Disentangling interacting psychological processes: A model-based approach to cognitive aging.

Andrew Heathcote^{1,2}

Roderick Garton³

Mark R. Hinder³

Angus Reynolds³

Charlotte Tanis²

Dora Matzke²

¹ Department of Psychology, University of Newcastle, Australia

² Department of Psychology, University of Amsterdam, The Netherlands

³ School of Psychological Sciences, University of Tasmania, Australia

Abstract

We propose a new cognitive-modelling methodology that disentangles the non-linear interactions among psychological mechanisms underpinning behaviour in a complex task. We apply the approach to the domain of cognitive aging and show that it affords simultaneous measures of the major mechanisms proposed to explain age-related deficits: capacity limits, processing speed, inhibition, and executive function. The measures' validities derive from the precisely specified roles of the model parameters and the model's comprehensive, detailed, and accurate account of the behaviour of our healthy older and younger participants. Because it derives constraint from all aspects of participant's behaviour, the analysis is powerful enough to provide clear results from a relatively small sample of participants who provide high quality data. When controlling for all other factors, we identify processing speed as the primary cause of age-related deficits. Executive control was better for older than younger participants when task demands were lower, but as demands increased this advantage reduced or disappeared, consistent with age-related capacity limitations. We conclude that such detailed cognitive models have the potential to be widely applied to provide an integrated, quantitative, and conceptually rich account of the many interacting psychological mechanisms that support behaviour in complex tasks.

A variety of broad explanations of cognitive aging have been proposed, from depletion of processing capacity (Craik & Byrd, 1982), through reduced inhibitory ability (Hasher & Zacks, 1988) and executive control (Braver & West, 2008), to a general decrease in the speed of cognitive operations (Salthouse, 1996). Traditional approaches to evaluate the influence of these factors within and across individuals typically use generic, linear, correlation-based methods and require many hours of testing from hundreds of participants each performing tens of different tasks (Salthouse et al., 2003). The validity of this approach rests on grouping different tasks according to assumed common psychological mechanisms. Summary measures that describe limited aspects of performance in each task are then used to separate task-specific variance from variance shared by the common mechanisms.

Here we develop and evaluate a more efficient approach grounded in a cognitive-modelling framework that is based on decades of research in experimental psychology and the neurosciences (Donkin & Brown, 2018; Matzke et al., 2018). It assumes that: 1) the broad mechanisms of interest in cognitive aging can be simultaneously active in a single complex task; and 2) their influence can be determined by estimating the parameters of a detailed quantitative model of the full array of psychological processes required to perform the task. The model's parameters summarize non-linear interactions among these processes within each individual. Hence, these parameters have clear psychological interpretations based on the explicitly defined roles they play in the model's information-processing architecture. The validity of the model itself in any given application is subject to direct verification through its ability to provide a comprehensive account of task behaviour. Because all aspects of performance constrain the model, this analysis is efficient, enabling clear conclusions in designs with a modest number of participants as long as each is well measured.

We demonstrate the utility of this approach in an experiment where younger and older participants performed many trials of a single complex task (see Figure 1). We show that a newly developed cognitive model (Tanis et al., 2022) provides an accurate and comprehensive description of their behaviour and a clear indication of which mechanisms are responsible for age-related changes in performance. Before reporting these results, we first motivate the task and the model based on the two major classes of deficits that have been proposed to cause cognitive aging: deficits in processing speed and in executive control.

Slowing with age has been described as “one of the best-documented and least controversial behavioural phenomena of aging” (Salthouse, 1996, p. 406). It occurs whether measured by repetitive psychomotor tasks, reduced accuracy for psychophysical choices based on briefly presented stimuli, psychometric paper-and-pencil tests of series of easy or hard decisions, the time course of psychophysiological responses, or choice response time (RT; Salthouse, 2000). All these findings have been explained by a reduction in the rate at which elementary cognitive operations are executed (Salthouse, 1996). Although not all age-related differences are determined by this single factor, processing speed is implicated in the age-related correlations among a wide variety of memory and cognitive abilities (Salthouse, 2000).

Inhibitory deficits, broadly affecting executive-control mechanisms thought to be localized in prefrontal cortex (Braver & West, 2008), provide another explanation of age-related slowing (Hasher & Zacks, 1988). A recent meta-analysis (Rey-Mermet & Gade, 2018) identified more than 250 studies aiming to measure the effect of aging on inhibitory processes. The studies used a wide variety of paradigms that compared RT in conditions placing more vs. less demand on inhibitory control of conflicting stimulus information as well as response inhibition paradigms, including simple versions of the stop-signal task used here (Matzke et

al., 2018). This meta-analysis found little evidence for deficits in inhibition related to the selective attention required when different aspects of a choice stimulus support conflicting responses, but it did support deficits in response inhibition.

Our participants performed an RT task, making binary button-press responses as soon as they believed they can accurately discriminate which of two colours dominates a stimulus. This discrimination is difficult (i.e., error prone even when decisions are not rushed) because the stimulus array has colours associated with each response, causing response conflict. Hence, good performance requires participants to use attention to select and compare appropriate evidence to resolve the conflict. We manipulated the strength of the colour majority within subjects to test if older participants were more affected by an increase in choice difficulty. To measure response inhibition, on 25% of the trials a stop signal appeared after the onset of the choice stimulus, instructing participants to change their goal from making a choice to withholding their response. As stop trials were less common than choice-only (“go”) trials, executive-control is required to maintain the stopping goal in the face of the habitual tendency to respond to the choice stimulus.

Figure 1. Display sequences with timings below each frame (in seconds [s]) for “go” (choice only) and stop trials. The position of the orange and blue squares in the choice stimulus changed randomly on each screen refresh (see https://youtu.be/kXZk_jCHjK4 for a video of the experiment). In the illustration, the cue indicates that there is a 70% chance that orange will be the dominant colour in the upcoming stimulus. In the “blockwise” cueing condition, the cue remained consistent across a block of 60 trials, whereas in the “trialwise” cueing condition, it changed unpredictably from trial to trial with equal likelihood across a block. The illustration depicts a correct (i.e., orange response to orange-dominated stimulus) and time-valid ($0.2s < RT < 2s$) go trial, with feedback showing the corresponding RT (in milliseconds [ms]). 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 0.05s after a stop trial where response inhibition was successful (making stopping more difficult on the next stop trial) and decreased by 0.05s if response inhibition failed. Note that for illustrative purposes this figure is not exactly to scale. ISI = inter-stimulus interval. Figure available at <https://tinyurl.com/2p8h7dkn> under CC-BY 2.0 license (<https://creativecommons.org/licenses/by/2.0/>).

The task also probed two further types of executive control through additional within-subject manipulations. Proactive control (Braver, 2012) was measured by cueing participants about which choice was most likely to be correct prior to each trial and comparing performance on

cue-congruent and cue-incongruent trials. The ability to flexibly switch between tasks (Monsell, 2003) as examined by changing the cue randomly between trials (“trialwise” cueing) vs. keeping it constant over a block of 60 trials (“blockwise” cueing).

Identifying specific executive-function deficits by contrasting conditions is difficult because of the general pattern of slowing, and of increasing differences between performance in easier and harder conditions, for older participants. We addressed this issue using a relative or “greater-than-proportional” standard for identifying specific deficits. For example, if older participants are slower than younger participants by a multiple of two in an easier condition, a *specific* deficit is only inferred if the multiple is more than two in a harder condition.

Methods

Open Practices Statement

The study was not preregistered. The sampling plan and data exclusions are described in the Experimental Methods section. Data and analysis code are made openly available at <https://osf.io/4bzav/> as described in the Data Preparation and Analysis section.

Experimental Methods

The sampling plan aimed for 20 participants younger than 30 years and 20 participants older than 60 years. All participants were self-screened for any relevant neurological, visual, or motor deficits, provided prior written consent to participate and to make their de-identified trial-level data publicly accessible, and were informed that the research procedures were approved by the UTAS Human Research Ethics Committee. Younger participants were University of Tasmania students, with 11 being first-year psychology students who received course-credit for participation and the remainder receiving AU\$30 in shopping vouchers. Older participants were healthy community-dwelling local residents who were recruited from a pool of respondents to advertisements for participants in age-related research and received AU\$40 in vouchers.

Data collection continued until the aim of the sampling plan was achieved, with two participants in the younger group being rejected during this process due to non-compliance with experimental instructions and associated near chance accuracy, and large proportions of very fast responses ($RTs < 0.2s$) and non-responses. The final group of 20 younger participants (70% female) had an age range of 18–28 years (mean = 21.4, $SD = 3.2$), excepting four participants whose age was verified as less than 30 but whose date of birth was not obtained. The final group of 20 older participants (65% female) had an age range of 63–76 years (mean = 68.3, $SD = 3.7$).

Participants completed both go-only and stop-signal versions of the dominant-colour judgement task (see Figure 1) in each of two sessions. For the sake of completeness, we also describe the go-only version, but our main analyses focus only on the stop-signal version. Every trial of both versions was preceded by a cue accurately indicating either that the correct response would be orange 70% of the time on the upcoming trial or that the correct response would be blue 70% of the time. The sessions themselves were distinguished by the biasing procedure: either blockwise or trialwise, with the order of administration counterbalanced within each age group. Accordingly, the biasing procedure was constant across the go-only and stop-signal versions of the task within a session. In the blockwise procedure, the same cue was used for all trials during a block, with half of the blocks having an orange cue and

half a blue cue. In the trialwise procedure, half of the cues in a block were orange and half blue, with their order being randomly shuffled.

The interval between the go-only and stop-signal versions within a session was participant controlled. Each version concluded within ~45–60 minutes, such that each session concluded within ~2 hours. The two sessions were separated for most participants by 1 to 7 days, up to a maximum of 14 days (mean = median = 5 days), depending on participants' availability.

After introductory verbal and on-screen instructions, each session commenced with an orientation block of 20 go-only trials, with equal frequency of the two targets (that is, of blue or orange as the dominant colour), and with very easy discriminations where the target colour filled 60% of the cells of the 20×20 stimulus array. This was followed by 2×40-trial practice blocks, and then 8×60-trial experimental blocks of go-only trials, within which the dominant colour filled either 52% or 54% of the stimulus array—for “hard” and “easy” stimulus discrimination, respectively—on an equal-frequency but randomly-ordered basis, and according to the cueing procedure for the session. On each trial, while keeping these proportions constant, the blue/orange color of each cell was randomly shuffled on every screen refresh.

The stop-signal version then commenced, with the same biasing procedure as implemented across the foregoing go-only blocks, starting with 2×40-trial practice blocks, followed by 8×60-trial experimental blocks. The stop trials differed from go trials by displaying, after a short delay (stop-signal delay; SSD), a thick red border around the stimulus. Also, the stimulus was displayed until 2 seconds had passed rather than being offset (as on go trials) as soon as a response was made. The proportion of stop trials within a stop-signal block was fixed at .25; that is, a stop signal was presented on 15 of the 60 trials per block. The trial on which the signal appeared was pre-set, independently of within-block conditions or performance.

Verbal and on-screen instructions at the start of the stop-signal version described that (1) it was important to withhold responding when the stop signal appeared but to respond as quickly and accurately as possible without trying to anticipate when the stop signal might appear; and (2) the stop signal could appear after the response was made. These instructions, which read as follows, were repeated on-screen at the start of each block: "During the next block there will be some trials where a red border will appear around the stimulus. If you see the red border, try to halt your response (do not press a key). Try not to slow your response by waiting for the border to appear. Sometimes the border will come too late for you to halt your response, and this is okay. Accurate and fast responding to the stimulus is more important.". These instructions continued by detailing the biasing procedure. For the blockwise biasing session, the colour that was most likely to be the dominant colour in that block was also specified.

SSD was set to 0.2s at the beginning of the first practice block in the stop-signal version and was updated after every stop trial, with the final updated value in a block setting the SSD on the first stop trial of the subsequent block. The adjustment after a stop trial was as follows. If no response was made before or after the stop signal (i.e., when response inhibition succeeded), SSD for the next stop trial was set to equal .05s more than the current SSD, or 2s, whichever was the lower value. If, however, a response was made before or after the onset of the stop signal (i.e., when response inhibition failed), SSD for the next stop trial was set to equal .05s less than the current SSD, or .05s, whichever was the higher value.

For the go-only version, and for go trials in the stop-signal version, one of four possible feedback messages was displayed for 0.5s according to the serially exclusive conditions of: (1) "too slow" given no response ahead of the RT maximum of 2s; (2) "too fast" if the RT was faster than 0.2s; (3) "incorrect" if the (time-valid) response value did not equal the target value; or (4) the (correct and time-valid) RT on that trial. No feedback message was provided at the end of a stop trial; the stop signal effectively took the place of the feedback message on these trials. The complete sequence of events on go and stop trials is presented in Figure 1. At the end of each block, feedback was provided on mean correct go RT (excluding RTs < .2s) and the accuracy of time-valid responses across all go trials.

Data Preparation and Analysis

Raw data and an R (R Core Team, 2021) script (*getAllDataForFitting.R*) to prepare the data for fitting are available on the Open Science Framework (OSF) at <https://osf.io/4bzav/>. Code for all manifest analyses is provided in the R markdown document *manifest.Rmd* with output in *Manifest.nb.html*. Models were fit (on the second scale) by hierarchical Bayesian methods using the DMC software (Heathcote et al., 2019) for R. The R markdown document *Model.Rmd* on the OSF provides R code that specifies the models and the prior distributions, performs goodness-of-fit and parameter analyses, and generates all figures featuring model parameters in the main paper and in the supplementary materials. Supplementary materials on the OSF also provide details of the model and the way in which it was fit.

For all analyses trials with RT < 0.2s (0.09%) were removed. Analyses detailed in the preparation script removed two participants from each group who did not perform in the stop-signal session as instructed, with three slowing across the course of the experiment and one younger participant starting very slow then speeding up. This resulted in atypical RT distribution shapes where, instead of the usual long right tail, RTs stacked up near the 2s upper bound of the response window. With these participants removed, there were 18 participants remaining in both the younger and older groups.

Results

Figure 2 illustrates the model and Figure 3 its fit to the data. The model for go trials assumes that it first takes some time to encode the choice stimulus into visual short-term memory (Smith & Ratcliff, 2009). Evidence is then repeatedly sampled into accumulators corresponding to each choice. Accumulators sum their evidence to average out the effects of moment-to-moment (diffusive) noise (Tillman et al., 2020). The two choice accumulators race until the first reaches its threshold, triggering the corresponding response. Participants set thresholds strategically based on expectations about the upcoming trial. RT is the sum of decision time (i.e., time to accumulate to threshold) and non-decision time (i.e., the time to initially encode the stimulus plus the time to execute a motor response).

Figure 3(a) shows that the model captures the large age-related slowing on go trials, not only in terms of an increase in mean RT (by 0.21s), but also in the variance and positive skew of RT distributions. Figures 3(b-h) show that the model also accurately captures all significant main and interaction effects on go trial mean RT and accuracy. Figures 3(b-d) show slowing for hard over easy responses (by 0.06s), and both trialwise over blockwise cueing and cue-incongruent over cue-congruent responses (both by 0.04s). In all three manipulations, the age-related speed deficit was larger for the slower conditions, and this also held in a greater-than-proportional sense, except for the Age×Bias interaction. Figure 3e shows a large decrease in accuracy with difficulty (by 11%), with older participants a little more accurate,

and more so for easy than hard choices (3% vs. 2%). Figure 3f shows that RT also increased with difficulty, more so for cue congruent (0.07s) than incongruent (0.04s) trials. Figures 3(g-h) show that bias effects on RT and accuracy were larger with blockwise than trialwise cueing (0.04s vs. 0.03s and 13% vs. 5%, respectively). Figure 3 (i-j) show that the model captures the increasing effect of stop-signal delay (SSD) on the probability and RT of responses that escaped inhibition, as well as a greater probability and faster responding for younger than older participants at the same SSD.

Figure 2. Graphical representation of the model and the corresponding likelihood equations. The left panel illustrates go trials with a choice stimulus presented at $t = 0$. Once the choice stimulus 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 across-trial distributions of 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. Equations 1 and 2 specify the likelihood (L) of making responses 1 (1=orange) and 2 (2=blue), respectively, at time t on a go trial, where θ_1 and θ_2 are the parameter vectors of the corresponding runners, and f_g and S_g are the Wald density and survivor functions for a go runner. The right panel 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 (the sum of a normal and exponential random variables), truncated below at 0.05s. Finishing time distributions for the go runners are assumed to be identical on stop and go trials. When the stop runner is an evidence-accumulation process, the stop-related model parameters are not identifiable (Matzke et al., 2020). This identifiability problem can be solved with the “hybrid” race model used here which combines a cognitive-process (Wald evidence-accumulation) and a descriptive (ex-Gaussian) framework for modelling the go and the stop processes, respectively (Tanis et al., 2022). On stop trials, the presentation of the choice stimulus is followed at $t = \text{SSD}$ by a stop signal that triggers the stop runner. The additional grey dashed line below the x-axis shows the finishing time of 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 (signal-respond RTs, solid lines) are faster than on go trials because slower go finishing times tend to lose the race to the stop runner. Equations 3 and 4 give the likelihood of orange and blue signal response trials, respectively, at time t , where θ_s is the parameter vector for the stop runner and S_s the ex-Gaussian survivor function. Equation 5 gives the likelihood of response inhibition, where f_s is the ex-Gaussian density function. We used a mixture-likelihood version of these equations that also accounts for lapses in goal maintenance (Matzke et al., 2019). Figure available at [tiny.cc/k8xpuz](https://creativecommons.org/licenses/by/2.0/) under CC-BY 2.0 license (<https://creativecommons.org/licenses/by/2.0/>).

The accurate descriptive account of both go and stop trials suggests that the model's parameters provide a valid and comprehensive characterization of participants' behaviour. Next, we interpret age-related differences in model parameters, describing them as “credible” if the associated 95% Bayesian credible intervals do not contain zero (Supplementary Materials provide credible intervals for all main effects and interactions).

Figure 3. Data (lines joining solid points) and model fit (open points with 95% credible intervals). (a) Go (i.e., choice) RT distribution percentiles on go trials averaged over conditions and participants as a function of age: greater variance and skew for older participants is indicated by the increasing spread of the percentiles for slower RTs. See the Supplementary Materials (<https://osf.io/4bzav/>; Figure S1) for go RT distribution plots separately for each within-subject condition broken down by correct and error responses. (b)-(h) All significant two-way interactions; note that no higher-order interactions were significant. See the Supplementary Materials for accuracy plots separately for each within-subject condition (Figure S2) and the results of all inferential tests on accuracy, mean RT, and the mean of the logarithm of RT (to test proportional effects) for the full $2 \times 2 \times 2 \times 2$ design. (i-j) The probability of responding on stop trials and median RT when doing so (SRRT = signal-response RT) as a function of stop-signal delay (SSD). Note that results are aggregated over ranges of SSDs so that there are approximately equal numbers of observations contributing to each point. See the Supplementary Materials for SRRT distribution and accuracy plots separately for each within-subject condition and the two-way interactions involving signal responses (Figure S3-S5), and inferential results for the probability of responding and SRRT for the full design.

Go Parameter Analysis

Figure 4 shows the main determinants of overall RT in the go task: non-decision time, and the two factors controlling decision time, processing speed and caution. Longer decision time for older participants explained a little over half of age-related slowing, with increased non-decision time for older participants explaining the balance. The first factor controlling decision time, the average of the go accumulator rates, is determined by both the strength of the stimulus encoding and participants' intrinsic speed of cognitive processing. As we always make comparisons that control for stimulus factors (e.g., the same stimuli for older and younger participants), we call this factor *processing speed*. The second factor, *caution* (the average of the two choice accumulators' thresholds) is set strategically by participants to manage the trade-off between decision speed and accuracy; greater caution gives more time to average out the effects of diffusive noise, and so increases accuracy, but at the cost of slowing. Approximately three quarters of the increase in decision time was due to reduced processing speed for older participants, with increased caution explaining the remainder.

Figure 4. Mean parameter estimates for younger and older participants (bars) with 95% credible intervals. The median differences between age groups (scored to be positive: older–younger for the left and right panels and younger–older for the middle panel), with associated 95% credible intervals, are given at the top of each panel.

We describe the difference between the rate for the accumulator that matches the stimulus (e.g., the “orange” accumulator for a predominantly orange stimulus) and the rate for the accumulator that mismatches the stimulus as “discrimination”. Discrimination is determined both by stimulus characteristics (e.g., proportion of blue vs. orange) and participants’ ability to selectively attend to evidence supporting an orange or a blue choice to provide input to the corresponding accumulator. Figure 5 shows that greater discrimination is the main determinant of better performance for easy than hard stimuli, with processing speed having little impact. Note that thresholds can play no role in explaining easy vs. hard effects as stimulus difficulty varied unpredictably from trial to trial (Donkin et al., 2011).

Because all participants were presented with the same stimuli, discriminability comparisons measure age-related selective-attention differences. Figure 5 shows that a small difference between age groups in absolute discrimination results in a large proportional advantage for older participants due to their substantially reduced processing speed. These selective-attention advantages for older participants interact with difficulty, being larger for easy stimuli in both an absolute (0.15 [0.02, 0.28]) and relative (0.17 [0.1, 0.24]) sense.

Figure 5. Mean accumulation-rate-based parameter estimates for the hard (top row) and easy (bottom row) stimulus difficulty condition for younger and older participants (bars), with 95% credible intervals. (a-b) Processing speed. (c-d) Discrimination, i.e., difference between matching and mismatching accumulation rates. (e-f) Relative discrimination, i.e., percentage of the maximum possible difference between matching and mismatching accumulation rates as a proportion of the average rate. The median differences between age groups (younger–older for the right column and older–younger for the other columns), with associated 95% credible intervals, are given at the top of each panel. See Supplementary Materials (Figure S6) for plots of the mean accumulation-rate-based parameter estimates for all within-subject conditions.

Information about whether an orange or blue stimulus is more likely on the upcoming trial can be used to make responses that are more optimal (i.e., faster, and more accurate, overall) by setting a lower threshold for the cue-congruent accumulator (Garton et al., 2019). We subtracted the threshold for the more likely response from that of the less likely response to measure proactive control over response bias. As shown in Figure 6, the greater flexibility required to set response bias in the trialwise condition resulted in much smaller bias than in the blockwise condition, but the bias manipulation had little effect on overall caution. Older participants were slightly more responsive to prior information in an absolute sense, but in a relative sense there was no evidence of an age difference.

Figure 6. Mean threshold-based parameter estimates for the trialwise (top row) and blockwise (bottom row) bias conditions for younger and older participants (bars), with 95% credible intervals. (a-b) Caution, i.e., the average of the two choice accumulators' thresholds. (c-d) Response bias (RB), i.e., threshold of the less likely response minus threshold of the more likely response. (e-f) Relative RB = absolute response bias divided by caution. The median difference between age groups (older-younger in all cases), with associated 95% credible intervals, are given at the top of each panel.

Figure 7 shows that the bias manipulation had little effect on overall processing speed. Absolute discrimination was greater for older participants in the blockwise condition, which constituted a credible interaction (0.27 [0.14,0.41]). Proportionally, discrimination was better for older participants in both cases, but more so in the blockwise condition, again a credible interaction (5.3 [1.7,9.1]).

Figure 7. Mean accumulation-rate-based parameter estimates for younger and older participants (bars) for the trialwise (top row) and blockwise (bottom row) bias condition, with 95% credible intervals. (a-b) Processing speed. (c-d) Discrimination, i.e., difference between matching and mismatching accumulation rates. (e-f) Relative discrimination, i.e., percentage of the maximum possible difference between matching and mismatching accumulation rates as a proportion of the average rate. The median difference between age groups (older–younger on the left and right and younger–older for the middle column), with associated 95% credible intervals, are given at the top of each panel. See Supplementary Materials (Figure S6) for plots of the mean accumulation-rate-based parameter estimates for all within-subject conditions.

All parameters, except for non-decision time, interact to determine a further aspect of go performance, accuracy. As detailed in the Supplementary Materials, slower processing speed explained around two thirds of older participants’ accuracy advantage; greater caution and discrimination explained equal parts of the remainder, with bias playing only a minor role.

Stop Parameter Analysis

The average time for the stop runner to complete, mean stop-signal RT (SSRT), measures inhibitory ability. For the stop-signal component of the task, the leftmost panel of Figure 8 shows that older participants had longer mean SSRTs. However, because the effectiveness of the stop runner depends on its speed *relative* to the go runner, it is also important to assess SSRT as a proportion of go RT. The second panel shows that in proportional terms, it is the younger participants who have an inhibitory deficit. The model also allows for estimating the prevalence of lapses in goal maintenance (Matzke et al., 2019), i.e., the percentage of trials where either the stop (“trigger failures”) or the choice runners (“go failures”) do not enter the race. The third panel shows that stop-goal maintenance was better in older participants, although trigger failures were relatively unlikely in both groups. The rightmost panel shows that go failures were very rare, but slightly more prevalent for older participants.

Figure 8. Mean stop parameter estimates for younger and older participants (bars) with 95% credible intervals. The median difference between age groups (older–younger on the left and right and younger–older for the two middle panels), with associated 95% credible intervals, are given at the top of each panel. SSRT = stop-signal RT. Relative SSRT = SSRT divided by average go (i.e., choice) RT.

Discussion

We showed that many different aspects of executive function and processing speed can be simultaneously measured by combining a stop-signal task with difficult perceptual choices requiring the flexible exercise of proactive bias control. A new cognitive model of the task (Tanis et al., 2022) strongly supported reduced processing speed, both peripherally (in non-decision time) and centrally (in the time to complete choice and inhibitory processes), as the main cause of age-related differences. These findings are consistent with the broad array of other types of speed measures (psychophysical, perceptual, psychomotor, and psychophysiological) that display aging deficits (Salthouse, 2000). The model’s validity was supported by the accurate and detailed account of the effects of aging on all aspects of behaviour, including increased RT variability and skew effects previously attributed to poorer executive control (West et al., 2002).

Even though response inhibition processes were slowed in the elderly, this was proportionally less than slowing in their choices, suggesting that inhibitory speed is better preserved with aging than decision speed. Overall, our results were largely inconsistent with inhibitory-deficit theories, as executive control often appeared *better* for older participants. Older participants displayed less goal-neglect in the stop-signal task, and in the choice task exercised more proactive control of response bias in an absolute sense, and the same level in a relative sense. With easier choice stimuli, and in the less demanding blockwise-bias condition, selective attention was better for older than younger participants, both in an absolute and relative sense.

Selective-attention advantages for older adults accord with findings that older participants often have better task motivation than younger undergraduate samples (Ryan & Campbell, 2021). However, these advantages were reduced or eliminated in conditions imposing a higher load, due to both harder choice stimuli and trialwise biasing. This is reminiscent of early conceptions of age-related capacity limitations (Craik & Byrd, 1982), extending the idea of capacity as processing speed (Kail & Salthouse, 1994) to include another accumulation-rate-based measure, selective attention. Capacity limits also affected proactive control; for the same choice task examined here, but without stop signals (Garton et al., 2019), analysis with an alternative racing evidence-accumulation model (Brown & Heathcote, 2008), found better proactive control for older participants. Applying our model to these data (see Supplementary Materials) confirmed the previous findings, and relative to our data with the additional stop

task, proactive control for older participants doubled in an absolute sense and became credible in a relative sense. Hence, the load imposed by the stop-signal task also reduced older participants' ability to exercise superior proactive control.

Our results are inconsistent with previous findings based on the Diffusion Decision Model (DDM; Ratcliff & McKoon, 2008), which has been extensively applied to aging data (Theisen et al., 2021). Although they agree that slowing is in part caused by increased non-decision time, the DDM consistently attributes the remainder of age-related slowing to more cautious responding by older participants. A corticostriatal network has been implicated in setting threshold to determine both caution (Forstmann et al., 2010) and bias (Mulder et al., 2012). Caution-based slowing has been attributed to reduced white matter integrity in this network's tracts, causing older participants to lack the flexibility to set their thresholds appropriately (Forstmann et al., 2011).

The DDM is a single accumulation process driven by the difference between evidence for each response. Although producing optimal behaviour in some circumstances (Bogacz, 2006) neural instantiations are likely to require more than one accumulator (Ratcliff et al., 2007), and the DDM's structure means that its mean accumulation-rate parameter cannot separate processing speed from selective attention and stimulus discriminability. Hence, if processing speed reduces, that must in large part be attributed to an increase in estimated caution. Importantly, our older participants were, if anything, better at controlling their thresholds than younger participants, which is inconsistent with an explanation of slowing in terms of reduced corticostriatal-network-mediated flexibility. This illustrates how a task-model combination that simultaneously measures a variety of aspects of an individual's cognitive processing can provide a basis for inferences about the causes of observed behaviour not otherwise be available. Here, simultaneous measurement supports the processing-speed account and make an explanation of age-related slowing solely in terms of a deficit in controlling the tradeoff between speed and accuracy unlikely.

Our analysis is not subject to reliability concerns about using tasks developed in experimental psychology to identify latent constructs via correlation-based methods (Hedge et al., 2018). Although our approach does not require large samples, each participant must perform a substantial number of trials (~100 per experimental condition here), consistent with growing calls to improve data quality in this regard (Smith & Little, 2018). Given large scale correlation-based studies of aging can require up to 6 hours testing (Salthouse et al., 2003), it would be practical to extend the present approach, which required 2 hours of testing, to obtain data from the same individuals in more than one task, with the added benefit that each task measures several latent constructs.

Future research using other complex tasks, and associated cognitive models, will be required to test the generality of our findings. Fortunately, the race model framework used here is increasingly applied to a range of complex tasks involving instrumental learning, prospective memory, task switching, and decision conflicts (Heathcote & Matzke, 2022) and new cognitive-model-based analyses have been developed that address earlier concerns about how to rigorously measure capacity limitations (Boag et al., 2019; Castro et al., 2019). New versions of correlation-based approaches are also being explored that can link cognitive-model parameters from different tasks through higher-order cognitive constructs (Wall et al., 2020). This will afford researchers insights from individual-difference analyses while maintaining the reliability and validity advantages of cognitive models.

In summary, our new approach to cognitive aging has provided a strong confirmation of the processing-speed deficit hypothesis (Salthouse, 1996; Salthouse et al., 2003; Salthouse, 2000; Kail & Salthouse 1994), both through its ability to measure cognitive processing speed and disentangle its effects from the speed of peripheral processes and decision caution, and through being able to simultaneously measure the speed of inhibitory processes. Thus, the combination of a complex task with detailed cognitive modelling represents a key step forward in understanding the roles that many simultaneously interacting factors play in cognitive aging. More broadly, this detailed cognitive modelling approach has the potential to disentangle the interacting effects of different psychological processes in any domain where high quality individual data can be obtained from participants performing suitable tasks.

References

- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1-48.
- Boag, R., Strickland, L., Loft, S. & Heathcote, A. (2019). Strategic attention and decision control support prospective memory in a complex dual-task environment. *Cognition*, 191, 1-24.
- 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.
- Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, 16, 106–113.
- Braver, T., & West, R. L. (2008). Working memory, executive control, and aging. In F. I. M. Craik & T. Salthouse (Eds.), *The handbook of aging and cognition* (3rd ed., pp. 311–372). New York, NY: Psychology Press.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57, 153–178.
- Castro, S., Strayer, D., Matzke, D. & Heathcote, A. (2019). Cognitive workload measurement and modelling under divided attention. *Journal of Experimental Psychology: Human Perception & Performance*, 45, 826–839.
- Craik, F. I., & Byrd, M. (1982). Aging and cognitive deficits: The role of attentional resources. In F. I. M. Craik & S. Trehub (Eds.), *Aging and cognitive processes* (pp. 191-211). Plenum, New York.
- 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. 349-377). John Wiley & Sons, Inc.
- Donkin, C., Brown, S., & Heathcote, A. (2011). Drawing conclusions from choice response time models: A tutorial using the linear ballistic accumulator. *Journal of Mathematical Psychology*, 55(2), 140–151.
- Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., Bogacz, R., & Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences of the United States of America*, 107, 15916–15920.
- Forstmann, B. U., Tittgemeyer, M., Wagenmakers, E.-J., Derrfuss, J., Imperati, D., & Brown, S. (2011). The speed-accuracy tradeoff in the elderly brain: A structural model-based approach. *Journal of Neuroscience*, 31, 17242–17249.
- Fox, J. & Weisberg, S. (2019). *An R companion to applied regression* (3rd edition). Thousand Oaks CA: Sage.

- Garton, R., Reynolds, A., Hinder, M. R. & Heathcote, A. (2019). Equally flexible and optimal response bias in older compared to younger adults. *Psychology and Aging*, 34, 821-835.
- Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. *Psychology of Learning and Motivation*, 22, 193-225.
- Heathcote, A., Lin, Y., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2019). Dynamic models of choice. *Behavior Research Methods*, 51, 961-985.
- Heathcote, A., & Matzke, D. (accepted 3/March/2022). Winner takes all! What are race models, and why and how should psychologists use them? *Current Directions in Psychological Science*.
- Hedge, C., Powell, G., & Sumner, P. (2018). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behavior Research Methods*, 50, 1166–1186.
- Kail, R., & Salthouse, T. A. (1994). Processing speed as a mental capacity. *Acta Psychologica*, 86, 199–225.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82, 1-26.
- Matzke, D., Curley, S., Gong, C. Q., & Heathcote, A. (2019). Inhibiting responses to difficult choices. *Journal of Experimental Psychology: General*, 148, 124-142.
- 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.
- Matzke, D., Love, J. & Heathcote, A. (2017). A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behavior Research Methods*, 49, 267–281.
- Matzke, D., Verbruggen, F., & Logan, G. (2018). The stop-signal paradigm. In E.-J. Wagenmakers & J.T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience, Volume five: Methodology* (4th ed., pp. 383-427). John Wiley & Sons, Inc.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7, 134–140.
- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience*, 32, 2335–2343.
- R Core Team (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>

- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, 97, 1756–1774.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20, 873–922.
- Rey-Mermet, A., & Gade, M. (2018). Inhibition in aging: What is preserved? What declines? A meta-analysis. *Psychonomic Bulletin & Review*, 25, 1695–1716.
- Ryan, A. D., & Campbell, K. L. (2021). The ironic effect of older adults' increased task motivation: Implications for neurocognitive aging. *Psychonomic Bulletin & Review*, 28, 1743–1754.
- Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review*, 103, 403–428.
- Salthouse, T. A. (2000). Aging and measures of processing speed. *Biological Psychology*, 54, 35–54.
- Salthouse, T. A., Atkinson, T. M., & Berish, D. E. (2003). Executive functioning as a potential mediator of age-related cognitive decline in normal adults. *Journal of Experimental Psychology: General*, 132, 566–594.
- Smith, P. L., & Little, D. R. (2018). Small is beautiful: In defense of the small-N design. *Psychonomic Bulletin & Review*, 25, 2083–2101.
- Smith, P. L., & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review*, 116, 283–317.
- 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.
- Tanis, C. C., Heathcote, A., Zrubka, M., & Matzke, D. (2022). A hybrid approach to dynamic cognitive psychometrics. <https://osf.io/rqdcw>
- Theisen, M., Lerche, V., von Krause, M., & Voss, A. (2021). Age differences in diffusion model parameters: A meta-analysis. *Psychological Research*, 85, 2012–2021.
- Tillman, G., Zandt, T. V., & Logan, G. D. (2020). Sequential sampling models without random between-trial variability: The racing diffusion model of speeded decision making. *Psychonomic Bulletin & Review*, 27, 911–936.
- Wall, L., Gunawan, D., Brown, S. D., Tran, M.-N., Kohn, R., & Hawkins, G. E. (2020). Identifying relationships between cognitive processes across tasks, contexts, and time. *Behavior Research Methods*, 53, 78–95.

West, R., Murphy, K. J., Armilio, M. L., Craik, F. I. M., & Stuss, D. T. (2002). Lapses of intention and performance variability reveal age-related increases in fluctuations of executive control. *Brain and Cognition*, 49, 402–419.

Acknowledgements

Grant support: AH (DP160101891, Révész Fellowship); CT (743086-UNIFY); DM (VI.Vidi.191.09); MH and DM (DP200101696); RG and AR (University of Tasmania PhD scholarships).

Author contributions

AH, CT, and DM developed the model. AH, RG, MH, and AR designed the experiment. RG and AR collected the data. AH and DM fit the models and analysed the data. AH wrote the paper and AH and DM integrated comments from all authors.

Competing interests

The authors declare no competing interests.