Winner takes all!

What are race models, and why and how should psychologists use them?

Andrew Heathcote[1,2] and Dora Matzke[2]

[1] School of Psychology, University of Newcastle, Australia
[2] Department of Psychology, University of Amsterdam

Contact Author:
Andrew Heathcote
andrew.heathcote@newcastle.edu.au
PO Box 15906, 1001 NK Amsterdam, The Netherlands

Interest in the processes that mediate between stimuli and responses is at the heart of most modern psychology and neuroscience. These quantities cannot be directly measured but instead must be inferred from observed responses. Race models, through their ability to account for both response choices and response times, have been a key enabler of such inferences. Examples appeared contemporaneously with the cognitive revolution, and since have become increasingly prominent and elaborated, so that psychologists now have a powerful array of race models at their disposal. We showcase the state-of-the-art for race models and describe why and how they are used.

Quantitative models facilitate the expression of psychological theories in a way that can be rigorously tested against data. This paper addresses an influential and widely-used type, *race models*. Figure 1 illustrates one of the most well-known examples, Logan and Cowan's (1984) "horse-race" model of the stop-signal task used to measure the ability to inhibit a response. The "horses" represent cognitive processes racing against each other; if the "go" runner wins, a response is made, and if the "stop" runner wins, it is withheld. The speed of the stop runner is the key measure of inhibitory ability, but definitionally that runner's finishing time is not observable. However, combining what can be observed with assumptions based on the horse-race model enables estimation of the stop runner's speed. Measuring the unobservable in this way exemplifies the role of race models in inferring psychological quantities that cannot be directly observed.

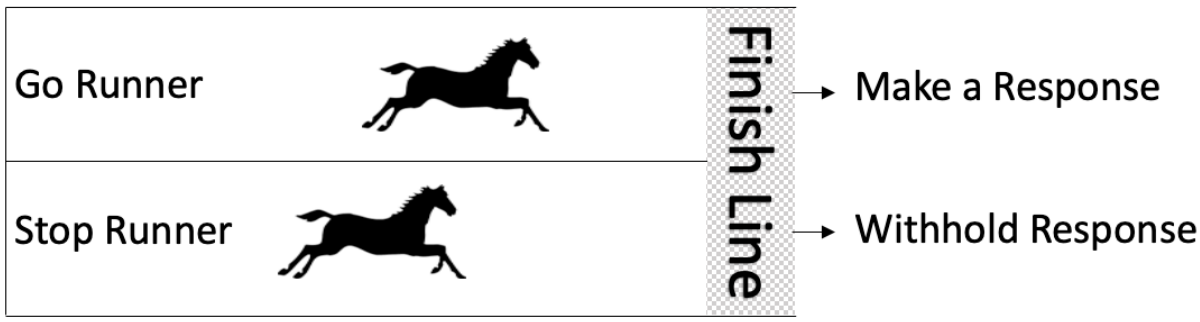

Figure 1. Logan and Cowan's (1984) horse-race model of response inhibition in the stop-signal task. On most trials ("go" trials) participants perform a choice task, making a response to a go stimulus, typically a binary choice, but on some trials, a stop signal (e.g., a tone) also occurs sometime after the go stimulus, indicating that they should withhold their response. Typically, the delay between the go and stop signals is varied, and effects on the probability of stopping and response times (RTs) when stopping fails, along with RTs when there is no stop signal, are observed. Combining the race model with these observables, along with some assumptions about the distribution of finishing times, provides an estimate of stop-signal RT, the time for the stop horse to run the race.

What are race models?

The defining characteristics of a race model are that: a) it contains one or more "runner" that take time to complete, b) where there is more than one runner, they may or may not interact in a way that affects their timing, and c) it acts according to a "winner-takes-all" rule that controls subsequent processing based on the runner, or set of runners, that finishes first. *Evidence-Accumulation Models* (EAMs) are the most widely-adopted special case of a race model, where the runners have a specific interpretation as processes that complete when they have accrued a threshold amount of evidence. We first examine several basic EAMs used to address simple decision tasks. We then describe how these basic EAMs can be used as building blocks to address a broader range of tasks and psychological processes.

Figure 2 illustrates three race models, two in which the runners are independent, the Linear Ballistic Accumulator (LBA) and the Racing Diffusion Model (RDM), and one which is a special case of the latter, the Wiener Diffusion Model (WDM), where the runners interact. In the LBA and RDM, the "horses" are evidence accumulators, and their positions are the current values of the sum of the evidence for the response that they represent. For example, a choice about whether a stimulus is moving left or right could have accumulators corresponding to left- and right-hand button-press responses, respectively, as shown in Figure 2. Accumulation rates vary from trial to trial but are on average higher for the accumulator that matches the stimulus than the accumulator that mismatches the stimulus.

In the LBA, accumulation occurs at a constant rate within a trial; in the RDM it is "diffusive", (i.e., varying randomly from moment to moment) with a constant average rate. In the LBA, the starting point of accumulation also varies from trial to trial. In both, the choice is determined by the first accumulator to reach a threshold. Response time (RT) is the sum of decision time (i.e., the time to move from the starting point to threshold) and non-decision time, typically the sum of the times to encode the choice stimulus and to produce a motor response once a winner is determined. The third column of Figure 2 shows the resulting RT distributions, with the area under each distribution representing the corresponding response probability.

The WDM is a special case of the RDM where moment-by-moment changes in the evidence in one accumulator correspond to equal and opposite changes in the other accumulator. Evidence for each response is reduced to a difference, and the race becomes a tug of war represented as one accumulator with two thresholds (right column, Figure 2). The basic WDM wrongly predicts that error and correct responses have identical RT distributions, so it has been augmented with the same types of trial-to-trial variability as in the LBA, in which case it has been called the Diffusion Decision Model (DDM; Ratcliff & McKoon, 2008). Both models remain limited to binary choices, but other diffusive race models like the RDM and variants with interactions among the runners, are not so limited.

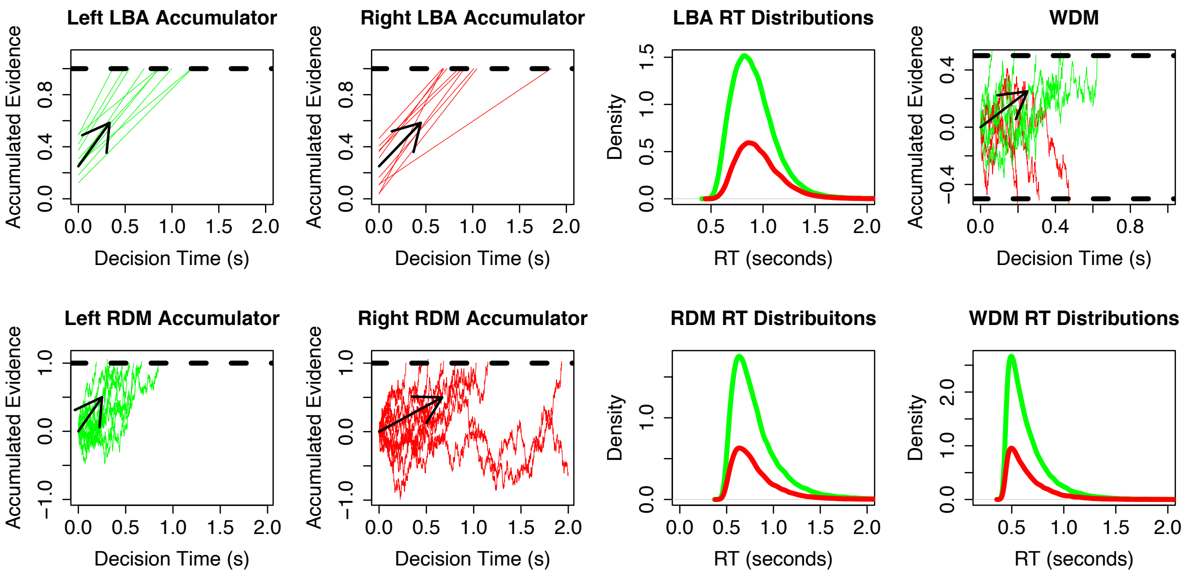

Figure 2. LBA, RDM, and WDM models for a choice between left and right response options where the left option is correct (accuracy ~73% in all cases), and corresponding response time (RT) distributions (s = seconds). Left evidence trajectories and RT distributions shown in green and right evidence trajectories/RT distributions in red. The LBA and WDM have one accumulator for each response option each with its own threshold, indicated by a dashed line. For the WDM, there is only one accumulator, but it has two thresholds: the upper dashed line corresponds to the left-response threshold and the lower dashed line to the right response threshold. In all cases, the RT distributions are depicted in their "defective" form where the area under the curve equals the probability of the corresponding response. For each accumulator, accumulation trajectories during 10 trials are plotted (straight lines in the LBA, jagged lines in the others), where arrows show the average accumulation rates. Each LBA accumulator has independent trial-to-trial variability in accumulation rates and starting points with Gaussian and uniform distributions, respectively. Both the RDM and WDM have normally distributed moment-to-moment variability in rates, which is independent between accumulators for the RDM. Parameters and code to plot this figure are given in an R script at https://osf.io/xzhe5. Figure available at https://tinyurl.com/52v7hwam under CC-BY 2.0 license (https://creativecommons.org/licenses/by/2.0/).

Our first recommended reading, Donkin and Brown (2018), provides a broad and inclusive review of most types of EAM. These vary along dimensions ranging from discrete vs. continuous evidence accrual to non-linearities due to "leakage" (i.e., evidence being lost from the total) and positive feedback. Bogacz et al. (2007) discuss motivations for EAMs in terms of optimality and their biological basis. Despite these differences, all applications interpret the core parameters similarly: rates are determined by arousal, attention, and stimulus characteristics, whereas thresholds are set to strategically control response caution and bias. Higher thresholds increase caution, slowing responding but increasing accuracy (i.e., the speed-accuracy trade-off) by overcoming start-point biases or averaging out diffusive variability. Bias occurs when the threshold for one response is lower than another.

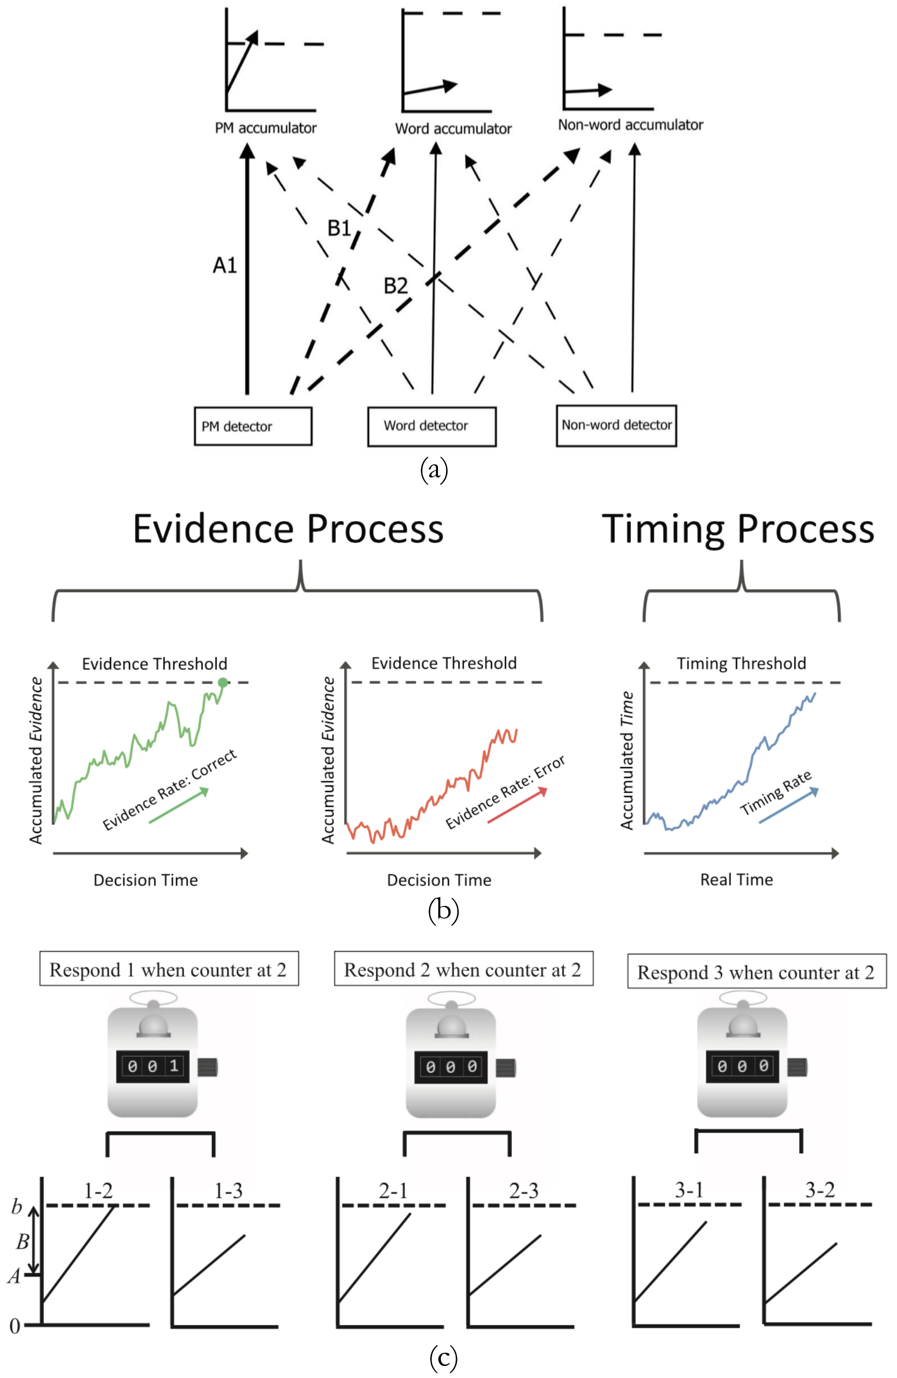

Figure 3. (a) Strickland et al.'s (2018) Prospective Memory Decision Control (PMDC) model. A standard binary (i.e., two accumulators corresponding to two response options) LBA makes a lexical decision (i.e., word vs. non-word choice) based on inputs from word and non-word detectors unless beaten by a third prospective memory (PM) accumulator that receives input from a detector for a PM stimulus attribute (e.g., "tor" as a sub-string of the lexical decision stimulus). Solid arrows indicate excitatory inputs and dashed lines inhibitory inputs. For example, the PM detector both excites the PM accumulator (A1) and inhibits the lexical-decision accumulators (B1 and B2), instantiating reactive control. Response thresholds instantiate proactive inhibition, in this example favouring PM responses by setting a lower threshold for the PM than word accumulator, as only words can contain the PM stimulus. (b) Hawkins and Heathcote's (2021) timed racing-diffusion model (TRDM). A standard binary RDM (the "Evidence Process") makes choices in the usual way unless beaten by a third accumulator (the "Timing Process") in which case a random choice is made. (c) van Ravenzwaaij et al.'s (2020) Advantage LBA (ALBA) model for a choice between three alternatives (1, 2, 3; the same approach can be generalized to any number of choices). Standard LBAs accumulate "advantage" evidence (e.g., 1-2 evidence has a mean equal to the mean for evidence for choice 1 minus the mean for the evidence for choice 2). The counters increment when an accumulator reaches threshold. Under a "win-all" stopping rule, each set of accumulators corresponding to a particular choice is connected to a counter that starts at zero and increments each time an accumulator in the set hits threshold. A response is triggered by the first counter to achieve a count equal to the number of accumulators in the set (e.g., in the figure, respond "1" if its counter is the first to register two counts).

Figure 3 shows examples where the LBA and RDM are used as components in models of complex tasks. Strickland et al.'s (2018) PMDC model (see Figure 3a for details) addresses a prospective memory (PM) task, where participants perform lexical decision (word vs. non-word) trials with letter strings, which in a small subset of randomly selected trials have an attribute (e.g., the sub-string "tor") that requires them to remember to make a different response. To do so, they augmented a binary-choice LBA with a third accumulator for the PM response. PMDC provided a novel cognitive-control-based perspective on PM, illustrating how race models can be used to instantiate both reactive (i.e., stimulus-driven) control, through feedforward inhibition of routine responses when a PM stimulus is detected, and proactive (i.e., anticipatory) control, through setting different thresholds across accumulators.

Hawkins and Heathcote's (2021) TRDM (Figure 3b) combines a binary-choice RDM with a leading model of time perception (Simen et al., 2016), a diffusion process with a constant input and a threshold set so that it is crossed on average at a target time. Previous attempts to model the effect of the passage of time on decision making relied on implicit timing mechanisms like decreasing thresholds. The TRDM provides a comprehensive account of these phenomena with an explicit and hence testable mechanism that is grounded in the literature on timing tasks.

van Ravenzwaaij et al.'s (2020) ALBA models responding contingent on more than one LBA threshold-crossing event. Accumulators correspond to "advantages", evidence for the presence of one stimulus over another (see Figure 3c for details). ALBA accounts for competition effects among choice options previously thought to rule out independent-race models. It extends to multiple choices, and predicts Hick's Law, a logarithmic increase in RT with the number of choices. ALBA uses dynamic logical "AND" functions where all accumulators in a group must finish to trigger a response. Other logical functions can be constructed in the same way (e.g., an "OR" function requires only one group member to finish), providing the building blocks for powerful and general-purpose computations.

Reynolds et al. (2021) proposed an alternative to the ALBA account of making more than two choices in a model that requires only two accumulators. They built on Vicker's (1979) hypothesis that equates confidence with a "balance-of-evidence", the difference between the winning and losing accumulator when the winner achieves threshold. Reynolds et al. added a set of thresholds to each accumulator to translate the balance-of-evidence into discrete responses. For example, a choice is rated as uncertain if the losing accumulator has passed all but the last threshold when the winner achieves threshold (i.e., the looser is not far behind the winner). Figure 4 illustrates this Multiple-Threshold Race (MTR) model and related approaches where a response is triggered as a function of the states of both accumulators, either based on the total number of thresholds passed, or when the thresholds themselves are defined jointly by both accumulators' states. Our second recommended reading, Kvam (2019), proposes a geometric perspective integrating these and other race models, where the states of each accumulator are represented as cartesian coordinates. Building on Smith's (2016) seminal circular-diffusion model, Kvam showed that with enough thresholds, or a non-linear joint threshold, race models can be extended to accommodate continuous responding. Kvam et al. (2022) used this approach to provide a unified account of discrete and continuous responding with respect to line length and colour matching judgements.

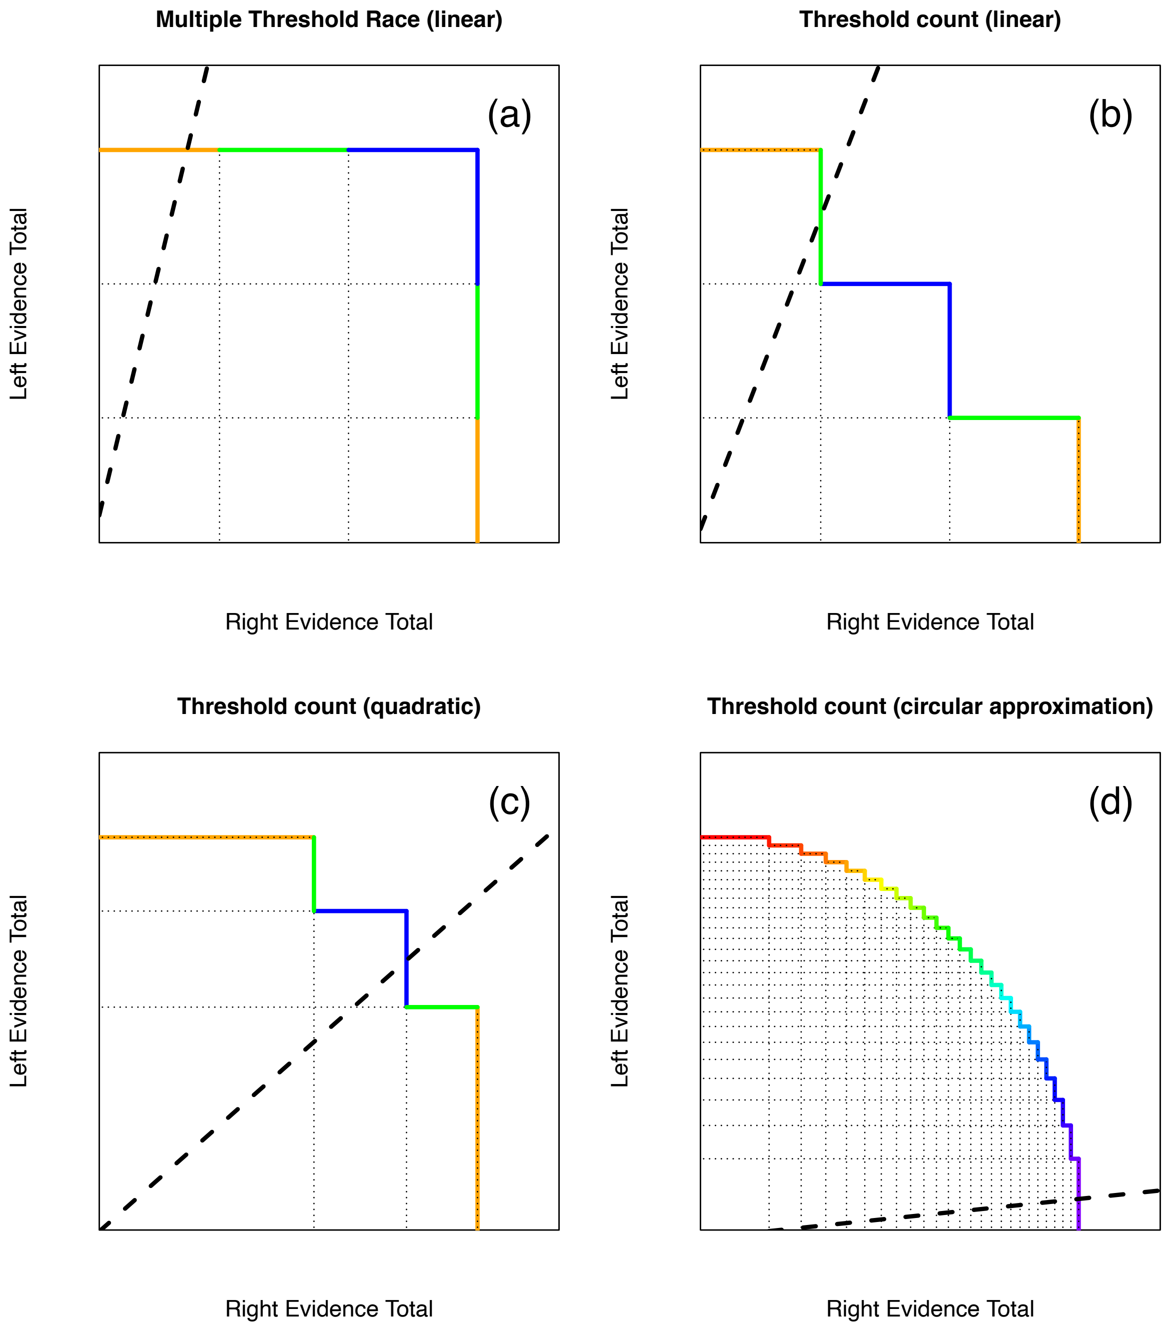

Figure 4. Geometric representations showing how two (left and right) racing accumulators can produce more than two choices by triggering a response based on a particular threshold-crossing event, with the response option chosen also contingent on previously occurring threshold-crossing events. The state of the right accumulator is represented on the x-axis and the state of the left accumulator on the y-axis, so each point on the plane represents the joint state of the two accumulators at a given time. The dashed lines represent examples how the joint state evolves over time when accumulation is linear and deterministic. Each accumulator has multiple thresholds: horizontal dotted lines represent thresholds for the first accumulator and vertical dotted lines the thresholds for the second accumulator. The colored line segments correspond to combinations of threshold-crossing events producing a particular choice. The four panels represent the relationship between Reynolds et al.'s (2021) multiple threshold race (MTR) models and Kvam's (2019) general geometric framework where a response is triggered, and an option chosen, based on an equation combining the x and y states. (a) An MTR with n = 3 thresholds per accumulator is able to choose among 6 options. Responding is triggered by the first time an accumulator's highest threshold is crossed, and the option chosen depends on how many of the losing accumulator's thresholds were previously crossed. For example, orange segments could correspond to a high confidence rating, green segments to medium confidence and blue segments to low confidence, with segments on the horizontal line corresponding to choosing left and segments on the vertical line to choosing right. In this example, the left accumulator crossed all three

thresholds before the right accumulator crossed any, corresponding to a high confidence response. Note that each accumulator could have different numbers of thresholds, thresholds can be spaced unevenly, and different segments could be mapped to the same option. (b) Threshold counting model with n = 3 evenly spaced thresholds: a response is triggered when n thresholds have been crossed. Continuing the confidence example, a medium confidence left response is illustrated, triggered when the first right threshold is crossed after two left thresholds were crossed previously. (c) A threshold counting model with quadratically spaced thresholds (i.e., even spacing on a squared scale). (d) Threshold counting with n = 25 quadratically spaced thresholds for choosing a color of the rainbow (a violet choice is illustrated). In the limit of large n, this approximates a circular boundary in Kvam's geometric framework (i.e., a response is triggered when $\sqrt{x^2 + y^2}$ exceeds the radius of the circle). Note that different spacings in the threshold-counting model can be used to approximate differently shaped geometric model boundaries. Figure available at https://tinyurl.com/yck8xrke under CC-BY 2.0 license (https://creativecommons.org/licenses/by/2.0/).

What are the advantages of using race models?

The expressive power of race models is being used in an increasing number of areas. What motivates researchers to quantitatively instantiate their psychological-process theories in this way? On reason is to avoid mistaken psychological inferences resulting from speed-accuracy tradeoffs. For example, Evans et al. (2018) applied the LBA to data from the Human Connectome Project and showed that performance correlations between twins that had been interpreted in terms of cognitive ability were more likely due to the heritability of response caution. Because correctly identifying underlying causes in this way is key for effective interventions, EAMs are increasingly used in areas ranging from computational psychiatry (e.g., ADHD; Weigard et al., 2018) to performance in time-pressured multi-tasking environments (e.g., Palada et al., 2019).

EAMs are also used as a second stage or “back-end” that enables decision-making effects to be disentangled from the effects of an initial “front-end” stage modelling non-decision phenomena. For example, Steyvers et al. (2019) combined a front-end accounting for trial-to-trial changes in task-set activation with an LBA back-end to model task-switching costs. The specificity of this approach supports clearer selection among theoretical positions, such as applied by White et al. (2011) to attentional-selection theories in vision, Osth and Farrell (2019) to primacy and recency effects in free recall, and Evans et al. (2019) to theories of multi-attribute choice. The latter case is instructive on the utility of the constraint afforded by RTs, as the selected model differed from earlier comparisons based only on choice data. Incorporating RTs also brings measurement advantages. For example, Jones et al.'s (2015) race model of healthcare preferences elicited by identifying the best and worst among a set of options provided the equivalent of more than doubling sample size relative to traditional choice-only models. Front-ends can also model variable processing times that add to a race model's decision time, such as in Provost and Heathcote's (2015) model of the mental-rotation-matching tasks, or can consist of a series of decision stages, as in Fific et al.'s (2010) test of different mental architectures for categorization.

Race models afford even greater expressive ability through probabilistic mixtures that model participants performing tasks in different ways on different trials, accounting for guessing in visual working memory tasks (Donkin et al., 2013), misapplication of complicated decision rules (Bushmakin, et al., 2017), and monkeys who responded with low caution in a putatively high decision-caution condition (Cassey et al, 2014). The ability to test for such effects can have applied as well as theoretical implications. For example, Matzke et al. (2017) used the mixture approach to account for the occurrence of goal neglect that causes the stop runner to fail to enter the race in response to a stop signal, where the finishing times of the go and stop runners are described by the sum of Gaussian and exponential random variables (i.e., ex-Gaussian distribution). Goal-neglect, rather than inhibitory deficits (i.e., slowing of the stop runner), was found to explain poorer performance in the stop-signal task by participants with schizophrenia.

How should race models be used?

Race-model parameters enable insights into the psychological causes of observed phenomena. Going beyond the ability of simple EAMs' to account for speed-accuracy tradeoffs, the race models in Figure 2 enable insights into more complex constructs. For example, Strickland et al. (2018) found that PMDC parameter estimates indicate that prospective-memory failure is mainly due to reactive control rather than proactive control or limited attention capacity.

Race models can also be used to unify theoretical constructs. For example, Hawkins and Heathcote (2021) found that people with greater precision (i.e., less moment-to-moment variability) in the diffusion process underlying a time-interval-production task also had higher precision in the timing component of the TRDM. Reynolds et al. (2020) used the MTR framework to unify race models with perhaps the most widely applied cognitive model, signal-detection theory, providing estimates of its discriminability and bias parameters that are informed by RT.

However, to be valid, parameter values must be uniquely identified by the data from which they are inferred. That is not always guaranteed, particularly when data quality (in terms of the number of observations per participant in each condition) is low. Our third recommended reading, Heathcote et al. (2018), addresses this issue using hierarchical Bayesian estimation. It describes the many advantages of a Bayesian approach and how to perform parameter-recovery studies that fit models to simulated data, enabling estimation accuracy and precision to be assessed by comparing data-generating and estimated parameters. We strongly recommend that parameter-recovery studies be carried out not only for new models but also for applications of existing models to new designs. If problems emerge, either the design or the model must be adjusted. Hierarchical estimation can help where data quality is limited, by constraining each participant's estimates through a group-level model of individual differences. With enough participants, well-estimated group-level parameters can be obtained, although limitations may remain for the participant-level estimates.

Carefully chosen parameterizations are also key for both estimation and interpretation. Ideally, rate parameters should be made a function of stimulus values. van Ravenzwaaij et al. (2020) proposed a simple and general way to do so, and applied it to determine rates (V) when choosing the brighter of left (L) vs. right (R) stimuli:

$$V_L = s + w_S(B_L + B_R) + w_D(B_L - B_R)$$
$$V_R = s + w_S(B_L + B_R) + w_D(B_R - B_L).$$

Parameter s quantifies cognitive-processing speed, and w_S and w_D the weights of the terms based on the sum (i.e., overall magnitude) and difference (i.e., advantage) of the subjective brightness of the two stimuli (B_L and B_R). Subjective brightness followed Weber's Law in being a logarithmic function of the objective luminance. Where objective and/or subjective stimulus values are not available, the processing-speed and magnitude terms are not separately identifiable. However, in any evidence-accumulation model, it can be useful to interpret differences of the estimated rates (i.e., advantages) as reflecting the aggregate impact of the discriminability of the stimuli corresponding to different choices and the quality of selective attention to the stimulus features that support that discrimination, and their sum as reflecting the aggregate impact of internal (e.g., processing speed) and external (e.g., stimulus magnitude) factors driving a response to occur.

Our final example shows how a single race model can provide a multi-faceted characterization of psychological processing that is theoretically revealing, in this case in the domain of healthy cognitive aging. Slowing with age is pervasive, so leading theories posit a reduction in cognitive processing speed with age as a broad explanatory factor, but age-related performance decrements have also been attributed to reduced executive function. Heathcote et al. (2022) investigated these theories using a novel stop-signal task (see Figure 5) requiring inhibition of both easier and harder binary choices as well as proactive control on go trials based on cues about the response most likely to be correct on the upcoming trial. They fit Matzke et al.'s (2017) stop-signal model with an ex-Gaussian stop runner but replaced the model of the go trials with RDM runners. The hybrid ex-Gaussian-RDM approach provides a rich array of measures including the key index of inhibitory ability, stop-signal RT (SSRT), as well as goal neglect, and measures based on RDM go processing, i.e., peripheral processing speed (non-decision time), central processing speed (average accumulation), selective attention (the rate advantage), caution, and proactive control of bias in response thresholds.

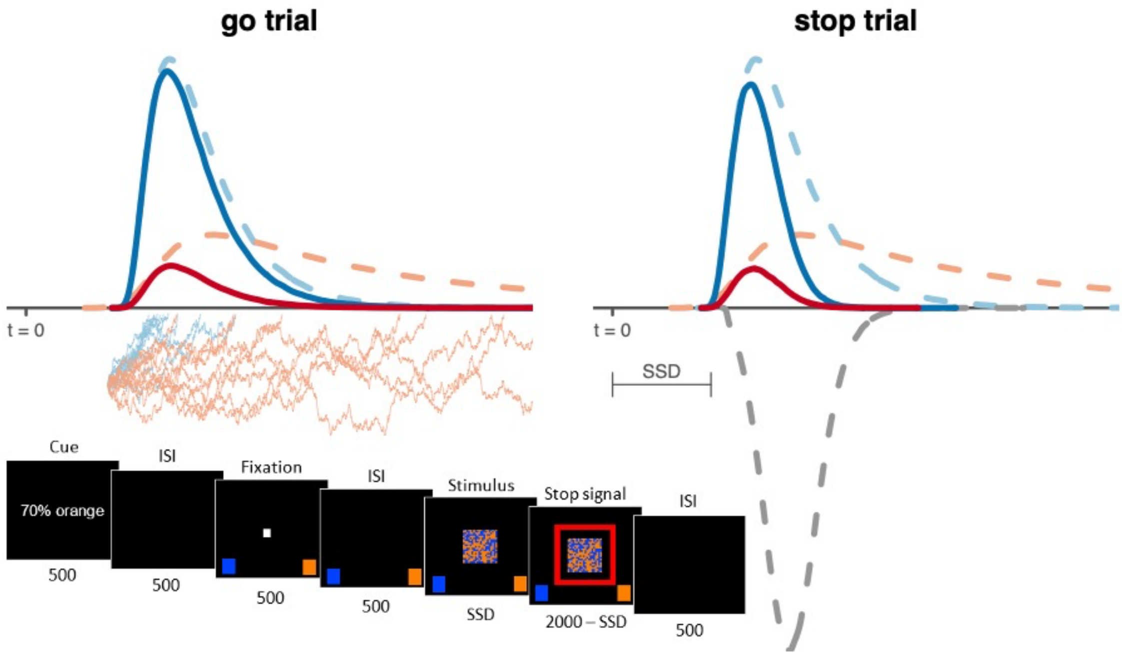

Figure 5. Heathcote et al.'s (2022) stop-signal task and the associated hybrid race model of go and stop trials. A stop trial is shown as a sequence of frames with timing in milliseconds (ms). The stop signal occurred on a randomly ordered 25% of trials. The go (i.e. choice) stimuli were 20 x 20 grids of blue and orange squares whose position changed randomly on each screen refresh (see https://youtu.be/kXZk_jCHjkM for a video of the task). Participants had to decide which color was dominant. Choice difficulty varied randomly from trial to trial, where the dominant color constituted 52% of the squares on harder trials and 54% on easier trials. In the illustration, the cue indicates that there is a 70% chance that orange will be the dominant color in the upcoming stimulus. Stop trials differed from go trials only in that a red square appeared around the choice stimulus at a stop-signal delay (SSD) determined by a staircase algorithm: SSD increased by 50ms after a stop trial where response inhibition was successful (making stopping more difficult on the next stop trial) and decreased by 50ms if response inhibition failed. ISI = inter-stimulus interval. The go-trial model on the upper left illustrates a choice stimulus presented at t = 0. Once it is encoded, two runners corresponding to the orange and blue response options start to race each other from the same initial level. The jagged lines illustrate noisy accumulating evidence totals for 10 races. The total that first crosses its response threshold triggers the associated response. In the illustration, the thresholds are assumed to be the same and correspond to the x-axis in the plot, but when there is response bias, the thresholds for each accumulator may differ. The dashed lines above the x-axis correspond to the distribution across trials of the finishing times of each runner, which in the illustration are longer for the orange accumulator as it has a lower average accumulation rate. The solid lines correspond to the distribution of the winning times for each accumulator, which are shorter than the finishing times because faster runners win races. The stop-trial model on the right shows the finishing and winning time distributions on stop trials. The finishing time of the stop runner is assumed to have an ex-Gaussian distribution, truncated below at 50ms (i.e., additional grey dashed line below the x-axis). The finishing time distributions for the go runners are assumed to be identical on stop and go trials. The presentation of the go stimulus is followed at t = SSD by a stop signal that triggers the stop runner. The go response is successfully inhibited when the stop runner finishes before either of the go runners. Hence, the distributions of winning times for the go runners when inhibition fails (solid lines) are faster than on go trials because slower go finishing times tend to lose out to the stop runner. Figure available at https://tinyurl.com/y5nmbee8 under CC-BY 2.0 license (https://creativecommons.org/licenses/by/2.0/).

Heathcote et al.'s (2022) primary finding is that ageing mainly affected processing speed, both central and peripheral, with only a small part of slowing due to caution increasing with age.

There was no evidence of any age-related deficit in executive-function measures except response inhibition, but again that was mediated by a slowing effect in the speed of the inhibitory runner. If anything, older participants had less goal neglect, and better selective attention and proactive threshold control. These results are surprising given that Ratcliff and McKoon (2008) summarize the results of almost a dozen applications of the traditional DDM to aging as showing “the slowdown is almost entirely due to older adults’ conservativeness [i.e., caution].” (p. 911). One explanation is that, in being driven purely by the difference in evidence between two options, the DDM is unable to represent stimulus magnitude and central processing speed effects, and so any reduction to the latter in older participants can only be attributed to increased caution. A more recent revision of the DDM can account for stimulus magnitude with the additional assumption that the mean and variability of rates are positively related (Ratcliff et al., 2018). Further research will be required to determine what this new type of DDM indicates about the causes of age-related slowing.

These results not only show how race models can provide theoretical insights, they also highlight the challenges with respect to parameter identifiability and interpretability. First, the descriptive ex-Gaussian account of the stop runner had to be retained in the hybrid model because Matzke et al. (2020) showed that stop-signal models in which all runners are evidence-accumulation processes result in extremely poor parameter recovery. In contrast, the hybrid model’s parameters are well recovered, even in Heathcote et al.’s (2022) complex design, with the key factor being that the ex-Gaussian does not require the estimation of a parameter corresponding to the distribution’s lower bound. Second, the strong theoretical divergence between the DDM and the hybrid model underlines the need to carefully consider the limitations in what a model can and cannot represent, and the potential for parameter tradeoffs, when interpreting any model-based estimate.

Future Directions

A range of emerging applications are using race models to integrate different data sources, tasks, and theoretical frameworks. For example, the LBA has been used to simultaneously link behavioral data to fMRI and EEG data through hierarchical Bayesian estimation (Turner et al., 2016). Similar methods have also been used to link performance in multiple tasks through shared LBA parameters (Wall et al., 2021). We believe that these "joint" models are more likely to be successful at linking data of different types and from different tasks than approaches such as structural equation modeling that ignore the details of cognitive processing, or implicitly assume details without making them explicit (and hence testable) as is done by joint models.

Our final recommended reading, Miletić, et al., (2019), generalizes and improves on previous attempts to integrate another very successful type of cognitive model, reinforcement learning, into the race-model framework. They used a simple delta rule to learn reward values for stimuli based on probabilistic feedback from forced choices. These values played the role of brightness in van Ravenzwaaij et al.'s (2020) equations, determining rates for RDM accumulators. In agreement with van Ravenzwaaij et al.'s results for perceptual choice, the learned values affected rates mainly through the advantage term, but the magnitude term also played a non-negligible role. In binary-choice tasks, the model outperformed previous approaches that combined the DDM with reinforcement learning, and then went beyond what can be achieved with the DDM, successfully addressing a three-choice task and reward-magnitude effects.

In closing we note that most of the race models discussed here have been analytically tractable enough to yield an easily computed likelihood, the key quantity required for fitting the models to data in a comprehensive way. However, requiring such tractability can limit the scope of potential applications. Fortunately, recent developments in approximate likelihood methods, combined with faster hardware (Lin et al., 2019) and deep neural-network approaches to Bayesian estimation (Radev et al., 2020) look set to make it increasingly practical to apply race models to a wider range of applications.

Acknowledgements

Preparation of this article was supported for AH by a Révész Visiting Chair at the Department of Psychology, University of Amsterdam, and Australian Research Council Discovery Project DP210100313. DM was supported by a Vidi grant (VI.Vidi.191.091) from the Netherlands Organization of Scientific Research (NWO).

References

Bogacz, R., Usher, M., Zhang, J., & McClelland, J. L. (2007). Extending a biologically inspired model of choice: Multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *362*(1485), 1655–1670.

Bushmakin, M. A., Eidels, A., & Heathcote, A. (2017). Breaking the rules in perceptual information integration. *Cognitive Psychology*, *95*, 1-16.

Cassey, P., Heathcote, A., & Brown, S. D. (2014). Brain and behavior in decision-making. *PLoS Computational Biology*, *10*(7), e1003700.

Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychological Review*, *120*, 873–902.

Evans, N. J., Holmes, W. R., & Trueblood, J. S. (2019). Response-time data provide critical constraints on dynamic models of multi-alternative, multi-attribute choice. *Psychonomic Bulletin & Review*, *26*, 901–933. http://doi.org/10.3758/s13423-018-1557-z

Evans, N. J., Steyvers, M., & Brown, S. D. (2018). Modeling the covariance structure of complex datasets using cognitive models: An application to individual differences and the heritability of cognitive ability. *Cognitive Science*, *42*(6), 1925–1944.

Fific, M., Little, D. R., & Nosofsky, R. M. (2010). Logical-rule models of classification response times: A synthesis of mental-architecture, random-walk, and decision-bound approaches. *Psychological Review*, *117*, 309–348.

Heathcote, A., Garton, R., Hinder, M., Reynolds, A., Tanis, C., & Matzke, D. (2022). Slower processing dominates executive-function deficits in cognitive aging. https://psyarxiv.com/udz7a

Heathcote, A., Suraev, A., Curley, S., Love, J., & Michie, P. (2015). Decision processes and the slowing of simple choices in schizophrenia. *Journal of Abnormal Psychology, 124*, 967-974.

Jones, L. G., Hawkins, G. E., & Brown, S. D. (2015). Using best-worst scaling to improve psychological service delivery: An innovative tool for psychologists in organized care settings. *Psychological Services, 12*, 20-27.

Kvam, P., Marley, A. A. J., & Heathcote, A. (2022). A unified theory of discrete and continuous responding, https://psyarxiv.com/dkybt/

Lin, Y-S., Heathcote, A., & Holmes, W. R. (2019). Parallel probability density approximation. *Behavior Research Methods*, *51*, 2777–2799.

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*(3), 295–327.

Matzke, D., Hughes, M., Badcock, J. C., Michie, P., & Heathcote, A. (2017). Failures of cognitive control or attention? The case of stop-signal deficits in schizophrenia, *Attention, Perception & Psychophysics*, *79*, 1078-1086.

Matzke, D., Logan, G. D. & Heathcote, A. (2020). A cautionary note on evidence-accumulation models of response inhibition in the stop-signal paradigm, *Computational Brain & Behavior*, *3*, 269-288.

Osth, A. F., & Farrell, S. (2019). Using response time distributions and race models to characterize primacy and recency effects in free recall initiation. *Psychological Review*, *126*, 578-609.

Palada, H., Neal, A., Strayer, D., Ballard, T., & Heathcote, A. (2019). Using response time modeling to understand the sources of dual-task interference in a dynamic environment. *Journal of Experimental Psychology: Human Perception and Performance*, *45*, 1331–1345.

Provost, A., & Heathcote, A. (2015). Titrating decision processes in the mental rotation task. *Psychological Review*, *122*, 735-754.

Radev, S. T., Mertens, U. K., Voss, A., Ardizzone L., & Köthe, U. (2020). BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*. doi: 10.1109/TNNLS.2020.3042395.
Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. http://www.mitpressjournals.org/doi/abs/10.1162/neco.2008.12-06-420
Ratcliff, R., Voskuilen, C., & Teodorescu, A. (2018). Modeling 2-alternative forced-choice tasks: Accounting for both magnitude and difference effects. *Cognitive Psychology*, *103*, 1–22. https://doi.org/10.1016/j.cogpsych.2018.02.002
Reynolds, A., Garton, R., Kvam, P., Griffin, V., Sauer, J., Osth, A., & Heathcote, A. (2021). A dynamic model of deciding not to choose. *Journal of Experimental Psychology: General*, *150*, 42-66.
Reynolds, A., Kvam, P., Osth, A., & Heathcote, A. (2020). Correlated racing evidence accumulator models. *Journal of Mathematical Psychology*, *96*, 102331.
Simen, P., Vlasov, K., & Papadakis, S. (2016). Scale (in)variance in a unified diffusion model of decision making and timing. *Psychological Review*, *123*(2), 151–181.
Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, *123*, 425–451.
Steyvers, M., Hawkins, G. E., Karayanidis, F., & Brown, S. D. (2019). A large-scale analysis of task switching practice effects across the lifespan. *Proceedings of the National Academy of Sciences, 116*, 17735-17740.
Strickland, L., Loft, S., Remington, R. W., & Heathcote, A. (2018). Racing to remember: A theory of decision control in event-based prospective memory. *Psychological Review, 125*, 851-887.
Turner, B. M., Rodriguez, C. A., Norcia, T. M., McClure, S. M., & Steyvers, M. (2016). Why more is better: Simultaneous modeling of EEG, fMRI, and behavioral data. *NeuroImage*, *128*, 96–115.
van Ravenzwaaij, D., Brown, S. D., Marley, A. J., & Heathcote, A. (2020). Accumulating advantages: A new approach to multialternative forced choice tasks. *Psychological Review, 127*, 186–215.
Vickers, D. (1979). *Decision processes in visual perception*. New York, NY: Academic Press.
Wall, L., Gunawan, D., Brown, S. D., Tran, M.-N., Kohn, R., & Hawkins, G. E. (2021). Identifying relationships between cognitive processes across tasks, contexts, and time. *Behavior Research Methods, 53*, 78-95.
Weigard, A., Huang-Pollock, C., Brown, S., & Heathcote, A. (2018). Testing formal predictions of neuroscientific theories of ADHD with a cognitive model–based approach. *Journal of Abnormal Psychology*, *127*(5), 529-539.
White, C. N., Ratcliff, R., & Starns, J. J. (2011). Diffusion models of the flanker task: Discrete versus gradual attentional selection. *Cognitive Psychology*, *63*, 210–238.

Recommended Readings

Donkin, C., & Brown, S. D. (2018). Response times and decision-making. In E.-J. Wagenmakers & J. T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience*, Volume five: Methodology (4th ed., pp. 349-382). John Wiley & Sons, Inc. Provides a broad and integrative historical overview of evidence-accumulation models, their applications from measuring latent constructs to theory development in the cognitive and neurosciences, and methods for estimation, evaluation, and inference.

Kvam, P. D. (2019). A geometric framework for modeling dynamic decisions among arbitrarily many alternatives. *Journal of Mathematical Psychology*, *91*, 14–37. A more advanced reading, which won the 2021 R. Duncan Luce Outstanding Paper Award from the Society for Mathematical Psychology. It provides a novel perspective on both deterministic and stochastic race models, placing them in an integrative geometric framework that encompasses both discrete and continuous responding.

Heathcote, A., Lin, Y-S, Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2018). Dynamic models of choice. *Behavior Research Methods*, *51*, 961-985. A practical introduction to estimation and inference for a variety of race models using hierarchical Bayesian Markov-Chain Monte Carlo sampling implemented in a collection of functions and associated tutorials in the R language.

Miletić, S., Boag, R. J., Trutti, A. C., Stevenson, N., Forstmann, B. U., & Heathcote, A. (2021). A new model of decision processing in instrumental learning tasks. *eLife, 10*:e63055. Evaluates how well the DDM and a variety of RDMs integrate with simple model-free reinforcement learning, finding that input equations and architectures from van Ravenzwaaij et al.'s (2020) "advantage" framework provide the best account of learning in value-based decision making.