulus onset), and follows a two-parameter Weibull distribution with parameters shape (k) and scale (λ), and PDF and CDF for $t > 0$,

$$f_{rhythmic}(t; k, \lambda) = \frac{k}{\lambda} \left(\frac{t}{\lambda}\right)^{k-1} e^{-\left(\frac{t}{\lambda}\right)^k},$$

$$F_{rhythmic}(t; k, \lambda) = 1 - e^{-\left(\frac{t}{\lambda}\right)^k}.$$

We do not estimate the 'shift' parameter of the Weibull distribution throughout as we assume this response process was reliably and unconditionally triggered at onset of the fixation cross, effectively setting the shift parameter to a value of 250 ms for all participants.² The shape parameter (k), commonly referred to as the 'failure rate' of the Weibull distribution, influences the likelihood of producing a rhythmic response as trial duration increases. Fig. 4 (left panel) illustrates a range of Weibull distributions that can be generated for different values of the shape parameter. When $k \leq 1$, shown in the two lightest blue lines, most rhythmic responses occur early, and rhythmic responses becomes less likely over time. When $k > 1$, shown in the two darkest blue lines, rhythmic responses occur around a clear modal time. The scale parameter (λ) of the Weibull distribution shifts that mode, 'stretching' or 'contracting' the effect of the shape parameter over trial time; this can be thought of as the time scale over which rhythmic responses occur, which will be sensitive to the overall speed of the task. That is, for a given k , the likelihood of a rhythmic response at time t is increased or decreased depending on the value of λ ; this effect is illustrated in Fig. 4 (right panel).

If the stimulus-unrelated response process reflects an insurance policy against mind wandering – that is, an unconditionally triggered process that ensures responding continues in a manner ostensibly similar to task-attending performance – then we expect the shape parameter (k) to be estimated at values greater than 1. This generates response distributions that more closely mimic typical RT distributions than shape parameters with values less than 1 (cf. left panel, Fig. 4); that is, unimodal where the mode is greater than 0. This contrasts to hypotheses based on time-minimization, which posit that participants might instead aim to minimize the overall amount of time invested in the external task (e.g., Hawkins, Brown, Steyvers, & Wagenmakers, 2012); this would manifest as values of the shape parameter less than 1 such that the stimulus-unrelated process predicts a peak at 0, the shortest RT possible. In this case the rhythmic process would more often win and so risk a response on a no-go trial, than a rhythmic process with the same average speed but a peak away from 0. Hence, shape parameters greater than 1 enable the model to better act as a habit-like response strategy or an insurance policy that protects against the negative side effects of mind wandering while reducing the probability of erroneous responses.

5.4.3. Parameter estimation

The final RRM reported here had 8 freely-estimated parameters, described in Table 1. We specified the RRM using the LBA likelihood functions for a single accumulator truncated to positive values (Heathcote & Love, 2012) as implemented in the *rtutils* package for the R statistical environment (Singmann, Brown, Gretton, & Heathcote, 2016) suitably modified to instantiate the RRM architecture.³ We used a hierarchical Bayesian framework to simultaneously estimate RRM parameters at the participant and group levels. We assumed mildly informative prior distributions for the group-level parameters, which allowed the estimated parameters to vary across a range of plausible values. Complete details of the prior distributions and parameter estimation routine can be found in

² The shift parameter moves the location of the starting point away from zero (where zero refers to onset of the fixation cross), representing the first time point with non-zero density. Since respondents could not provide a response prior to stimulus onset (250 ms post fixation cross), we assumed the shift was a constant fixed at 250 ms. We examined the robustness of this assumption in preliminary explorations of the RRM, where the shift parameter was reliably estimated at or marginally greater than 250 ms (i.e., at stimulus onset), for all participants. In all model applications reported in Experiment 1 the shift parameter was subsequently constrained to 250 ms for simpler estimation.

³ If one wishes to implement the DDM version of the RRM, the required likelihood functions are also contained in the *rtutils* package for R.

Fig. 4. Illustration of various Weibull distributions. The left panel plots a family of Weibull distributions across a range of shape parameters (k) with a fixed value of the scale parameter ($\lambda = .5$). The right panel plots a family of Weibull distributions across a range of scale parameters (λ) with a fixed value of the shape parameter ($k = 2.5$).

Table 1
Parameters of the Rhythmic Race Model (RRM) freely estimated in Experiment 1.

Latent process	Parameter	Description
Stimulus-related decision process	$d_{go go}$	Processing speed: go responses to non-targets.
	$d_{go nogo}$	Processing speed: go responses to targets.
	$d_{nogo nogo}$	Processing speed: nogo responses to targets.
	A	Variability in start point of evidence accumulation.
	s	Variability in processing speed.
	t_0	Non-decision time.
Stimulus-unrelated rhythmic response process	k	Relative timing of rhythmic responses (early, late).
	λ	Time scale of rhythmic responses.

Appendix D. We also conducted a parameter recovery study demonstrating that the parameters of the RRM are identifiable in realistically-sized data sets (**Appendix E**).

5.5. Model comparison

To determine whether the RRM provides an improved account of the data relative to existing methods for dealing with unusual RT data, we quantitatively compared it to five competitor models.⁴ For ease of exposition in this section, we refer to the RRM as the Racing Weibull. We compared the Racing Weibull to a contaminant mixture model, derived from Ratcliff and Tuerlinckx (2002). Specifically, we assumed that a response was generated from the LBA model with probability $1 - p$ and from a contaminant model with probability p . We assumed that contaminant responses consisted of a nogo (withheld) response on half of the trials and a go (observed) response on the other half of the trials, consistent with the assumptions of the standard contaminant mixture model. We estimated two forms of the contaminant mixture model that differed with respect to the predicted distribution of go responses: the first assumed a uniform distribution where the minimum and maximum were specified separately for each participant as the range of that participant’s observed RTs (Ratcliff & Tuerlinckx, 2002), and the second assumed a Weibull distribution that was parameterized in precisely the same form as the RRM; we refer to these as the Mixture Uniform and Mixture Weibull models, respectively, where the mixture parameter (p) was freely estimated at the participant and group levels (details in Appendix D). The Mixture Uniform model is of interest as it provides a comparison to the standard method for dealing with outlier RT data in the decision making literature. The Mixture Weibull model is of interest as it provides a test of the necessity of the proposed racing architecture of the RRM. For completeness, we filled the 2×2 matrix that crossed *psychological process* (race vs. mixture) with *contaminant distribution form*

⁴ Our model comparison was conducted amongst a family of LBA models so that each model differed only in its assumptions about the stimulus-unrelated rhythmic response process. One could conduct the parallel set of model comparisons with a DDM implementation of the RRM.

Table 2

Experiment 1 DIC model comparison values for the models under investigation, including the number of participant-level parameters estimated in each model (*N*). The model with the lowest DIC (the DIC-preferred model) is shown in bold.

Model	<i>N</i>	DIC
LBA with Racing Weibull (RRM)	8	−26,984
LBA with Racing Uniform	6	−22,672
LBA with Mixture Weibull	9	−26,912
LBA with Mixture Uniform	7	−25,881
LBA with Response Bias	7	−19,757
LBA	6	−19,396

(Weibull vs. Uniform), by also estimating a Racing Uniform model. The minimum and maximum of the Racing Uniform distribution were specified in precisely the same way as the Mixture Uniform model.

Finally, we also considered two variants of the regular two-choice LBA model with accumulators corresponding to the go and nogo responses. The first variant assumed asymmetric starting points for the go accumulator relative to the nogo accumulator, which translates to asymmetric response thresholds for the two LBA accumulators. Asymmetric response thresholds allows for response bias in a manner that could account for the very fast responses (lower threshold for the go relative to nogo accumulator) and provide an approximation to a normative account of the increased prior probability of a go response (details in [Appendix D](#)). The second variant was a vanilla two-choice LBA with no modifications, which was included primarily as a baseline comparison rather than a serious competitor model. Nevertheless, it is a complete version of the LBA model studied in [McVay and Kane \(2012\)](#). It was specified, parameterized and estimated exactly as described above for the RRM, with the obvious exception that all rhythmic response related components were removed.

We performed model comparison using the Deviance Information Criterion (DIC; [Spiegelhalter, Best, Carlin, & van der Linde, 2002](#)). The model with the lowest DIC is preferred.

5.6. Modeling results

For all models reported here, we did not exclude a single trial of data from any participant. This approach forces our models to account for the full range of data produced in an experimental session.

5.6.1. Model comparison

The DIC model comparison values for the six models are shown in [Table 2](#). The Racing Weibull (RRM) had a DIC approximately 70 units lower than the next best model, which was the Mixture Weibull. A DIC difference of 10 units between models is generally considered strong evidence in favor of the lower-DIC model (e.g., [Pratte, Rouder, & Morey, 2010](#)), so our result indicates quite strong evidence in favor of the RRM over all four competitor models. The summed-across-participants DICs shown in [Table 2](#) obscure an interesting result: only one participant had strong evidence in favor of the Mixture over Racing Weibull model – the participant with the most atypical RT distribution (participant 13 in [Fig. 4](#)) – but across all other participants the evidence was very strongly in favor of the Racing over Mixture Weibull model, with a DIC difference over 430 units (when excluding participant 13). Furthermore, the two variants of the regular LBA model provided a much poorer account of the data than any of the models with a contaminant distribution. Together, these outcomes provide strong evidence in favor of a racing process over a mixture, and for a Weibull-distributed component over the conventionally assumed Uniform distribution. These results therefore provide evidence in favor of an unconditionally triggered response process in the SART over previous conclusions that performance in the SART, and subsequently mind wandering, is solely associated with increased across-trial variability in drift rate (e.g., [McVay & Kane, 2012](#)).

5.6.2. RRM goodness of fit to data

The RRM provides an excellent quantitative account of all aspects of the individual participant data. [Fig. 5](#) shows individual participant RT distributions, collapsed across go and nogo trials, in the same format as [Fig. 2](#). The posterior predictive distribution is overlaid on each participant's observed responses. The figure shows that the model precisely captures two of the challenging RT trends in SART data: very fast responses and a shallow leading edge to the RT distribution; [Fig. 1 of the supplementary material](#) separates the predicted responses by latent component (stimulus-unrelated vs. stimulus-related). Other models under comparison – including the regular LBA previously studied in the mind wandering literature – fail to provide an adequate account of these data (see [Supplementary Material](#)).

To display the RRM's ability to quantitatively capture the proportion of go and nogo responses, and the shape of their corresponding RT distributions, [Fig. 6](#) shows defective CDFs of data (circular symbols connected with lines) and the posterior predictive distribution (crosses) separately for each participant. The data fall within the uncertainty region of the posterior predictive distribution, and are therefore captured by the model, when the circular symbols fall within the (*x*, *y*) region covered by the crosses. Orange lines show RTs on go trials (i.e., responses to stimuli 1–2, 4–9; correct responses), which reach (or are very close to) 1 on the *y*-axis for all participants; this means that go trials were virtually always responded to with a button press. The proportion of trials where participants gave a go response on a go trial is shown in the lower right of each panel for data, $p(\text{go}|\text{go}, D)$, and as predicted by

Fig. 5. Individual participant response time distributions (gray histograms) with overlaid posterior predictive distribution of the rhythmic race model (blue density curves) from Experiment 1. The posterior predictive distribution shows a sample of 25 predicted distributions of response times that were generated from 25 independent random samples of parameter values from the posterior distribution of the parameters. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the model, $p(\text{go}|\text{go}, M)$.

The blue lines in Fig. 6 show RTs from nogo trials (i.e., responses to stimulus 3; commission errors), clearly demonstrating the RRM's ability to capture the third trend in SART RT data: the distribution of error (nogo) RTs is much faster than the distribution of correct (go) RTs. Furthermore, the height (y-axis position) of the blue lines varies considerably across participants, indicating a large range in participants' ability to withhold an inappropriate response; for example, participant 12 had a low probability of making a commission error, while participant 2 had a very high probability. Again, the proportion of trials where participants gave a go response on a nogo trial is shown in each panel as $p(\text{go}|\text{nogo}, D)$, and as predicted by the model, $p(\text{go}|\text{nogo}, M)$. The RRM precisely captures the vastly differing abilities of participants to withhold an inappropriate response. The standard method for dealing with unusual RT data in the literature (Mixture Uniform, Table 2) misfit these data (see Supplementary Material).

The RRM quantitatively captured all aspects of the observed RT distributions, and by extension the three patterns in SART RT data that pose problems for standard evidence accumulation models (i.e., models similar to Model 2; DIC shown in Table 2, goodness of fit shown in Supplementary Material). This can be seen in Figs. 5 and 6 where there is a very close quantitative match between data (lines with circles) and the posterior predictive distribution (crosses) for all participants. This result is due to the stimulus-unrelated rhythmic response process. Firstly, the rhythmic runner allows the RRM to predict non-zero probability of termination very early in a trial – at values lower than the estimated non-decision time parameter of the stimulus-related decision process (group-level mean .12s, Table 2) – which predicts very fast responses.

Secondly, the RRM predicts a shallow leading edge of the RT distribution as the result of some responses having been generated from the slower stimulus-related decision process (go accumulator reaching threshold before the nogo accumulator or the rhythmic runner) and other responses generated from the faster stimulus-unrelated rhythmic responding process (the rhythmic runner being faster than the time for the go or nogo accumulators to reach threshold), effectively 'smoothing' the leading edge of the predicted distribution of the evidence accumulation model. Thirdly, the RRM predicts that the distribution of error (nogo) RTs is very fast relative to correct (go) RTs. This final prediction occurs because the RRM only predicts withheld responses when the nogo accumulator finishes first – before the go accumulator and the rhythmic runner. For nogo trials with an observed response, the model assumes that a high drift rate was sampled in the go accumulator or the rhythmic runner was fast, where either outcome generates a response before the nogo accumulator can reach threshold. Since the rhythmic runner has non-zero probability of early termination, it frequently terminates before the nogo accumulator can reach threshold, thus predicting faster error RTs (nogo trials) than correct

Fig. 6. Defective cumulative distribution functions showing data and the posterior predictive distribution of the rhythmic race model for individual participant data from Experiment 1. Points along the x-axes show the 10th, 30th, 50th (i.e., median), 70th and 90th percentiles of the RT distributions for go (correct; orange) and nogo (error; blue) responses. Points along the y-axes show the cumulative probability of a response. Data are shown as circular symbols connected with lines. The posterior predictive distribution for the 5 percentiles of each RT distribution is shown with crosses, where the center (respectively, horizontal and vertical arms) of the cross indicates the expectation (respectively, 95% credible intervals). Where the crosses cannot be seen, the 95% credible interval is smaller than the overlaid data symbol. The lower right of each panel displays the proportion of go (i.e., observed) responses on go and nogo trials observed in data for each participant, $p(\text{go}|\text{go},D)$ and $p(\text{go}|\text{nogo},D)$, respectively, and the same quantities as predicted by the model, $p(\text{go}|\text{go},M)$ and $p(\text{go}|\text{nogo},M)$.

RTs (go trials). Although the RRM predicts faster mean error RTs versus correct RTs, it also predicts approximately equivalent RT distributions for very fast responses, since most of the predicted responses arise from the fast rhythmic process which is blind to stimulus identity; this trend was also observed in data.

5.6.3. RRM parameter estimates

We report the 95% highest density interval (HDI) of the group-level marginal posterior distributions. As in the data analysis section, when comparing differences in the value of two estimated parameters, we computed the 95% HDI of the difference between the two distributions and concluded a difference in the parameters if the interval did not include 0.

The parameter estimates of the RRM provide insight into performance in the SART. Regarding the key parameters of the stimulus-related decision process, there was very little difference between the three drift rates (Table 3), and all pairwise comparisons between the three drift rates included 0. There are a few ways to interpret these estimates. At the group level, the stimulus information driving the input to the go accumulator was as strong on nogo trials (stimulus 3; error response) as it was on go trials (stimuli 1–2, 4–9; correct response). This result was consistently reflected at the participant level, too; the difference distribution between $d_{\text{go}|\text{go}}$ and $d_{\text{go}|\text{nogo}}$ included 0 for every participant. Similarly, at the group level there was again no difference in the strength of the accumulator inputs for the go and nogo accumulators on nogo trials. However, this group-level effect represents an average over different participant effects: the difference between $d_{\text{nogo}|\text{nogo}}$ and $d_{\text{go}|\text{nogo}}$ included 0 for 12 participants, whereas $d_{\text{nogo}|\text{nogo}} > d_{\text{go}|\text{nogo}}$ for 6 participants and $d_{\text{nogo}|\text{nogo}} < d_{\text{go}|\text{nogo}}$ for 1 participant. These differences suggest that different participants approached the task in different ways: the majority of participants (63%) reflected the group level trend, whereas a substantial minority of participants (32%) had greater sensitivity to the imperative nogo stimulus (i.e., $d_{\text{nogo}|\text{nogo}} > d_{\text{go}|\text{nogo}}$).

Table 3
Parameter estimates of the Rhythmic Race Model (RRM) in Experiment 1.

Latent process	Parameter	Posterior median	Posterior HDI (lower, upper)
Stimulus-related decision process	$d_{go go}$	3.50	3.25, 3.78
	$d_{go nogo}$	3.42	3.18, 3.70
	$d_{nogo nogo}$	3.60	3.33, 3.90
	t_0	.12	.11, .14
	A	.73	.58, .92
	s	.87	.73, 1.03
Stimulus-unrelated rhythmic response process	k	1.94	1.54, 2.45
	λ	.95	.62, 1.51

Regarding the parameters of the stimulus-unrelated process, if the unconditionally triggered rhythmic process reflects an insurance policy against the negative side effects of mind wandering then we expect the shape parameter (k) to be estimated at values greater than 1, as this generates response distributions that closely align with typical RT distributions compared to shape parameters less than 1 (cf. left panel, Fig. 4). This hypothesis was supported at the group level (Table 3); the majority of participants also showed this trend (15/19). The remaining participants' HDIs crossed 1 (3/19), meaning that it was unclear whether their stimulus-unrelated process had a predicted mode equal to or greater than 0, and a single participant whose rhythmic process had a predicted mode of 0. The interpretation for this final participant (13 in Figs. 5 and 6) is that they produced many very early responses, which they timed to occur just after stimulus-onset; this was the same (and only) participant for whom the Mixture Weibull model provided the most parsimonious account of their data.

5.7. Self-reported mind wandering and the RRM

If the RRM provides a process-level account of task performance in the presence of mind wandering, then the model's predictions ought to relate to self-reported mind wandering. We tested this hypothesis by determining the relative contribution of the stimulus-related and stimulus-unrelated processes to participants' behavioral data as a function of their thought probe responses.

Our data analysis (cf. Fig. 1) indicated that in the trials preceding participants' self-report that their attention was focused off-task they made faster responses that were more variable. To determine whether the RRM could account for these response patterns, we worked backwards to gain insight into the model's prediction at the single-trial level. If we denote the likelihood of the datum from trial j having been generated by the stimulus-related decision process conditioned on the stimulus-unrelated rhythmic response process not responding $L(\text{decision}|D_j)$, and the rhythmic response process conditioned on the decision process not responding $L(\text{rhythm}|D_j)$, ignoring notation for participants and parameters for now, then we can transform the two likelihoods into a relative likelihood. This estimates the probability that the response on trial j was generated from the stimulus-related decision process, $p(\text{decision}_j)$, given the likelihood of the datum under the decision and rhythmic response processes:

$$p(\text{decision}_j) = \frac{L(\text{decision}|D_j)}{L(\text{decision}|D_j) + L(\text{rhythm}|D_j)},$$

where $p(\text{rhythm}_j) = 1 - p(\text{decision}_j)$. Some trials must be assigned with probability 1 to one response distribution or the other given the structure of the model. For example, unobserved responses (i.e., the nogo accumulator crossed threshold first) have 0 likelihood under the rhythmic distribution, and RTs faster than the estimated non-decision time have 0 likelihood under the decision distribution. For all other responses, their relative speed is more or less likely under one distribution or the other (conditioned, of course, on the competing distribution not having terminated). With just this simple approach, we found that participants with greater mean $p(\text{decision})$ across all trials had lower mean rates of self-reported mind wandering ($r = -.44$, $BF_{10} = 2$),⁵ shown in the upper right of Fig. 7. Although the statistical evidence is weak (below the conventional threshold of $BF_{10} = 3$), this result is in the direction that participants who report more mind wandering might be more likely to also have a more effective insurance policy.

A more targeted test of the relationship between task performance and the RRM can be assessed through the correlations between self-reported mind wandering and the model parameters.⁶ As in the analysis of $p(\text{decision})$, we used each participant's mean self-report to the set of thought probes experienced across the task, and correlated this value with the posterior median of each participant's RRM parameter estimates. Mean self-reported mind wandering was strongly positively correlated with decision-to-decision variability in processing speed of the stimulus-related decision process (s ; $r = .66$, $BF_{10} = 19.4$). The association only emerged in the architecture of the RRM; it was absent in the regular LBA ($r = .08$, $BF_{10} = .50$). Parameter s reflects the consistency of task-related responding, so the association suggests that participants who more often reported off-task thoughts made more variable decisions (i.e., greater variability in the stimulus-related decision process), which is consistent with the findings of McVay and Kane (2012).

⁵ We report non-parametric rank-order correlation coefficients and their corresponding statistical tests throughout the manuscript so as to circumvent issues with potential assumption violations of the parametric correlation coefficient, including outliers.

⁶ We thank a reviewer for this suggestion.

Fig. 7. Analysis of thought probe data. The barplot (upper left) shows the distribution of thought probe responses, averaged across participants. The scatter plot (upper right) shows the across-participant relationship between the probability of a stimulus-related decision as predicted by the RRM ($p(\text{decision})$; y-axis) and self-reported task-unrelated thoughts (x-axis), where each point represents each participant's mean across both measures. The violin plot (lower row) shows the posterior distribution of the group-level effect from the Bayesian data analysis of $p(\text{decision})$, in the same format as Fig. 1. The estimated probability of a stimulus-related decision (y-axis) is shown as a function of each probe response (x-axis). The window size, the number of trials preceding each thought probe response over which $p(\text{decision})$ was averaged, is shown with the shade of the violin plots. All other details are as described in Fig. 1.

This finding can be interpreted through the lens of the continuous dimension representation introduced earlier: gradual fluctuations in attention across trials may cause changes to drift rate across trials (i.e., variability in drift rate, s). The more extreme those attentional fluctuations, the greater the variability in drift rate, which was picked up in our data.

There was also a reasonably strong negative correlation with mean processing speed to the nogo stimuli ($d_{\text{nogo}|\text{nogo}}$; $r = -.58$, $BF_{10} = 6.5$), indicating that people who more effectively identified the infrequent targets (nogo stimuli) were on average less likely to report off-task thoughts. All other correlations were modest at best ($|r| < .37$), with ambiguous evidence or evidence in favor of the null hypothesis ($BF_{10} < 1.2$). This indicates, for example, that mean efficiency at identifying frequent non-targets ($d_{\text{go}|\text{go}}$) or parameters of the rhythmic process (k , λ) were not associated with mind wandering.

The correlation between mean probe rating and $p(\text{decision})$ was lower in magnitude than the correlation between mean probe rating and s or $d_{\text{nogo}|\text{nogo}}$, suggesting that these two parameters were not synergizing to increase $p(\text{decision})$. Indeed, the opposite conclusion is more likely: combining the effect of all RRM parameters into the computation of $p(\text{decision})$ actually decreased the strength of the association (i.e., the correlation with $p(\text{decision})$ was weaker than the correlation with s or $d_{\text{nogo}|\text{nogo}}$). In terms of the model architecture, a larger value of s means the stimulus-related decision process is less likely to win the race, on average, implying that the stimulus-unrelated rhythmic response process is more likely to win the race, which together reduce $p(\text{decision})$. This suggests that it is the nature of changes in the decision process that determine $p(\text{decision})$, and hence the association with self-reported mind wandering, rather than the rhythmic response process. This argues for variability in the stimulus-related decision process as a potential indicator of mind wandering. This interpretation is consistent with our interpretation of the stimulus-unrelated process as an insurance policy against mind wandering rather than a process-level representation of mind wandering per se.

Given the association with RRM parameters can only be assessed across participants, for a finer-grained analysis we tested whether responses given at different positions along the thought probe scale were systematically related to the RRM-predicted probability of a stimulus-related decision in the trials preceding the probe response. The hypothesis was that higher values of

$p(\text{decision})$ would precede on-task reports, and lower values of $p(\text{decision})$ would precede off-task reports. To this end, we analyzed the RRM-predicted $p(\text{decision})$ in precisely the same form as the Bayesian data analyses of mean RT, RTCV and omission and commission errors reported earlier. Details are provided in [Appendix B](#). In brief: separately for each participant we computed mean $p(\text{decision}_j)$ over the $j = 1 \dots w$ trials preceding each probe response, where w corresponds to the window size as defined above, and treated these averaged values of $p(\text{decision}_j)$ as the outcome variable for the analysis.

There was a monotonically decreasing probability that responses were generated from the stimulus-related decision process when shifting from self-reported on-task through to off-task responding, shown in the lower row of [Fig. 7](#). The values in [Fig. 7](#) can be interpreted to mean that approximately 4 in every 5 trials that preceded a probe response of completely on-task (1) were more likely to have been generated under the stimulus-related decision process (conditioned on the stimulus-unrelated process not terminating) than the stimulus-unrelated process (conditioned on the stimulus-related process not terminating). For probe responses of completely off-task (4), the preceding trial was closer to equal probability of having been generated by the stimulus-related or stimulus-unrelated processes. Unlike the across-participant analysis above, Bayes factors are not available for this within-participant analysis, so we instead report the HDI of the posterior parameter estimates.

Using the same analysis approach as before, the 95%HDI of the difference in $p(\text{decision}_j)$ between on-task and off-task reports (i.e., positions 1 vs 4 on the probe scale) excluded 0 for all window sizes. Furthermore, the difference between on-task and off-task reports appeared to have subtle variation with window size: for SART trials that immediately preceded a probe trial, there was a median difference of $p(\text{decision}_j) \approx .25$ between end points of the subsequent probe report scale, which reduced to a difference of $p(\text{decision}_j) \approx .2$ by a window size of 10.

These results suggest there are temporal dynamics in play: if there is a distinct off-task process, as hypothesized in perceptual decoupling theory, then the interpretation through the lens of the RRM is that observed responses are almost equally likely to have been generated from the slower stimulus-related decision process or the faster stimulus-unrelated rhythmic response process. In contrast, on-task responding appears to be generated from something closer to a purer stimulus-related decision process, according to the introspective thought probe responses. It is important to note that these potentially time-varying trends are an emergent property of the stationary RRM that was applied to the data; $p(\text{decision})$ is not a parameter of the model, and the model has no notion of latent states that might vary across trials. Nevertheless, we argue that when the RRM is applied in the manner we have described here, it is useful in identifying quantitative trends in data that will subsequently serve as targets for a latent, time-varying, process-level quantitative model of mind wandering. We provide a proof-of-concept example of such an approach below.

Contributions from both the stimulus-related and stimulus-unrelated processes explain why an increasing likelihood of off-task responding is accompanied by decreasing mean RT but increasing RTCV relative to on-task responding, as reported in our data and previously observed in the literature: the faster rhythmic response process and the slower decision process are almost equally likely to generate the response. The resulting distribution of RTs is then comprised of faster responses, on average, that are also more variable. Compare this to RTs observed during on-task responding, where the rhythmic response process is less likely to win the competitive race, which are therefore comprised of RTs largely generated from the slower decision process; slower, on average, and less variable.

To quantitatively illustrate this point, we classified individual participant responses as more likely to have been ‘decision’ responses (i.e., trials where $p(\text{decision}) \geq .5$) or ‘rhythmic’ responses (i.e., trials where $p(\text{decision}) < .5$). Relative to the overall sample, where mean RT was .325s and RTCV was .29, ‘rhythmic’ responses were much faster on average than ‘decision’ responses ($M = .175\text{s}$ vs .352s, $BF_{10} > 1000$) and considerably more variable (RTCV = .45 vs .20, $BF_{10} > 1000$). This result is particularly impressive in light of the extant mind wandering literature showing that off-task responding is more variable than on-task responding ([Bastian & Sackur, 2013](#); [Cheyne et al., 2009](#); [Leszczynski et al., 2017](#); [McVay & Kane, 2009, 2012](#); [Mrazek et al., 2012](#); [Stawarczyk et al., 2011](#)), and occasionally faster on average, too ([McVay & Kane, 2009, 2012](#)). These trends emerged from the RRM – through the assignment of each response as having been generated by a ‘decision’ or ‘rhythmic’ latent process – without any knowledge of self-reported ratings of mind wandering.

Furthermore, the overall sample violates the widely observed relation that the mean and standard deviation (SD) of RT distributions are linearly related ([Wagenmakers & Brown, 2007](#)): participants’ mean RT and SD were not correlated ($r = -.03$, $BF_{10} = .49$).⁷ However, when trials were split into RRM-classified ‘decision’ or ‘rhythmic’ responses and the correlation recalculated, the linear relation appeared in the predicted direction for both types of responses ($r = .64$, $BF_{10} = 11.2$ and $r = .60$, $BF_{10} = 8.8$, respectively). This illustrates that considering responses generated from the two latent processes in isolation conforms to the expected linear mean-SD RT relation, yet the aggregated responses generated from both latent processes does not. Conventional evidence accumulation models conform to the linear mean-SD relation ([Wagenmakers & Brown, 2007](#)), which provides yet another line of evidence indicating that conventional evidence accumulation models cannot account for the trends in SART data.

5.8. Temporal dynamics and the RRM

As a final exploratory test, we provide a proof-of-concept example that outlines how the across-trial dynamics implied in the lower panel of [Fig. 7](#) might be incorporated in the RRM. The approach we outline here is descriptive: we covary a selected parameter of the RRM by a selected property of the task. This is not a process model of across-trial dynamics in latent states, nor is it intended to be. Rather, we see this descriptive approach as a stepping stone that highlights potential parameters to target in the model (i.e., latent

⁷ Unlike the correlation coefficients reported earlier in the main text and later in Experiment 2, the correlation coefficients reported in this paragraph are parametric, consistent with the hypothesized linear relationship between the mean and SD ([Wagenmakers & Brown, 2007](#)).

cognitive processes) that are influenced by changes across time in the task (across-trial dynamics in latent states).

We tested whether the shape (k) parameter of the stimulus-unrelated rhythmic response process had sequential dependencies across trials, covarying it by the uninterrupted run length of non-target (go) trials. The motivation for examining this feature was to test whether longer runs of the same stimulus type (where the correct response is to press the single response button) were associated with greater precision of the rhythmic response component; this would manifest as increasing k with increasing run length (cf. left panel, Fig. 4). One might expect this association since the timing properties become more salient with each stimulus repetition, allowing more precise entrainment to the task timing. We chose to descriptively explore dynamics in the shape parameter of the rhythmic response process rather than the two parameters that were more strongly correlated with mean self-reported mind wandering, across-trial variability in drift rate (s) and mean drift rate to target stimuli ($d_{\text{nogo|nogo}}$), for two reasons. First, variability in drift rate (s) is defined across trials, so to covary this parameter across trials makes little conceptual sense. Second, covarying mean processing speed to the rare target stimuli, which appeared only 1 in every 9 trials on average, would make it difficult to observe an effect, since this parameter appeared in the likelihood function in only 80 of the 720 trials.

Formally, the RRM was specified and estimated as described above for the DIC-best model, with the addition of a single parameter covariate – the uninterrupted run length of non-target (go) trials. This was calculated directly from each participant's data. The covariate was set to a value of 0 for the first non-target trial in the experiment, and was reset to 0 following each target (nogo) or thought probe trial. For all other trials, the covariate was incremented by 1 for each successive non-target trial in the sequence. The coefficient estimated on the covariate had prior distributions of $N(0, 1)$ and $\Gamma(1, 1)$ for the group-level mean and standard deviation parameters, respectively.

The RRM with estimated across-trial dynamics had a DIC of $-27,125$, which was ~ 140 units better than the stationary RRM reported above (cf. Table 2). This quantitative improvement in fit indicates there are sequential dependencies in data related to runs of the same trial type. The group-level mean of the coefficient on the run of non-target trials had a posterior median of .052, 95% HDI [.005, .104]. This estimate indicates that as the uninterrupted sequence of non-target trials grows longer – that is, as the task becomes more monotonous at a local scale – the latent rhythmic process becomes more precisely centered around its mean (i.e., less variable). We do not interpret this result in further detail, as it is purely descriptive and had a post hoc motivation.⁸ Nevertheless, we believe this model exploration is sufficient to demonstrate two important points: time-varying trends can be successfully incorporated into the RRM, and we highlight a potential pattern in data that could be a target for constraining a process-based model of mind wandering (i.e., increasingly longer sequences of the same stimulus and/or response type).

Taking the stationary and dynamic results together, it appears the stimulus-unrelated rhythmic response process captures substantial and meaningful variance in data, and that the parameter estimates of our novel cognitive process model contain information about the extent to which a participant is attending to their ongoing task. This link has often been stated in qualitative terms. However, to our knowledge, we have provided the first direct test of the quantitative link between self-reported task-unrelated thoughts and parameter estimates of a cognitive model of decision making.

6. Experiment 2

In Experiment 1 we demonstrated that the RRM provided a good account of all stationary aspects of SART data (behavior and self-reported mind wandering). In Experiment 2 we conducted a different test that demonstrated the necessity of the RRM's stimulus-unrelated rhythmic response process. For this test, we re-analyzed previously collected though unpublished data. This data set was withheld during the RRM development and testing described in Experiment 1, thus providing a valid generalization test of the model.

A key feature of the RRM is that the unconditionally triggered rhythmic response process is sensitive to non-stimulus related timing components of a task, such as changes in the inter-stimulus interval, and potentially other factors including participant boredom, but it is not sensitive to stimulus-related changes. It follows that if we manipulate a timing property of the experiment that is independent of the decision process (i.e., not related to the imperative decision stimulus) then it is likely to affect the parameters of the rhythmic response process. Furthermore, if all aspects of the imperative decision stimulus remain unchanged, then we expect to observe no change in the parameters of the stimulus-related decision process. Such a pattern of selective effects would provide further evidence that the rhythmic response process accounts for aspects of observed performance that are generated outside the decision process. It would also provide support for the convergent and divergent validity of the model.

7. Method

Participants completed the SART under two different conditions. In the first condition, which we refer to as the fixed inter-stimulus interval (ISI) condition, participants completed the SART as described in Experiment 1, with a constant ISI between trials. In the second condition, which we refer to as the random ISI condition, participants completed the SART as described in Experiment 1 with one exception: there was a variable ISI between trials.

Our guiding motivation was that participants entrain to the sequence of external events across trials, which directly affects the unconditionally-triggered rhythmic response process. In the fixed ISI condition, every trial had equal duration, so participants had a very reliable cue (the fixation cross) from which to prepare their response, unconditionally triggering a timed response process. The random ISI condition interrupted this ability to entrain to the task timing as the external metronome that cued the precise time until

⁸ We thank a reviewer for suggesting this model exploration.

stimulus onset (the fixation cross) was no longer predictable in the same way that it was in the fixed ISI condition. That is, even though the random ISI condition had equal trial duration on average and the fixation cross was displayed for the same duration as the fixed ISI condition, there was less reliable timing from one trial to the next. We propose that the variable trial duration led to lower precision in the timing component of the rhythmic response process, because the cue that unconditionally triggered that process was variable. A lower precision (more variable) rhythmic process can arise from a lower shape (k) parameter; lower k also increases the positive skew of the latent rhythmic distribution, which results in faster rhythmic responses relative to the stimulus-related decision process. Larger values of the Weibull scale parameter (λ) also increase variability of the rhythmic process, but they also drastically change the mean of the distribution, which is less plausible in this context given the experimental design imposed an upper limit on the trial duration. We therefore expected the fixed and random ISI conditions to differ in the shape alone or both shape and scale parameters of the rhythmic response process, but not the parameters of the decision process of the RRM.

Behaviorally, lower precision of the latent rhythmic process in the random relative to fixed ISI condition increases the predicted variability of the observed outputs of the rhythmic process. By definition, variability in the rhythmic process means that some trials will have a rhythmic process that is slower than average and therefore likely to be censored by the decision process, which would terminate the trial. On other trials, the rhythmic process will be faster than average, and is more likely to be faster than the decision process and will hence terminate the trial. With lower precision in the rhythmic process of the random ISI condition, there will be even greater variability – the slower rhythmic processes will be even slower (and still censored) and the faster rhythmic processes will be even faster (and even less likely to be censored) – which will manifest in behavioral data as a faster leading edge to the RT distribution in the random ISI condition compared to the fixed ISI. A faster leading edge will also increase the probability of errors of commission. Together, these behavioral effects can be interpreted as a less effective insurance policy.

Unless noted otherwise, methodological details were as described in Experiment 1.

7.1. Participants

Twenty-five undergraduate psychology students (19 females, 6 males, $M_{age} = 23.1$, $SD_{age} = 4.0$) from the University of Amsterdam participated in the experiment for course credit. The fixed and random ISI conditions were manipulated in separate blocks with order counterbalanced across participants. However, we only analyzed the first block (i.e., condition) completed by each participant. We did this because preliminary analysis of the data indicated there were very strong order effects: self-reported mind wandering was greater in the second block (condition) compared to the first block, independent of which condition was completed first ($M_{block1} = 2.64$ vs $M_{block2} = 2.96$ on a 5-point scale, $BF_{10} > 1000$). This interaction meant that condition effects (random vs fixed ISI) were potentially conflated with time-on-task effects (block 1 vs 2). The latter do not always appear, though they have been previously reported in the mind wandering literature, where longer time-on-task is associated with increased mind wandering (e.g., Boayue et al., in press; Smallwood, McSpadden, Luus, & Schooler, 2008; Stawarczyk et al., 2011). By analyzing the first block of trials, we effectively changed the experiment to a between-subjects design with 13 and 12 participants in the random and fixed ISI conditions, respectively. The lower sample size was compensated by a high per-participant trial count, which produced a clear pattern of results.

7.2. Materials

7.2.1. SART

Participants completed the experiment in a testing laboratory fitted with eye tracking hardware; eye movements were recorded but are not analyzed here, so we do not mention them further. All stimulus details were as described in Experiment 1. Participants completed 640 go trials and 80 nogo trials in each of the random and fixed ISI conditions for a total of 1440 SART trials. 720 trials per participant, from the first block, were included in the analysis.

A trial began with a centered fixation cross that was displayed for .25 s. The fixation cross was then replaced with the digit stimulus for the trial, also displayed for .25 s, which was followed by a blank display. The participant could provide a response from the time the fixation cross was presented. The blank display was shown for .9 s (fixed ISI condition) or for .6 s plus a random time sampled from a uniform distribution with minimum 0 and maximum .6 s (random ISI condition). The random and fixed ISI conditions therefore had equal mean ISI though considerable variability in the random condition. Once the ISI period elapsed, the next trial commenced with the display of the fixation cross.

7.2.2. Thought probes

Thought probes were presented as in Experiment 1 except that responses were given on a 5-point Likert scale with response options 1–5 where 1 was labeled as “off-task” and 5 was labeled as “on-task”. For consistency with Experiment 1, we report reverse-scored thought probe responses below (i.e., 1 indicates “on-task” and 5 indicates “off-task”). Participants responded to 20 thought probes in each condition.

7.3. RRM parameter estimation

Parameters of the RRM were estimated as described in Experiment 1 with two exceptions. Firstly, to examine whether the random ISI manipulation led to a faster leading edge of the RT distribution, Experiment 2 gained more sensitivity to very fast responses by permitting participants to respond prior to stimulus onset (i.e., negative response times with respect to stimulus onset). To model these data we added to all observed RTs the duration of the fixation cross, which was displayed for .25s immediately prior to stimulus

onset, ensuring the data conform to the strictly positive support of the RRM. This means that all modeled RTs between 0s and .25s were responses given prior to stimulus onset. To allow for the possibility of such pre-stimulus responses we fixed the shift parameter of the stimulus-unrelated response process to 0. Similarly, to prevent the stimulus-driven LBA from giving a response prior to stimulus onset, we added .25s to the estimated non-decision time. Secondly, in Experiment 2 we assumed separate location (μ) and scale (σ) parameters for the random and fixed ISI conditions for all RRM parameters; effectively, parameter estimation for the two conditions was conducted separately. We took 5,000 posterior samples from each of 40 MCMC chains with a burnin period of 3,000 samples. During the first 2,000 samples of the burnin period we conducted the migration update step for participant-level parameters with probability .05.

8. Results

The SART RT distributions observed in Experiment 2 precisely followed the same three trends reported in Experiment 1 (compare Figs. 6 and 8). SART RTs preceding self-reported off-task versus on-task responding were also faster, on average, and more variable, with a similar magnitude as in Experiment 1 (cf. Fig. 1).

8.1. RRM goodness of fit to data

The RRM provides an excellent fit to all aspects of the individual participant data – see Fig. 8. The model again provided a precise quantitative account of the proportion of go responses in go and nogo trials. It also captured the patterns of very fast responses, shallow leading edge of the RT distributions, and the distribution of error (nogo) RTs that were very fast relative to correct (go) RTs.

The individual participant data suggests that the random ISI condition might have produced a faster leading edge of the distribution than the fixed ISI condition. These differences can be better seen in the group-averaged data shown in Fig. 9. For all percentiles of the RT distribution faster than the median RT, the random ISI condition always responded faster on average than the fixed ISI condition. The effect is clearest in the quantile-quantile plots shown in the lower half of Fig. 9: for the fastest responses, the data (dots) and model (lines) are always below the identity line. Interestingly, however, once the fastest responses have passed, the RT distributions in the random and fixed ISI conditions begin to align for both go and nogo responses to the point where they are almost identical for moderate and slower responses. The observation that the very fastest RTs differ between conditions, but the slowest responses do not, suggests that the rhythmic response process was more likely to be the generator of the observed differences than the decision process (i.e., where the stimulus-unrelated process terminated prior to the stimulus-related process).

8.2. RRM parameter estimates

The posterior distribution of the parameters followed very similar trends to those reported in Experiment 1 – see Table 4. The primary comparison of interest in Experiment 2 was whether one or both of the parameters of the stimulus-unrelated rhythmic response process differed between the random and fixed ISI conditions while none of the decision-related parameters differed.

The hypothesis was borne out in data. The shape parameter took on a lower value in the random ISI condition relative to the fixed ISI condition (95% HDI of the difference distribution [.04, 2.15]); nevertheless, the shape parameter was greater than 1 on average for both conditions, indicating a predicted mode greater than 0. There was no corresponding difference in the scale parameter (λ ; 95% HDI [–.25, .08]). This implies that the latent rhythmic process had greater variance (lower precision) for participants in the random ISI condition: relative to the fixed ISI condition, the latent rhythmic process was both slower and faster across trials. Since slower rhythmic processes were likely to be censored by the decision process, but the fastest processes were not, the observed behavioral result was a faster leading edge to the RT distribution. This can be seen in the upper panel of Fig. 9 where the fastest responses for the random ISI condition are shifted leftward along the x-axis relative to the fixed ISI condition. The greater proportion of fast responses gave rise to the higher probability of a commission error (go response on nogo trials), shown in Fig. 9. Interestingly, unlike Experiment 1, all participants had values of the k considerably greater than 1, which reflects the finding that no participant RT distributions were peaked at 0 (cf. participants 2 and 13 of Experiment 1, Fig. 2). We believe this occurred because Experiment 2 allowed participants to make responses prior to stimulus onset – negative RTs – whereas Experiment 1 did not, so these very rapid responses were spread around rather than massed at the time of stimulus onset.

Importantly, there were no differences between the fixed and random ISI conditions for any of the decision-related parameters (all 95% HDIs included 0). This latter result provides support for the divergent validity of the model; a manipulation that was independent of the decision process did not influence parameter estimates of the stimulus-related decision process.

8.3. Self-reported mind wandering and the RRM

Mean rates of self-reported mind wandering were similar across conditions ($M_{\text{random}} = 2.58$ vs $M_{\text{fixed}} = 2.70$, $BF_{10} = .39$), and so too was the RRM-estimated probability of a stimulus-related decision ($M_{\text{random}} = .753$ vs $M_{\text{fixed}} = .759$, $BF_{10} = .37$). We then collapsed across the two groups and assessed whether the (non-parametric) correlations between self-reported mind wandering and RRM metrics generalized from Experiment 1 to Experiment 2, using two different approaches. The ‘independent’ approach used the same methods as Experiment 1, which is a strict criterion for replication as it makes the assumption that the first test never occurred (i.e., the knowledge gained in Experiment 1 does not contribute to the replication test in any statistical sense; we refer to this as $BF_{10, \text{ind.}}$). The ‘replication’ approach assumed a framework of cumulative science where the effect in Experiment 2 was evaluated with respect

Fig. 8. Defective CDFs showing data and the posterior predictive distribution of the rhythmic race model for individual participant data from Experiment 2. All details are as described in Fig. 6 with the exception that data and posterior predictive distributions for the random and fixed ISI conditions are shown in the left and right columns, respectively.

Fig. 9. Defective CDFs showing group-averaged data with overlaid group-averaged posterior predictive distribution of the rhythmic race model from Experiment 2 (upper panel), and quantile-quantile plots showing the difference in RT distributions for go and nogo responses between the random and fixed ISI conditions (lower panels). Upper panel: Go and nogo RTs are shown in dashed and solid lines, respectively. Quantile averaged data and posterior predictive distribution are shown in red shades (random ISI) and blue shades (fixed ISI), respectively. The mean proportion of go (i.e., observed) responses on go and nogo trials, $p(\text{go}|\text{go}, D)$ and $p(\text{go}|\text{nogo}, D)$, respectively, and the same quantities as predicted by the model, $p(\text{go}|\text{go}, M)$ and $p(\text{go}|\text{nogo}, M)$, are shown for the random and fixed ISI conditions. Lower panels: Percentiles of the RT distributions from the random and fixed ISI conditions are displayed as a quantile-quantile plot for go responses (left) and nogo responses (right). Data and posterior predictive distribution are shown in black dots and green lines, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

to the information that was gained about the effect in Experiment 1, which can be efficiently computed with the so-called replication Bayes factor ($BF_{10, \text{rep}}$; Ly, Etz, Marsman, & Wagenmakers, in press). This second approach is completely Bayesian in the sense that the posterior distribution over the parameters in Experiment 1 forms the prior distribution over the parameters for Experiment 2 (“today’s posterior is tomorrow’s prior”, Lindley, 1970).

Table 4
Parameter estimates of the Rhythmic Race Model (RRM) in Experiment 2.

Latent process	Parameter	ISI Condition	Median	HDI (lower, upper)
Stimulus-related decision process	$d_{go go}$	Random	3.63	3.29, 3.98
		Fixed	3.8	3.38, 4.31
	$d_{go nogo}$	Random	3.83	3.44, 4.27
		Fixed	3.97	3.33, 4.68
	$d_{nogo nogo}$	Random	4.35	4.03, 4.70
		Fixed	4.55	3.94, 5.21
	t_0	Random	.07	.06, .09
		Fixed	.09	.07, .11
	A	Random	.18	.08, .40
		Fixed	.28	.11, .68
	s	Random	1.04	.86, 1.23
		Fixed	1.11	.93, 1.33
Stimulus-unrelated rhythmic response process	k	Random	3.67	3.25, 4.11
		Fixed	4.75	3.88, 5.83
	λ	Random	.89	.77, 1.03
		Fixed	.81	.71, .92

Participants with greater mean $p(\text{decision})$ across all trials tended to have lower mean rates of self-reported mind wandering, though the association was weaker than in Experiment 1 ($r = -.15$) and provided ambiguous support, tending toward the null hypothesis ($BF_{10,ind.} = .53$, $BF_{10,rep.} = .82$). As in Experiment 1, there were stronger associations between mean self-reported mind wandering and decision-to-decision variability in drift rate (s ; $r = .38$) and mean processing speed to targets ($d_{nogo|nogo}$; $r = -.30$), although the evidence about these associations was equivocal ($BF_{10,ind.} = 1.8$, $BF_{10,rep.} = 1.7$; and $BF_{10} = 1.0$, $BF_{10,rep.} = .26$, respectively). The direction and relative magnitudes of the associations reported in Experiment 1 were consistent with those observed in Experiment 2 though did not lead to convincing levels of evidence, even when evaluated with respect to the results of Experiment 1, so should be treated with caution until tested in more highly-powered studies.

We also examined the probability of a stimulus-related decision in the responses preceding each thought probe response (cf. Fig. 7). These results followed similar, though more variable, trends as Experiment 1. Due to the relatively low participant numbers in each condition, it was not feasible to further analyze these data.

9. General discussion

We developed and tested an integrated cognitive process model of performance in the presence of mind wandering. The model assumes that decision behavior is generated by a competitive race between two latent processes: a stimulus-related decision process and a stimulus-unrelated, unconditionally triggered rhythmic response process. The rhythmic race model, or RRM, provided an excellent account of all stationary features of individual participant performance in two SART experiments, including choice proportions and response time distributions; the RRM is the first model to achieve this level of quantitative precision in the study of mind wandering.

The RRM also provided novel insights into task performance in the presence of mind wandering. In particular, the RRM's parameters, which were estimated without any reference to self report measures, were meaningfully associated with self-reported task-unrelated thoughts: participants predicted by the model to make more variable decisions were more likely to report lower levels of task-related attention throughout the task; this relationship only emerged when the cognitive model was endowed with the competing rhythmic process. Furthermore, the RRM's architecture allows us to infer whether individual trials are more likely to be on-task or off-task; these predictions were also corroborated against participants' self reports, though the association requires further validation in future studies. Finally, from the architecture of the RRM we generated the novel hypothesis that randomizing (vs. fixing) the inter-stimulus interval would increase the variability of the latent rhythmic component, which has the behavioral consequence of generating even faster responses. This effect was observed in data. An alternative interpretation is that randomizing (vs. fixing) the time between successive trials can actually increase the extremity of anticipatory responses, which is inconsistent with common lore surrounding randomizing trial timings. Taken together, the RRM provides a comprehensive and theoretically coherent account of task performance in the presence of mind wandering.

9.1. Mind wandering and the generation of a rhythmic race

We have provided a new model for understanding performance when people might be distracted from their ongoing task, or engaging in mind wandering. Such variations in performance as a function of task-unrelated thoughts could have many possible generators, which has led to a range of accounts from different theoretical proposals including perceptual decoupling (Schooler et al., 2011; Smallwood & Schooler, 2015), executive resources (Smallwood & Schooler, 2006; Teasdale et al., 1995), and executive control (McVay & Kane, 2009, 2012). We argue that deeper insight into the underlying causes of mind wandering will be gained once cognitive theories of mind wandering are implemented and tested in quantitative frameworks (for an excellent introduction to this

argument, see [Lewandowsky & Farrell, 2011](#)). We see our ‘bottom-up’ development of the RRM as the first step toward this larger goal of developing a quantitative theory of mind wandering. In this broader theory, existing qualitative theories of mind wandering could be implemented as quantitative ‘front-end’ components – where task-related and task-unrelated thoughts are generated – that drive an RRM ‘back-end’ component – where latent mind wandering tendencies are transformed into model parameters that generate overt responses. Seen from this perspective, the RRM serves our aim of developing a cognitive model of performance in the presence of mind wandering, rather than a cognitive model of mind wandering per se. We believe the RRM represents the start of what is hopefully a fruitful new direction for the field.

We see a number of ways in which the field can start moving toward this broader goal, by linking across current disparate theories. For example, although it was not framed as such, there already exists a quantitative implementation of the executive control theory of mind wandering ([Van Vugt et al., 2015](#)); indeed, [Van Vugt et al. \(2015\)](#) and [Peebles and Bothell \(2004\)](#) both proposed models based on competition between latent processes, similar in spirit to the RRM. [Van Vugt et al.’s \(2015\)](#) ACT-R-based model provides an interesting proposal about the generation of task-unrelated thoughts – sequential retrieval of items from declarative memory – and subsequent return to task-related thoughts – when a retrieved item reminds the system to attend to the ongoing task. A shortcoming of this model, however, is its limited quantitative account of behavioral data, even though it encodes similar assumptions of competition between latent components as the RRM. By re-considering the response output process of the ACT-R model, one could take [Van Vugt et al.’s \(2015\)](#) model as a ‘front-end’ and our RRM as a ‘back-end’ to form a fully quantitative theory of mind wandering, which can be tested against fine-grained features of the data. This approach combines the strength of two traditionally separated fields of cognitive modeling: ACT-R and its excellent account of higher-level features of a person operating in their environment, and evidence accumulation models and their precise account of detailed features of performance. There is a precedent to such approaches, too; for example, in understanding task learning and retrieval dynamics (e.g., [Van Maanen, Van Rijn, & Taatgen, 2012](#)).

In addition to informing cognitive theories of mind wandering, the precise account of behavioral data provided by the RRM provides a new avenue for investigating the neural origins of mind wandering. The neural causes and correlates of mind wandering have been of great interest (e.g., [Christoff et al., 2009](#); [Christoff, Irving, Fox, Spreng, & Andrews-Hanna, 2016](#); [Mittner et al., 2016](#); [Weissman et al., 2006](#)). For the most part, however, the field has been lacking a coherent theoretical explanation that quantitatively links neural recordings to observed behavior (for an exception, see [Mittner et al., 2014](#)). The RRM has the potential to be highly informative in this line of research. For example, it could provide a ‘back-end’ behavioral theory that can aid discrimination between discrete state- and continuous dimension-based theories of mind wandering, as discussed earlier. Furthermore, with a precise quantitative account of behavior, we can test clear hypotheses about the relationship between specific neural measures of mind wandering (e.g., activity in the default mode network) and latent components of cognitive processing (i.e., parameters of the cognitive model, such as the relative speed of the stimulus-unrelated process); we can also investigate the reverse direction – using parameters derived from the RRM as covariates in neural analyses. These approaches will lead to a much better understanding of the relationship between neural functioning and behavioral performance in the presence of mind wandering.

Our proposal has links to the interval timing literature, where the canonical task asks participants to repeatedly reproduce experimenter-specified intervals with button presses. We believe that the similarity between repeatedly reproducing externally-specified intervals and our latent process that generates a response at a stereotypical interval following onset of the fixation cross provides an interesting avenue to pursue links between literatures. It also provides an avenue to consider potential neural implementations of the rhythmic response process, since neurally-based theories of interval timing as an accumulation process have received substantial attention ([Simen et al., 2011](#); [Simen et al., 2016](#)).

9.2. The rhythmic race model in a broader context

We emphasize that even though we tested the RRM in the LBA framework, its use is not restricted to this evidence accumulation model. To the contrary, the RRM’s stimulus-related decision process could take the form of any evidence accumulation model that provides an adequate account of ‘pure’ decision behavior (i.e., performance in the absence of mind wandering), such as the diffusion decision model, leaky competing accumulator, poisson counter, or any other combination of independent or interacting accumulator or random walk/diffusion models. [Appendix C](#) outlines how to make such a change. Our hope is that the primary contribution of this manuscript is the proposal that performance is generated by a competitive race between a stimulus-related process and a stimulus-unrelated process. We are not committed to the parametric form of the stimulus-related process, and wish to sidestep debate on this issue; one can substitute their preferred generative evidence accumulation process into the RRM.

In a similar vein, although we tested the RRM in the SART, its use is not restricted to this context. The model could be used to understand performance in any domain where evidence accumulation is assumed to be an underlying principle driving behavior, regardless of whether mind wandering is the focus of study. This speaks to the generality of our approach: even if one is not interested in studying mind wandering, the principle of the rhythmic race provides an integrated approach for dealing with outlier or contaminant data without implementing ad-hoc exclusion criteria; we reiterate that we did not exclude a single trial of data from any participant in either experiment. This point will not be lost on researchers familiar with evidence accumulation models: the response time distributions we observed were well outside the explanatory scope of conventional evidence accumulation models (cf. [Fig. 2](#)). With the addition of a simple, unconditionally triggered response process, the framework accounted for a range of challenging data. Furthermore, by incorporating a mechanism that accounts for data outside the scope of conventional evidence accumulation models, those outlying data do not bear undue influence on the parameter estimates of the stimulus-related process. In other words, the RRM can provide a purer estimate of the stimulus-related decision process.

Appendix A. Hierarchical Bayesian data analysis of Experiment 1 data

We focused on 4 outcome variables that are prominent in the SART literature, each of which are calculated in the trials leading up to a thought probe trial: mean response time (RT), coefficient of variation in response times (RTCV), number of omission errors, and number of commission errors. One challenge in analyzing these outcome variables as a function of another dependent variable – that is, probe responses – is that it leads to an uneven number of observations in each cell of the design (possible probe responses) across participants. For example, participant 1 might respond to thought probes with 1 (on-task) on almost all probe trials and respond with 2, 3 and 4 on just a single trial each, while participant 2 might evenly distribute their responses across the 4-points of the scale. In this example, participant 1's outcome variables for probe response 1 ought to contribute more heavily to the group mean than participant 2, but vice versa for probe responses 2–4. With such uneven numbers, there are shortcomings in conventional analysis techniques, including the potential to overweight the relative impact of participant cell means.

To avoid these issues, we developed a custom Bayesian analysis that appropriately accounted for the uneven cell numbers across participants. Specifically, we used a hierarchical Bayesian approach where we estimated the group-level mean and standard deviation for each outcome variable at each position on the thought probe scale (i.e., 1–4), conditioned on the participant-level estimates; this approach treats thought probe responses as a fixed effect. The hierarchical approach allowed us to share information across participants in a principled manner, which naturally accounts for some participants contributing more information to the group-level parameter estimates for each position on the thought probe scale than others.

To illustrate, we continue with the example above, assuming that we are considering mean RT in the SART trials preceding each probe trial. Participant 2 responded approximately the same number of times with each position on the thought probe scale, so their data contains an approximately equivalent amount of information for their estimate of mean RT for each position on the thought probe scale. In contrast, participant 1 almost always responded to probe trials with a response of 1, so their data contain a high degree of information for their estimate of mean RT for position 1 on the thought probe scale, but almost no information about positions 2–4 on the scale (a single trial each). For these latter values, participant 1's estimate of mean RT would closely follow the prior distribution for positions 2–4 on the thought probe scale, which in a hierarchical model is the group-level distribution for those parameters. The group-level distributions are influenced by the other participants' parameter estimates, which lead to shrinkage for participant 1's estimates toward the group-level distribution, as it should since there was next-to-no information in participant 1's responses for these positions on the thought probe scale.

A.1. Mean RT

Let y_{ij} be participant i 's mean RT over a window of w trials preceding their response to thought probe item j , denoted p_{ij} . The analysis was then defined as:

Data level:

$$y_{ij} \sim N\left(\alpha_{i,p_{ij}}, \epsilon_i\right)$$

Group level:

$$\alpha_{ik} \sim N(\alpha_k^\mu, \alpha^\sigma)_{(0,)} \\ \epsilon_i \sim N(\epsilon^\mu, \epsilon^\sigma)_{(0,)}$$

Prior distributions:

$$\alpha_k^\mu \sim N(.3, .1) \\ \epsilon^\mu \sim N(.1, .1) \\ \alpha^\sigma, \epsilon^\sigma \sim \Gamma(1, .1)$$

where i denotes participants, j denotes probe trials, and k denotes positions on the thought probe scale (i.e., 1–4). $N(a, b)$ refers to a Normal distribution with mean a and standard deviation b ; subscript $_{(0,)}$ refers to a truncated distribution with lower bound at 0; $\Gamma(c, d)$ refers to the Gamma distribution with shape c and scale d . The analysis was repeated for window sizes $w = 3, 4, \dots, 10$. Posterior distributions of the group-level mean parameters (α_k^μ) for selected window sizes are shown in the upper row of Fig. 1 of the main text.

A.2. RT Coefficient Of Variation (RTCV)

The RTCV analysis was identical to the mean RT analysis except that y_{ij} denoted the RTCV from participant i over the window of w trials preceding thought probe item j , and we changed the group-level prior distribution for the mean parameter to $\alpha_k^\mu \sim N(.2, .1)$ to account for the smaller scale of the outcome variable. The results of this analysis are shown in the second row of Fig. 1 of the main text.

A.3. Probability of omission and commission errors

The choice response analyses were similar to the mean RT and RTCV analyses except that y_{ij} now denotes the number of omission/commission errors from participant i that occurred within the previous n_{ij} trials. For the analysis of omission errors, n_{ij} denotes the number of non-target (i.e., go) trials that occurred within the w trials preceding probe p_{ij} . For the analysis of commission errors, n_{ij} denotes the number of target (i.e., nogo) trials that occurred within the w trials preceding probe p_{ij} . We assume y_{ij} follows a binomial distribution with rate parameter θ_{ik} , and counts n_{ij} :

Data level:

$$y_{ij} \sim \text{Binomial}(\theta_{i,p_{ij}}, n_{ij})$$

Group level:

$$\theta_{ik} \sim \text{Beta}(\alpha_k, \beta_k)$$

$$\alpha_k = \theta_k^\mu \cdot \lambda$$

$$\beta_k = (1 - \theta_k^\mu) \cdot \lambda$$

Prior distributions:

$$\theta_k^\mu \sim \text{Beta}(1, 1)$$

$$\lambda \sim \Gamma(1, .01)$$

where i , j and k again denote participants, probe trials and positions on the thought probe scale, respectively. $\text{Binomial}(a, b)$ refers to a Binomial distribution with probability of success a (i.e., an omission/commission error) from b trials; $\text{Beta}(c, d)$ refers to a Beta distribution with shape parameters c and d . For simpler interpretation, we use the mean-precision parameterization of the Beta distribution: θ_k^μ represents the group-level mean of each of k Beta distributions, and λ represents the precision of this estimate (specifically, λ represents the estimated number of observations informing our belief about θ_k^μ). For simplicity, we assume a common λ across the k positions on the thought probe scale, which mirrors the assumption of a common group-level standard deviation (σ^2) in the analysis of mean RT and RTCV. Posterior distributions of the group-level probability of omission and commission errors (θ_k^μ) for selected window sizes are shown in the third and fourth rows of Fig. 1 of the main text.

Appendix B. Hierarchical Bayesian data analysis of $p(\text{decision})$ in Experiment 1

We analyzed the probability of a response generated from the stimulus-related decision process, relative to the stimulus-unrelated rhythmic response process, $p(\text{decision})$, using an almost identical hierarchical Bayesian analysis as described in Appendix A.

Data level:

$$y_{ij} \sim \text{Beta}(\alpha_{i,p_{ij}}, \beta_{i,p_{ij}})$$

Group level:

$$\alpha_{ik} = \theta_{ik} \cdot \lambda_i$$

$$\beta_{ik} = (1 - \theta_{ik}) \cdot \lambda_i$$

$$\theta_{ik} \sim \text{Beta}(\alpha_k^\mu, \beta_k^\mu)$$

$$\lambda_i \sim \Gamma(\text{shape}, \text{rate})$$

Prior distributions:

$$\alpha_k^\mu, \beta_k^\mu \sim \mathcal{U}(0, 1000)$$

$$\text{shape}, \text{rate} \sim \mathcal{U}(0, 100)$$

$$\theta_k^\mu = \frac{\alpha_k^\mu}{\alpha_k^\mu + \beta_k^\mu}$$

Appendix C. Diffusion Decision Model (DDM) specification of the Rhythmic Race Model (RRM)

The RRM assumes a race between a stimulus-related decision process and a stimulus-unrelated rhythmic responding process. In this appendix we describe an RRM where the stimulus-related decision process is specified as a DDM, in contrast to the Linear Ballistic Accumulator (LBA) implementation from the main text. In this RRM, a go response can be generated in two ways:

1. When the decision variable crosses the go boundary and the rhythmic runner is slower than the time taken to cross the go boundary.
2. The rhythmic runner is faster than the time taken for the decision variable to cross the go or nogo boundaries.

To specify the RRM architecture we denote the probability density function (PDF) that the decision variable crossed the *go* or *nogo* boundaries at time t , without having previously crossed either boundary, as $f_{go}(t)$ and $f_{nogo}(t)$, respectively, with corresponding cumulative distribution functions (CDF) $F_{go}(t)$ and $F_{nogo}(t)$. Similarly, we denote the PDF and CDF of the rhythmic runner producing a response at time t as $f_{rhythmic}(t)$ and $F_{rhythmic}(t)$, respectively. The PDF of a *go* response at time t is therefore

$$PDF_{go}(t) = f_{go}(t) \times (1 - F_{rhythmic}(t)) + f_{rhythmic}(t) \times (1 - F_{go}(t) - F_{nogo}(t)).$$

A *nogo* response only occurs if the rhythmic runner is slower than the time taken for the decision variable to cross the *nogo* boundary. Since there is no observed response when the *nogo* boundary is crossed (i.e., it is an implicit boundary that corresponds to withholding a response), we integrate over the predicted distribution of response times for a *nogo* response. The probability of a *nogo* response is therefore

$$p(nogo) = \int_{t=0}^{\infty} f_{nogo}(t) \times (1 - F_{rhythmic}(t)) dt.$$

Appendix D. Parameter estimation for the Rhythmic Race Model (RRM)

We used a hierarchical Bayesian framework to simultaneously estimate the RRM parameters at the participant and group levels. Parameters $\theta_i \in \{d_{go|go}, d_{go|nogo}, d_{nogo|nogo}, A, s, t0, k, \lambda\}$ for participant i , where $i \in \{1, 2, \dots, N\}$, were transformed to the real line and drawn hierarchically from Normally-distributed group-level hyper distributions with mean θ^μ and standard deviation θ^σ , for $\theta = (\theta_1, \dots, \theta_n)$ where $n = 8$ was the number of model parameters at the participant (i.e., data) level. We assumed mildly informative prior distributions for the group-level parameters, which allowed the estimated parameters to vary across a range of plausible values. Formally, the RRM was defined as:

$$\begin{aligned} &\text{Data level:} \\ &(RT_i, response_i) \sim \text{RRM}(\theta_i) \\ &\text{Group level:} \\ &\log(\theta_i) \sim N(\theta^\mu, \theta^\sigma), \quad \forall \theta_i \in \{\theta_j\} \setminus \{A_i\} \\ &\text{logit}(A_i) \sim N(A^\mu, A^\sigma) \end{aligned}$$

$$\begin{aligned} &\text{Prior distributions:} \\ &d_{go|go}^\mu, d_{go|nogo}^\mu, d_{nogo|nogo}^\mu \sim N(\log(4), \log(2)) \\ &A^\mu \sim N(0, 1.3) \\ &s^\mu \sim N(\log(.5), .5) \\ &t0^\mu \sim N(\log(.2), .5) \\ &k^\mu, \lambda^\mu \sim N(0, 1) \\ &d_{go|go}^\sigma, d_{go|nogo}^\sigma, d_{nogo|nogo}^\sigma, A^\sigma, k^\sigma, \lambda^\sigma \sim \Gamma(1, 1) \\ &s^\sigma, t0^\sigma \sim \Gamma(1, .5) \end{aligned}$$

where $N(a, b)$ denotes the Normal distribution with mean a and standard deviation b and $\Gamma(c, d)$ denotes the Gamma distribution with shape c and scale d . In addition to these assumptions, the Mixture Uniform and Mixture Weibull models freely estimated the probability of a contaminant response (p) at the participant and group levels; specifically, $\text{logit}(p_i) \sim N(p^\mu, p^\sigma)$, where $p^\mu \sim N(-1.5, 1.5)$ and $p^\sigma \sim \Gamma(1, 1)$. In the response bias model, we set the response threshold for the *go* accumulator as a scaling parameter of the model, $b_{go} = 1$, and freely estimated the response threshold of the *nogo* accumulator; $\log(b_{nogo}) \sim N(b_{nogo}^\mu, b_{nogo}^\sigma)$, where $b_{nogo}^\mu \sim N(-.5, 1)$ and $b_{nogo}^\sigma \sim \Gamma(1, 1)$.

Parameters were estimated by differential evolution Markov chain Monte Carlo (DE-MCMC; Turner, Sederberg, Brown, & Steyvers, 2013) sampling, using the default settings unless noted otherwise (see Turner et al., 2013). We took 5000 posterior samples from each of 30 MCMC chains with a burnin period of 3000 samples. During the first 1500 samples of the burnin period we conducted a “migration” update step for participant-level parameters with probability .05, which helps to remove chains stuck in low likelihood regions of the parameter space (for details, see Turner et al., 2013). Convergence was monitored through visual inspection and the multivariate potential scale reduction factor ($\hat{R}$; Brooks & Gelman, 1998).

Appendix E. Parameter recovery in the Rhythmic Race Model (RRM)

We confirmed that the parameters of the RRM can be reliably recovered from data in a parameter recovery study modeled on the structure of Experiment 1. We first randomly sampled a parameter vector from each participant’s posterior distribution of the RRM parameters estimated from the real data, and used these parameter vectors to generate a synthetic data set of the same size as the real data in terms of the number of trials per participant (640 *go* trials, 80 *nogo* trials) and the number of participants (19). We then estimated the RRM from these simulated data using identical methods to those applied to the real data (Appendix D). We independently repeated the procedure 25 times. This produced 25 independent hierarchical estimation exercises that assessed parameter recovery from a total of 475 participant-level posterior distributions, where the focus of the recoverability analysis was at the participant level.

Fig. E.1. Parameter recovery for the RRM. Model parameters are shown in separate panels, data-generating parameter values are shown along the x-axes, and the median of the estimated posterior distributions are shown on the y-axes. Perfect parameter recovery is indicated along the diagonal. Parameters are shown in the transformed space in which they were estimated (logit scaling for A, log scaling for all other parameters; [Appendix D](#)).

[Fig. E.1](#) shows the results of the parameter recovery study. Each panel shows a separate parameter with the data-generating values shown on the x-axis and the median of the estimated participant-level posterior distribution on the y-axis. Overall, all parameters of the model were well recovered, shown by the points largely falling closely to the identity line. These results demonstrate that the parameters of the RRM can be recovered in typically-sized data sets.

Appendix F. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.cogpsych.2019.05.002>.

References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Bastian, M., & Sackur, J. (2013). Mind wandering at the fingertips: Automatic parsing of subjective states based on response time variability. *Frontiers in Psychology*, 4, n/a. <https://doi.org/10.3389/fpsyg.2013.00573>.
- Boayue, N. M., Csifcsak, G., Aslaksen, P., Turi, Z., Antal, A., Groot, J., ... Mittner, M. (2019). Increasing propensity of mind wandering by transcranial direct current stimulation: A registered report. *European Journal of Neuroscience* in press.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review*, 113, 700–765.
- Brooks, S. P., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7, 434–455.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology*, 57, 153–178.
- Brown, K. W., & Ryan, R. M. (2003). The benefits of being present: Mindfulness and its role in psychological well-being. *Journal of Personality and Social Psychology*, 84, 822–848.
- Bussemeyer, 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.
- Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., & Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, 14, 1462–1467.
- Cheyne, J. A., Solman, G. J., Carriere, J. S., & Smilek, D. (2009). Anatomy of an error: A bidirectional state model of task engagement/disengagement and attention-related errors. *Cognition*, 111, 98–113.
- Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 8719–8724.
- Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews Neuroscience*, 17, 718–731.
- Dang, J. S., Figueroa, I. J., & Helton, W. S. (2018). You are measuring the decision to be fast, not inattention: The sustained attention to response task does not measure sustained attention. *Experimental Brain Research*, 236, 2255–2262.
- Donkin, C., Brown, S., & Heathcote, A. (2009). The over-constraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, 16, 1129–1135.
- Duncan, J., Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: The organization of goal-directed behavior. *Cognitive Psychology*, 30, 257–303.
- Esterman, M., Noonan, S. K., Rosenberg, M., & Degutis, J. (2013). In the zone or zoning out? Tracking behavioral and neural fluctuations during sustained attention. *Cerebral Cortex*, 23, 2712–2723.

- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, 53, 134–140.
- Estes, W. K., & Maddox, W. T. (2005). Risks of drawing inferences about cognitive processes from model fits to individual versus average performance. *Psychonomic Bulletin & Review*, 12, 403–408.
- Evans, N. J., & Hawkins, G. E. (2019). When humans behave like monkeys: Feedback delays and extensive practice increase the efficiency of speeded decisions. *Cognition*, 184, 11–18.
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Science*, 105, 17538–17542.
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, 67, 641–666.
- Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *The Journal of Neuroscience*, 35, 485–494.
- Giambra, L. M. (1995). A laboratory method for investigating influences on switching attention to task-unrelated imagery and thought. *Consciousness and Cognition*, 4, 1–21.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574.
- Gomez, P., Ratcliff, R., & Perea, M. (2007). A model of the go/no-go task. *Journal of Experimental Psychology: General*, 136, 389–413.
- Gouraud, J., Delorme, A., & Berberian, B. (2017). Autopilot, mind wandering, and the out of the loop performance problem. *Frontiers in Neuroscience*, 11, n/a. <https://doi.org/10.3389/fnins.2017.00541>.
- Gunzelmann, G., Gross, J. B., Gluck, K. A., & Dinges, D. F. (2009). Sleep deprivation and sustained attention performance: Integrating mathematical and cognitive modeling. *Cognitive Science*, 33, 880–910.
- Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2012). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic Bulletin & Review*, 19, 339–348.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35, 2476–2484.
- Hawkins, G. E., Marley, A. A. J., Heathcote, A., Flynn, T. N., Louviere, J. J., & Brown, S. D. (2014). Integrating cognitive process and descriptive models of attitudes and preferences. *Cognitive Science*, 38, 701–735.
- Hawkins, G. E., Mittner, M., Boekel, W., Heathcote, A., & Forstmann, B. U. (2015). Toward a model-based cognitive neuroscience of mind wandering. *Neuroscience*, 310, 290–305.
- Hawkins, G. E., Mittner, M., Forstmann, B. U., & Heathcote, A. (2017). On the efficiency of neurally-informed cognitive models to identify latent cognitive states. *Journal of Mathematical Psychology*, 76, 142–155.
- Heathcote, A. & Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in Psychology*, 3, n/a. doi:<https://doi.org/10.3389/fpsyg.2012.00292>.
- Heathcote, A., Suraev, A., Curley, S., Gong, Q., & Love, J. (2015). Decision processes and the slowing of simple choices in schizophrenia. *Journal of Abnormal Psychology*, 124, 961–974.
- Heitz, R. P. (2014). The speed-accuracy tradeoff: History, physiology, methodology, behavior. *Frontiers in Neuroscience*, 8, n/a. <https://doi.org/10.3389/fpsyg.2014.00150>.
- Hiatt, L. M., & Traflet, J. G. (2015). A computational model of mind wandering. In D. C. Noelle, (Ed.). *Proceedings of the 37th annual conference of the cognitive science society*. Cognitive Science Society.
- Killingsworth, M. A., & Gilbert, D. T. (2010). A wandering mind is an unhappy mind. *Science*, 330, 932.
- Kruschke, J. (2011). *Doing Bayesian data analysis: A tutorial with R and BUGS*. Oxford: Academic Press.
- Lee, M. D., & Webb, M. R. (2005). Modeling individual differences in cognition. *Psychonomic Bulletin & Review*, 12, 605–621.
- Leszczynski, M., Chaeib, L., Reber, T. P., Derner, M., Axmacher, N., & Fell, J. (2017). Mind wandering simultaneously prolongs reactions and promotes creative incubation. *Scientific Reports*, 7, n/a. <https://doi.org/10.1038/s41598-017-10616-3>.
- Lewandowsky, S., & Farrell, S. (2011). *Computational modeling in cognition*. California: Sage.
- Lindley, D. V. (1970). *Bayesian statistics: A review*. Society for Industrial and Applied Mathematics Philadelphia.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492–527.
- Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, 121, 66–95.
- Luce, R. D. (1986). *Response times*. New York: Oxford University Press.
- Ly, A., Etz, A., Marsman, M., & Wagenmakers, E.-J. (2019). Replication bayes factors from evidence updating. *Behavior Research Methods* in press.
- Manly, T., Davison, B., Heutink, J., Galloway, M., & Robertson, I. H. (2000). Not enough time or not enough attention? Speed, error and self-maintained control in the Sustained Attention to Response Test (SART). *Clinical Neuropsychological Assessment*, 3, 167–177.
- Manly, T., Robertson, I. H., Galloway, M., & Hawkins, K. (1999). The absent mind: Further investigations of sustained attention to response. *Neuropsychologia*, 37, 661–670.
- Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: The default network and stimulus-independent thought. *Science*, 315, 393–395.
- McVay, J. C., & Kane, M. J. (2009). Conducting the train of thought: Working memory capacity, goal neglect, and mind wandering in an executive-control task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 196–204.
- McVay, J. C., & Kane, M. J. (2010). Does mind wandering reflect executive function or executive failure? Comment on Smallwood and Schooler (2006) and Watkins (2008). *Psychological Bulletin*, 136, 188–207.
- McVay, J. C., & Kane, M. J. (2012). Drifting from slow to D'oh!: Working memory capacity and mind wandering predict extreme reaction times and executive control errors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 525–549.
- Mittner, M., Boekel, W., Tucker, A. M., Turner, B. M., Heathcote, A., & Forstmann, B. U. (2014). When the brain takes a break: A model-based analysis of mind wandering. *The Journal of Neuroscience*, 34, 16286–16295.
- Mittner, M., Hawkins, G. E., Boekel, W., & Forstmann, B. U. (2016). A neural model of mind wandering. *Trends in Cognitive Sciences*, 20, 570–578.
- Mooneyham, B. W., & Schooler, J. W. (2013). The costs and benefits of mind-wandering: A review. *Canadian Journal of Experimental Psychology*, 67, 11–18.
- Mrazek, M. D., Smallwood, J., & Schooler, J. W. (2012). Mindfulness and mind-wandering: Finding convergence through opposing constructs. *Emotion*, 12, AA2–AA8.
- Nunez, M. D., Srinivasan, R., & Vandekerckhove, J. (2015). Individual differences in attention influence perceptual decision making. *Frontiers in Psychology*, 8, n/a. <https://doi.org/10.3389/fpsyg.2015.00018>.
- Nunez, M. D., Vandekerckhove, J., & Srinivasan, R. (2017). How attention influences perceptual decision making: Single-trial EEG correlates of drift-diffusion model parameters. *Journal of Mathematical Psychology*, 76, 117–130.
- Pachella, R. G. (1974). The interpretation of reaction time in information-processing research. In B. H. Kantowitz (Ed.). *Human information processing: Tutorials in performance and cognition* (pp. 41–82). Hillsdale (NJ): Lawrence Erlbaum Associates.
- Peebles, D., & Bothell, D. (2004). Modelling performance in the Sustained Attention to Response Task. *Proceedings of the 6th international conference on cognitive modeling* (pp. 231–236). Mahwah, NJ: Lawrence Erlbaum.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2, n/a. <https://doi.org/10.3389/neuro.11.010.2008>.
- Pratte, M. S., Rouder, J. N., & Morey, R. D. (2010). Separating mnemonic process from participant and item effects in the assessment of ROC asymmetries. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 224–232.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59–108.
- Ratcliff, R., Huang-Pollock, C., & McKoon, G. (2018). Modeling individual differences in the go/no-go task with a diffusion model. *Decision*, 5, 42–62.

- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111, 333–367.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20, 260–281.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9, 438–481.
- Reed, A. V. (1974). Speed-accuracy trade-off in recognition memory. *Science*, 181, 574–576.
- Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (1997). 'Oops!': Performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia*, 35, 747–758.
- Schooler, J. W., Smallwood, J., Christoff, K., Handy, T. C., Reichle, E. D., & Sayette, M. A. (2011). Meta-awareness, perceptual decoupling and the wandering mind. *Trends in Cognitive Sciences*, 15, 319–326.
- Seli, P., Cheyne, J. A., & Smilek, D. (2012). Attention failures versus misplaced diligence: Separating attention lapses from speed-accuracy trade-offs. *Consciousness and Cognition*, 21, 277–291.
- Seli, P., Cheyne, J. A., & Smilek, D. (2013). Wandering minds and wavering rhythms: Linking mind wandering and behavioral variability. *Journal of Experimental Psychology: Human Perception and Performance*, 39, 1–5.
- Seli, P., Jonker, T. R., Cheyne, J. A., & Smilek, D. (2013). Enhancing SART validity by statistically controlling speed-accuracy trade-offs. *Frontiers in Psychology*, 4, n/a. <https://doi.org/10.3389/fpsyg.2013.00265>.
- Seli, P., Jonker, T. R., Solman, G. J., Cheyne, J. A., & Smilek, D. (2013). A methodological note on evaluating performance in a sustained-attention-to-response task. *Behavior Research Methods*, 45, 355–363.
- Simen, P., Balci, F., deSouza, L., Cohen, J. D., & Holmes, P. (2011). A model of interval timing by neural integration. *The Journal of Neuroscience*, 31, 9238–9253.
- Simen, P., Vlasov, K., & Papadakis, S. (2016). Scale (in)variance in a unified diffusion model of decision making and timing. *Psychological Review*, 123, 151–181.
- Singmann, H., Brown, S., Gretton, M., & Heathcote, A. (2016). *rtlists: Response time distributions* [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=rtlists> (R package version 0.6-6).
- Smallwood, J. (2010). Why the global availability of mind wandering necessitates resource competition: Reply to McVay and Kane (2010). *Psychological Bulletin*, 136, 202–207.
- Smallwood, J., Davies, J. B., Heim, D., Finnigan, F., Sudberry, M., O'Connor, R., & Obonsawin, M. (2004). Subjective experience and the attentional lapse: Task engagement and disengagement during sustained attention. *Consciousness and Cognition*, 13, 657–690.
- Smallwood, J., McSpadden, M., Luus, B., & Schooler, J. W. (2008). Segmenting the stream of consciousness: The psychological correlates of temporal structures in the time series data of a continuous performance task. *Brain and Cognition*, 66, 50–56.
- Smallwood, J., McSpadden, M., & Schooler, J. W. (2007). The lights are on but no one's home: Meta-awareness and the decoupling of attention when the mind wanders. *Psychonomic Bulletin & Review*, 14, 527–533.
- Smallwood, J., & Schooler, J. W. (2006). The restless mind. *Psychological Bulletin*, 132, 946–958.
- Smallwood, J., & Schooler, J. W. (2015). The science of mind wandering: Empirically navigating the stream of consciousness. *Annual Review of Psychology*, 66, 487–518.
- Smilek, D., Carriere, J. S., & Cheyne, J. A. (2010). Failures of sustained attention in life, lab, and brain: Ecological validity of the SART. *Neuropsychologia*, 48, 2564–2570.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society B*, 64, 583–639.
- Stawarczyk, D., Majerus, S., Maj, M., Van der Linden, A., & D'Argembeau, M. (2011). Mind-wandering: Phenomenology and function as assessed with a novel experience sampling method. *Acta Psychologica*, 136, 370–381.
- Swenson, R. G. (1972). The elusive tradeoff: Speed versus accuracy in visual discrimination tasks. *Perception & Psychophysics*, 12, 16–32.
- Teasdale, J. D., Dritschel, B. H., Taylor, M. J., Proctor, L., Lloyd, C. A., Nimmo-Smith, I., & Baddeley, A. D. (1995). Stimulus-independent thought depends on central executive resources. *Memory & Cognition*, 23, 551–559.
- Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, 18, 368–384.
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 108, 550–592.
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2008). A Bayesian approach to diffusion models of decision-making. In V. M. Sloutsky, B. C. Love, & K. McRae (Eds.). *Proceedings of the 30th annual conference of the cognitive science society* (pp. 1429–1434). Cognitive Science Society.
- Van Maanen, L., van Rijn, H., & Taatgen, N. A. (2012). RACE/A: An architectural account of the interactions between learning, task control, and retrieval dynamics. *Cognitive Science*, 36, 62–101.
- Van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2012). A diffusion model decomposition of the effects of alcohol on perceptual decision making. *Psychopharmacology*, 219, 1017–1025.
- Van Vugt, M., Taatgen, N., Sackur, J., & Bastian, M. (2015). Modeling mind-wandering: A tool to better understand distraction. In N. Taatgen, M. Van Vugt, J. Borst, & K. Mehlhorn (Eds.). *Proceedings of the 13th international conference on cognitive modeling* (pp. 252–257). The Netherlands: Groningen.
- Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review*, 7, 424–465.
- Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, 7, 208–256.
- Vickers, D. (1979). *Decision processes in visual perception*. London: Academic Press.
- Wagenmakers, E.-J., & Brown, S. (2007). On the linear relation between the mean and the standard deviation of a response time distribution. *Psychological Review*, 114, 830–841.
- Wald, A., & Wolfowitz, J. (1948). Optimal character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19, 326–339.
- Walsh, M. M., Gunzelmann, G., & Van Dongen, H. P. (2017). Computational cognitive modeling of the temporal dynamics of fatigue from sleep loss. *Psychonomic Bulletin & Review*, 24, 1785–1807.
- Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses of attention. *Nature Neuroscience*, 9, 971–978.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, 41, 67–85.