The Over-Constraint of Response Time Models: Rethinking the Scaling Problem

Chris Donkin, Scott D. Brown & Andrew Heathcote

School of Psychology,
The University of Newcastle

Address correspondence to:

Chris Donkin
School of Psychology
University of Newcastle
Callaghan, NSW 2308
Australia

Phone: +61249216959

Email: chris.donkin@newcastle.edu.au

Word Counts

Abstract: 137

Body: 4231

Captions: 238

Footnotes: 1 (18 words)

References: 33 (713 words)

Figures: 3

Tables: 1

Abstract

Theories of choice response time provide insight into the psychological underpinnings of simple decisions. Evidence accumulation (or sequential sampling) models are the most successful theories of choice response time. These models all have the same ‘scaling’ property— that a subset of their parameters can be multiplied by the same amount without changing their predictions. This property means that a single parameter must be fixed to allow estimation of the remaining parameters. We show that the traditional solution to this problem has over-constrained these models, unnecessarily restricting their ability to account for data and making implicit, and therefore unexamined, psychological assumptions. We show that versions of these models which address the scaling problem in a minimal way can provide a better description of data than their over-constrained counterparts, even when increased model complexity is taken into account.

Many psychological experiments involve a choice between two alternatives. Despite their apparent simplicity, there are many complicated empirical regularities associated with the speed and accuracy of such choices. Response time (RT) distributions take on characteristic shapes, which differ systematically depending on whether the associated response is correct or incorrect and depending on any number of experimental manipulations of stimulus properties or of instructions to the participants. A range of theories have been proposed to account for both choice probability and response time when making simple decisions (for reviews see Luce, 1986; Ratcliff & Smith, 2004). Over the past 40 years, evidence accumulation (or “sequential sampling”) models have dominated the debate about the cognitive processes underlying simple decisions (e.g., Busemeyer & Townsend, 1993; Ratcliff, 1978, Ratcliff & Smith, 2004; Smith, 1995; Stone, 1960; Usher & McClelland, 2001; Van Zandt, Colonius, & Proctor, 2000).

More recently, evidence accumulation models have been applied more widely, for example, as general tools to measure cognition in the manner of psychometrics (Schmiedek, Oberauer, Wilhelm, Suß & Wittmann, 2007; Wagenmakers, van der Maas & Grasman, 2007; Vandekerckhove, Tuerlinckx, & Lee; submitted), and as models for the neurophysiology of simple decisions (e.g.: Forstmann, Dutilh, Brown, Neumann, von Cramon, Ridderinkhof, & Wagenmakers, 2008; Ho, Brown, & Serences, 2009; Smith & Ratcliff, 2004). In light of this growing influence it is especially important that users of these models are not misled by implicit - and hence unexamined - assumptions.

Evidence accumulation models all share a basic framework wherein, when making a decision, people repeatedly sample evidence from the stimulus. This evidence is accumulated until a threshold amount is reached, which triggers a decision response. These models naturally predict the response made (depending on which response has accumulated the most evidence) and the latency of the response (depending on how long the evidence took to

accumulate). We illustrate these models using the example of a lexical decision task, where a participant must decide whether a string of letters is a valid word (e.g. “DOG”) or not (e.g. “DXG”). The participant samples information from the stimulus repeatedly, and finds some evidence that suggests that the stimulus is a word and other evidence to suggest that the stimulus is not a word. The participant accrues this information, waiting until there is enough evidence for one of the two options before responding. Their choice corresponds to the response with the most evidence, and the time taken for this evidence to be accumulated is the response latency.

Over the past four or five decades, dozens of evidence accumulation models have been proposed, and all of them share a mathematical “scaling property”: one can multiply a subset of their parameters by an arbitrary amount, without changing any of the model’s predictions. To avoid complications arising from the scaling property, just one parameter of the model must be constrained arbitrarily. We show that the conventional approaches – which have been universally applied to solve the scaling problem – have actually over-constrained the models by fixing more than one parameter. This over-constraint has been largely unrecognized by the field, and so it is equivalent to making a tacit, untested, psychological assumption. Further, we show that this tacit assumption can sometimes have important consequences: when the scaling problem is solved in a minimal way, the models can sometimes provide better account for data.

Overview of the models

There are two major classes of evidence accumulation models: single accumulator models (Busemeyer & Townsend, 1993; Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff & Tuerlinckx, 2002; Smith, 1995; Stone, 1960) and models that have one accumulator for each possible response (Brown & Heathcote, 2005, 2008; Smith & Van Zandt, 2000; Smith & Vickers, 1988; Townsend & Ashby, 1983; Usher & McClelland, 2001; Van Zandt, Colonius

& Proctor, 2000; Vickers, 1970). The customary method for solving the scaling problem differs between the two classes of models, even though the principle is the same. To simplify our discussion we choose a specific model from each class: the single accumulator diffusion model (Ratcliff & Tuerlinckx, 2002) and the multiple accumulator linear ballistic model (LBA; Brown & Heathcote, 2008). We have chosen these two models largely for convenience, as both have easy-to-implement computer code that is freely available (see Donkin, Averell, Brown & Heathcote, in press, and Voss & Voss, 2007). The general point that we make, however, applies to all evidence accumulation models.

[Insert Figure 1 about here]

Continuing the lexical decision example, the diffusion model assumes that participants sample evidence from the stimulus continuously, and this evidence stream updates an evidence total, say x , illustrated as a function of time by the irregular line in Figure 1. The accumulator begins the decision process in some intermediate state, say $x=z$. Evidence that favors the response “word” decreases the value of x , and evidence that favors the other response (“non-word”) increases the value of x . The evidence accumulation process continues until sufficient evidence favors one response over the other, causing the total (x) to reach one of its two boundaries (the horizontal lines at $x=0$ and $x=a$ in Figure 1). The choice made by the model depends on which boundary is reached (a for a “non-word” response or 0 for a “word” response) and response time equals the accumulation time plus a constant, T_{er} , that represents the time taken by non-decision processes, such as encoding the stimulus and producing the response.

Depending on the stimulus, evidence tends to accumulate more towards one boundary or another, and the average rate of this accumulation is called the “drift rate”, which we will label v . The evidence accumulation process also varies randomly from moment-to-moment during the accumulation process, and the amount of this variability is another parameter of the

model, s . Recent applications of the diffusion model include three extra parameters, but these are not important for our purposes, so we delay their introduction until later. When experimental conditions differ only in stimulus characteristics that vary randomly from trial to trial, all parameters except the drift rate are conventionally assumed constant over conditions (Ratcliff, Gomez & McKoon, 2004).

The LBA is a multiple accumulator model, meaning that it assigns a separate evidence accumulator to each possible response: for example, in lexical decision, one accumulator gathers evidence in favor of the response “word” and the other gathers evidence for the “non-word” response, as illustrated in Figure 1. The activity level in each accumulator begins at a value that is randomly sampled (separately for each accumulator) from the interval $[0, A]$. Evidence accumulation is noiseless (“ballistic”) and linear with a slope that we again call the “drift rate”, v . When the evidence accumulated for either response reaches a threshold, b , a response is made. Like the diffusion model, the LBA assumes that non-decision processing takes fixed time, T_{er} . The drift rates are assumed to vary from trial to trial according to normal distributions with means v_W for the “word” accumulator and v_{NW} for the “non word” accumulator, and a common standard deviation, s .

Scaling Properties

Consider just one of the evidence accumulators from the LBA. The accumulator begins a trial with some activity, say x_0 , between 0 and A , and increases at a rate of v units per second (v is the drift rate for this accumulator). Evidence accumulation ends when the threshold b is reached, which will take $(b-x_0)/v$ seconds. If all of these model parameters were multiplied by a common amount the predicted response time would remain unchanged; for example, if the parameters were doubled then the evidence accumulation process would travel twice as quickly, but would also have to travel twice as far. This scaling property is true of all

evidence accumulation models – all parameters that affect evidence accumulation can be multiplied by any fixed amount without altering the model's predictions.

The scaling property makes it impossible to estimate unique model parameters from data unless the value of one parameter is fixed arbitrarily. In single accumulator models, including Ratcliff's diffusion, this has always been done by fixing the variability of the diffusion process at either $s=0.1$ or $s=1$ (e.g., Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff & Tuerlinckx, 2002; Smith & Ratcliff, 2004; Van Zandt, Colonius & Proctor, 2000; Voss, Rothermund, & Voss, 2004). In fact, the diffusion coefficient is usually referred to as *the* “scaling parameter” of the model, even though any other parameter could equally well be fixed to avoid scaling problems.

For the LBA and other multiple accumulator models, problems due to the scaling property have been avoided by fixing the sum of the drift rates for the two accumulators to a constant (Brown & Heathcote, 2005, 2008; Forstmann, et al., 2008; Ho, et al., 2009; Ratcliff & Smith, 2004; Smith & Van Zandt, 2000; Townsend & Ashby, 1983; Usher & McClelland, 2001). For example, in the word vs. non word model above, one might fix $v_W + v_{NW} = 1$. Mathematically speaking, any other parameter constraint would do just as well to solve the scaling problem, for example the boundary separation, or one of the drift rates, could be fixed. It is simply a matter of convention that the field has settled on the sum-of-drift-rates constraint for multiple accumulator models, and the diffusion noise constraint for single accumulator models.

The scaling properties just described are simple and well-understood. However, the situation is complicated in practice because evidence accumulation models are almost never used to analyze just one experimental condition in isolation. Instead, data are collected from multiple experimental conditions, which are analyzed together¹. This allows some parameters to be fixed across experimental conditions, depending on what psychological assumptions one

is willing to make. For example, when experimental conditions differing only in stimulus properties are randomly ordered from trial to trial, parameters that are assumed to be under the strategic control of the participant (such as boundary separation) are often fixed. This is justified by the notion that such parameters take time and effort to change, and hence that such changes are unlikely to occur between stimulus onset and the response (Ratcliff, 1978).

When parameters are fixed across conditions, changes in parameters for one condition naturally alter the predictions for other conditions. So, when multiple conditions are analyzed simultaneously, the scaling properties of the models can be constrained by fixing a single parameter in *only one* of the conditions. For example, suppose one conducts an experiment with five different levels of difficulty defined by different stimuli. In this design, scaling problems are avoided if a single parameter is constrained in only one of the five conditions. If Ratcliff's diffusion model is used, one can set $s=1$ in just one of the five conditions, or in the LBA, one could set the sum of the drift rates equal to one in just one of the five conditions.

However, these are not the constraints that have been used in practice. In all of the studies that we have reviewed (including our own), researchers have constrained "scaling parameters" independently in *all* conditions. To continue our example, they have either fixed s for all five conditions (in single accumulator models, like the diffusion) or fixed the sum of the drift rates for all five conditions (in multiple accumulator models). Avoiding estimation problems due to scaling properties requires fixing just one parameter value, but researchers have always fixed one parameter value *per experimental condition*.

These so-called 'scaling parameters' (i.e. s or the sum of drift rates) have also come to be treated quite differently from the other model parameters, as "fixed, not free" (p. 440, Ratcliff & Tuerlinckx, 2002). Mathematically, dealing with scaling parameters entails two independent decisions – firstly, one model parameter is arbitrarily selected for constraint, and secondly that parameter can be (and always has been) held to its fixed value across *all*

experimental conditions. The first decision has no theoretical consequence, because constraints on different types of parameters are mathematically equivalent. The second decision can have theoretical consequence. Suppose, for example, that one had decided to fix the boundary separation parameter to solve the scaling problem. Keeping this parameter fixed across conditions that differ in instructions that emphasize the speed or accuracy of responses is clearly inappropriate. In practice, however, it seems that the second decision has automatically followed the first; that choosing a scaling parameter has pre-determined its role as fixed across experimental conditions.

Our central message is that the parameter type chosen as a scaling parameter *can* vary across conditions while still satisfying the scaling constraint property. As we will now demonstrate, this can be important since using minimal constraints changes two things: the predictions and the psychological interpretations of the models. By re-analyzing previously published data we show that allowing parameters previously treated in the conventional scaling manner to vary across conditions can have a substantial, and sometimes useful, effect. We show that when a minimal rather than conventional scaling solution is used, the resulting improvement in fit is sufficient to justify the extra parametric freedom, even when a very strict complexity penalty is employed. We then briefly address the psychological implications of our proposal.

Re-analysis of Gould, Wolfgang and Smith's (2007) data

Gould, Wolfgang and Smith (2007) investigated the effect of cueing and localization in a stimulus detection task. We focus on one of their cueing conditions, the one providing the greatest challenge for evidence accumulation models. Gould et al. manipulated the difficulty of stimulus detection by varying contrast over five levels. This resulted in ten different RT distributions – one for correct responses and one for incorrect responses from each of the five contrast levels. Figure 2 summarizes these ten distributions using quantile probability (QP)

plots, with data (averaged across subjects) represented as filled circles connected by solid lines. QP plots have proven very important in discriminating between models of choice RT (Brown & Heathcote, 2008; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004). The x-axis measures response probability, and the y-axis shows the latencies associated with five quantiles of the RT distributions (10%, 30%, 50% - the median, 70% and 90%).

An example may make this clearer; consider just the data from the very easiest stimulus condition. Average response accuracy for this condition was 97.3%, so the quantile estimates from the distribution of correct responses are plotted as five filled circles vertically above $x=.973$, and the quantile estimates from the distribution of incorrect responses are plotted above $x=.027$ (i.e., $1-.973$). The 10% quantile estimate for the correct RT distribution was 409msec (i.e., 10% of correct responses were faster than 409msec), so the first filled circle above $x=.973$ is at $y=409$. The procedure is repeated for the remaining quantile estimates for both the correct and incorrect distributions. The resulting plot allows one to assess how the RT distribution changes with response accuracy and between correct and incorrect responses. The solid lines in Figure 2 join each of the quantiles across contrast values for correct and error responses.

[Insert Figure 2 about here]

As is typical, correct responses for difficult decisions were slower than easy decisions (as we move from the right to the center of the plots) at all five quantiles. Incorrect responses were slower still (on the left half of the plots) with the possible exception of errors in easiest condition (far left). What is unusual in these data is the amount of change in the fastest response times (the 10% quantile estimate). The 10% quantile estimate for correct responses was 87 ms slower in hardest condition relative to the easiest condition. Most previous studies have observed that the fastest decision times change by at most 30ms across the range of a QP plot (e.g., Ratcliff et al., 2004; Ratcliff & Smith, 2004).

Model fits – Conventional scaling

Figure 2 shows fits of the diffusion and LBA models when employing the over-constrained conventional solution to the scaling problem (i.e., fixing one parameter across all experimental conditions). To fit both models we followed the usual convention of assuming that only drift rate varies between conditions, and that all other parameters were equal across all conditions. A constant non-decision time (T_{er}) assumes encoding and response production time does not vary across conditions. Constancy of the strategic parameters is justified because stimulus conditions varied randomly from trial to trial. We also assumed that responding was unbiased; for the diffusion model this means evidence accumulation begins half way between the bounds ($z=a/2$), and that response boundaries are equal in the LBA.

The simple diffusion model described earlier would fit these data with five free parameters. However, in practice the diffusion model uses another three free parameter types related to trial-to-trial variation: the starting point of the evidence accumulation process varies according to a uniform distribution on $[z-s_z, z+s_z]$, the drift rate varies according to a normal distribution $N(v, \eta)$, and non-decision time (T_{er}) varies according to a uniform distribution on $[T_{er}-s_T, T_{er}+s_T]$. These additions make a total of ten free parameters ($T_{er}, s_T, a, \eta, s_z, v_1, v_2, v_3, v_4, v_5$), as they are assumed the same across conditions. The LBA fits also used ten free parameters for each participant ($T_{er}, s_T, A, b, s, v_1, v_2, v_3, v_4, v_5$). This is one more parameter than has previously been used in applications of the LBA, as non-decision time, T_{er} , is usually assumed fixed (i.e., $s_T=0$). We allowed non-decision time to vary here for equivalence with the diffusion model (although s_T was estimated at about zero for the LBA). The sum of average correct and incorrect drift rates in the LBA was set at 1 for all five stimulus contrast levels. The diffusion model was constrained by having the diffusion coefficient fixed at $s=1$ across all five conditions.

Parameters were estimated using the method of quantile maximum probabilities (Heathcote & Brown, 2004). Model predictions were evaluated using the LBA code provided by Donkin et al. (in press) and the diffusion model code provided by Voss and Voss (2007). The Bayesian information criterion (BIC) was calculated at the best-fitting parameters for each participant: the BIC statistic for N observations grouped into bins is:

$$BIC = -2(\sum_i Np_i \ln(\pi_i)) + M \ln(N)$$

where p_i is the proportion of observations in the i^{th} bin, and π_i is the proportion of observations in the i^{th} bin as predicted by the model. M is the number of parameters of the model used to generate predictions. The BIC is composed of two parts, the first is a measure of misfit, and a second part, $M \ln(N)$, penalizes a model for its complexity as indicated by the number of estimated parameters. When comparing two models, the model with the smaller BIC is considered to have provided a better fit to the data, after complexity has been taken into account. We use BIC because it imposes a larger complexity penalty than alternatives such as the Akaike Information Criterion (AIC), and so it provides a more stringent test of whether the models benefit from the extra parameter variation allowed by imposing minimal constraints to solve the scaling problem.

[Insert Table 1 about here]

Parameter estimates and BIC values are shown in Table 1 (we focus on averaged data for brevity). As Figure 2 shows, both models provided poor accounts of the data – the diffusion under predicts the shift in RT distribution across conditions, while the LBA fails to capture the faster errors that occur in easy conditions. The conclusion we draw is that standard applications of both models fail to provide convincing accounts of these data.

Model Fits – Minimally Constrained

For the minimally constrained version of Ratcliff's diffusion model, we fixed the diffusion coefficient in the highest contrast condition at $s_I=1$ and freely estimated diffusion

coefficients for the other four contrast conditions (s_2, s_3, s_4, s_5). For the minimally constrained version of the LBA, we fixed the sum of correct and error drift rates to be one in the easiest condition, and estimated this sum in the other four conditions (i.e. $\Sigma v_2, \Sigma v_3, \Sigma v_4, \Sigma v_5$). Table 1 reports the estimated parameters and the BIC values for the minimally constrained fits, shown in Figure 3. The quality of fit was greatly improved, with both models providing a much better account of the data than the conventionally constrained versions.

[Insert Figure 3 about here]

BIC values were better in the minimally constrained versions of both models, suggesting that the improvement in fit outweighed the cost of adding four additional parameters. Using methods outlined by Wagenmakers and Farrell (2004), BIC values can be converted to model selection probabilities (see Raftery, 1995, for a discussion of conventions for interpreting such probabilities). The BIC improvement provided very strong evidence ($p > .99$, Raftery, 1995) favoring both minimally constrained models over their conventionally constrained counterparts. The improvement in the diffusion model seems to have come from predicting a larger shift in RT distribution across conditions and no longer predicting such extreme skewness for difficult decisions. The minimally constrained LBA was better able to accommodate the fast errors. As before, estimated drift rates for both models decrease in a sensible manner with decreasing stimulus contrast (Table 1). For both models, the scaling parameter also decreased with decreasing stimulus contrast (the sum of the drift rates for the correct and incorrect response accumulators in the LBA, and the diffusion variability coefficient in the diffusion model).

Discussion

All evidence accumulation models require a “scaling property” to be fixed, before parameters can be estimated. To this end, researchers must choose a parameter to constrain, but this choice is logically independent of the subsequent decision of whether to further

constrain that parameter across experimental conditions. In practice, however, these two decisions have never been separated – the parameter chosen to satisfy the scaling property has always also been constrained across experimental conditions. This is a nontrivial assumption, because the scaling parameters of the models could plausibly be driven by stimulus characteristics which often differ between conditions. A re-analysis of one such case, from Gould et al. (2007), showed that separating these two decisions was justified by improved fits to data for both the LBA and diffusion models, even allowing for a very stringent model complexity penalty.

For multiple accumulator models, such as the LBA, the assumption of a constant sum for correct and incorrect drift rates across conditions implies that increasing the stimulus evidence in favor of one response will equally increase the evidence against the other response. However, it seems reasonable that some stimulus manipulations could decrease the evidence available for *both* responses. Contrast is plausibly one such manipulation: as contrast decreases there may be less evidence supporting either response, and our parameter estimates from Gould et al.'s (2007) data were consistent with this interpretation. For single accumulator models, such as Ratcliff's diffusion, the conventional constraints imply that the variability in evidence accumulation is independent of the mean rate of accumulation. This assumption might be reasonable if, for example, the decision signal arises from one set of processes whereas *all* decision noise arises from an independent set of processes. However, our parameter estimates suggest that the diffusion model may better account for the effects of decreasing stimulus contrast by assuming some dependence between decision signal and decision noise.

Although we have focussed on the diffusion and LBA models, the same arguments apply to all evidence accumulation models. Other multiple accumulator models have been similarly over-constrained, particularly the many variants of Usher and McClelland's (2001)

leaky competing accumulator model, including the racing diffusion of Ratcliff, Cherian and Segraves (2003) and the ballistic accumulator (Brown & Heathcote, 2005). The Poisson counter models (Smith & Van Zandt, 2000; Townsend & Ashby, 1983; Ratcliff & Smith, 2004; Van Zandt, Colonius & Proctor, 2000) have all been similarly over constrained by their own conventional solutions for the scaling problem.

General Discussion

The way in which the parameters of evidence accumulation models are constrained across conditions is based on careful argument and empirical evidence. For example, Ratcliff (1978) proposed that strategic parameters (e.g., boundary separation) should not differ among conditions whose order is randomized within blocks of trials. In contrast, parameters related to the quality of evidence provided by the stimulus (e.g., drift rate) should vary whenever stimulus properties change (see also: Ratcliff & Rouder, 1998; Voss et al., 2004). However, the “scaling parameters” of evidence accumulation models have always been fixed across all conditions, even though these parameters may most naturally be interpreted as ones influenced by stimulus properties. A review of the literature reveals neither careful argument nor empirical evidence to justify this extra constraint – it appears to have been a result of a misunderstanding scaling property (this is certainly true on our own part). Our results show that this over-constraint may not have always been benign; it can restrict the models’ ability to account for data, and it makes implicit psychological assumptions.

It is possible that experts in the field may have been aware of this additional assumption being made when scaling parameters were fixed across conditions. For example, more than 25 years ago, Weatherburn (1978), Pike and Dalgleish (1982) and Weatherburn and Grayson (1982) discussed whether or not scaling parameters might vary in earlier instances of multiple accumulator models. This discussion, however, did not include actually trying out such models, and our literature review suggests that the implications of their

discussion have since gone unrecognized (we could find no citations of these papers in the past 14 years). Our aim is to ensure that the ever-expanding group of researchers who use response time models are aware of the implicit assumptions made when fixing the scaling parameter constant across conditions.

References

- Brown, S., & Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Review*, 112, 117-128.
- Brown, S.D., & Heathcote, A.J. (2008). The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology*, 57, 153-178.
- Bussemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100, 432– 459.
- Donkin, C., Averell, L., Brown, S.D., & Heathcote, A. (in press). Getting more from accuracy and response time data: Methods for fitting the Linear Ballistic Accumulator model. *Behavior Research Methods*.
- Forstmann, B.U., Dutilh, G., Brown, S.D., Neumann, J., von Cramon, D.Y., Ridderinkhof, K.R., & Wagenmakers, E.J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Science* 105, 17538-17542.
- Gould, I. C., Wolfgang, B. J., & Smith, P. L. (2007). Spatial uncertainty explains exogenous and endogenous attentional cuing effects in visual signal detection. *Journal of Vision*. 7(13):4, 1-17.
- Heathcote, A. & Brown, S.D. (2004). Reply to Speckman and Rouder: A theoretical basis for QML. *Psychonomic Bulletin & Review*, 11, 577-578.
- Ho, T.C., Brown, S.D. & Serences, J.T. (2009). Domain general mechanisms of perceptual decision making in human cortex. *Journal of Neuroscience*.
- Luce, R. D. (1986). *Response times*. New York: Oxford University Press.
- Pike, A.R., & Dalgleish, L. (1982). Latency-probability curves for sequential decision models: A comment on Weatherburn. *Psychological Bulletin*, 91, 384-388.

- Raftery, A. (1995). Bayesian model selection in social research. *Sociological Methodology*, 25, 111-163.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59–108.
- Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two choice decisions. *Journal of Neurophysiology*, 90, 1392-1407.
- Ratcliff, R., Gomez, P. & McKoon, P. (2004). A diffusion model account of the lexical decision task, *Psychological Review*, 111, 159-182.
- Ratcliff, R., & Rouder, J. N. (1998). Modelling response times for two choice decisions. *Psychological Science*, 9, 347–356.
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111, 333–367.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating the parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9, 438–481.
- Schmiedek, F., Oberauer, K., Wilhelm, O., Suß, H.M. & Wittmann, W.W. (2007). Individual differences in components of reaction time distributions and their relations to working memory and intelligence, *Journal of Experimental Psychology: General*, 136, 414-429.
- Smith, P. L. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review*, 102, 567–591.
- Smith, P.L., & Ratcliff, R. (2004). Psychology and Neurobiology of Simple Decisions. *Trends in Neurosciences*, 27, 161-168.
- Smith, P. L., & Van Zandt, T. (2000). Time-dependent Poisson counter models of response latency in simple judgment. *British Journal of Mathematical and Statistical Psychology*, 53, 293–315.

- Smith, P.L., & Vickers, D. (1988). The accumulator model of two-choice discrimination. *Journal of Mathematical Psychology*, 32, 135-168.
- Stone M. (1960). Models for choice-reaction time. *Psychometrika*, 25, 251–260.
- Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. Cambridge, England: Cambridge University Press.
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 108, 550–592.
- Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, 7, 208–256.
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (submitted). Hierarchical diffusion models for two-choice response times. *Psychological Methods*.
- Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, 13, 37–58.
- Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, 32, 1206-1220.
- Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavioral Research Methods*, 39, 767–775.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192-196.
- Wagenmakers, E.-J., van der Maas, H.L.J., & Grasman, R.P.P.P. (2007). An EZ-diffusion model for response time and accuracy. *Psychonomic Bulletin & Review* 4, 3-22.
- Weatherburn, D. (1978). Latency-Probability functions as bases for evaluating competing accounts of the sensory decision process. *Psychological Bulletin*, 85, 1344-1347.
- Weatherburn, D., & Grayson, D. (1982). Latency-Probability functions: A reply to Pike and Dalgleish. *Psychological Bulletin*, 91, 389-292.

Footnotes

¹ Applications of Wagenmakers et al.'s (2007) EZ estimation technique, such as by Schmiedek et al. (2007), are an exception.

Table Captions

Table 1 Parameter estimates and BIC from fits to average data from Gould et al.'s (2007) cued+FID condition (D = diffusion model, L= LBA model)

Tables

	T_{er}^a	s_T^a	a	η^b	s_z	v_1^b	v_2^b	v_3^b	v_4^b	v_5^b	s_2	s_3	s_4	s_5	BIC
D	.352	.083	1.46	1.76	.000	5.17	4.01	2.61	1.12	0.37	-	-	-	-	39406
	.364	.099	1.26	1.59	.046	4.78	3.68	2.36	1.12	.378	1.01	.923	.843	.788	39305
	T_{er}^a	s_T^a	b	s^b	A	v_I^b	v_2^b	v_3^b	v_4^b	v_5^b	$\sum v_2^b$	$\sum v_3^b$	$\sum v_4^b$	$\sum v_5^b$	BIC
L	.144	.000	.356	.276	.047	.974	.881	.757	.621	.538	-	-	-	-	39403
	.168	.002	.287	.219	.082	.797	.720	.609	.475	.378	.923	.800	.692	.648	39166

Note: ^a indicates parameters whose units are in ‘seconds’. ^b have units ‘per second’, while other parameters have arbitrary units.

Figure Captions

Figure 1 Graphical representations of a single decision made by the diffusion model and the LBA model.

Figure 2 Quantile probability plots for Gould et al.'s (2007) data and fits (averaged across participants). The data are shown as filled circles joined by solid lines. In the left and right panels the dotted lines are the fits of the diffusion and LBA models, respectively.

Figure 3 Quantile probability plots and fits averaged over participants for the minimally constrained versions of the diffusion model (left panel), the LBA (right panel) for Gould et al.'s (2007) data. The data are shown as filled circles and solid lines.

Figures

Figure 1

Figure 2

Figure 3