AN ACCUMULATION OF PREFERENCE: TWO ALTERNATIVE DYNAMIC MODELS FOR UNDERSTANDING TRANSPORT CHOICES

Thomas O. Hancock (Corresponding Author)
Choice Modelling Centre and Institute for Transport Studies, 34-40 University Road, University of Leeds, Leeds, LS2 9JT, United Kingdom.
tratoh@leeds.ac.uk

Stephane Hess
Choice Modelling Centre and Institute for Transport Studies, University of Leeds.
S.Hess@leeds.ac.uk

Anthony A. J. Marley
Department of Psychology, University of Victoria, Victoria, Canada.
Institute for Choice, University of South Australia, Adelaide, Australia.
ajmarley@uvic.ca

Charisma F. Choudhury
Choice Modelling Centre and Institute for Transport Studies, University of Leeds.
C.F.Choudhury@leeds.ac.uk

ABSTRACT

Interest in behavioural realism has gradually led to the introduction of alternatives to random utility maximisation (RUM) as a paradigm for discrete choice models, with notable interest for example in random regret minimisation (RRM). These models have however continued to rely on a framework where a single *value function* of some form is calculated once for each alternative in each choice setting, and the choice probabilities are calculated by comparing these value functions across alternatives. In contrast, research in mathematical psychology has used a more dynamic approach, where the preference value of each alternative updates over time in a given situation while the decision maker is deliberating about the choice to make. These *accumulator* models are well suited to accommodating a variety of context effects, and have been shown to give good performance for data collected in laboratory-based settings. The present paper considers two such accumulator models, namely decision field theory (DFT) and the multi-attribute linear ballistic accumulator (MLBA), and makes a number of methodological improvements to address limitations that have thus far prevented their use in travel behaviour research. This includes the ability to capture the influence of socio-demographics, the presence of underlying preferences for specific alternatives, or dealing with attributes that have opposite effects on choice probabilities. We offer what we believe to be the first in-depth simultaneous comparison of DFT and MLBA with typical discrete choice models, and also for the first time test both DFT and MLBA on a revealed preference dataset. We find that both models outperform typical RUM and RRM implementations for both estimation and out-of-sample prediction across our datasets, including in a large scale simulation experiment.

1. INTRODUCTION

Whilst mainstream choice modelling has been grounded in firm economic foundations [McFadden, 1974], the modelling of decision-making behaviour in other fields has been implemented with very different aims and objectives. Since work in the 1970s [Tversky, 1972, Tversky and Kahneman, 1973, Tversky, 1977], the field of behavioural economics has considered choice from an economic viewpoint whilst simultaneously demonstrating that decision-makers are subject to biases, heuristics and context effects that result in choices being made that are not the most likely under traditional choice models. Choice modellers have long had an interest in increasing the behavioural realism of their models, with recent methodological advances aimed at incorporating alternative behavioural ideas such as random regret minimisation (RRM) [Chorus et al., 2008, Chorus, 2010], heuristics [Swait, 2001] and satisficing [González-Valdés and Ortúzar, 2017].

Moving away from the traditional random utility maximisation (RUM) framework however entails a number of disadvantages, notably an inability to perform welfare analysis [Hess et al., 2018]. This means that careful consideration is required before we move to alternative models. In this context, the question then arises whether, if we are willing to move away from RUM, we should move to models that are substantially different from it, rather than still staying within a logit framework as is the case for random regret minimisation. This observation leads us to look further afield in the present paper, and in particular at the work in mathematical psychology, where researchers have tended to try and build models to mathematically represent context effects such as the attraction, compromise and similarity effects [Roe et al., 2001, Trueblood et al., 2013b, Noguchi and Stewart, 2014] as well as decision-making under time pressure [Busemeyer and Townsend, 1993].

It is notable that very few papers as yet have tested whether models developed in mathematical psychology can be used for predicting choices in general (i.e. outside laboratory settings). Some notable exceptions include Hawkins et al. [2014] who applied the linear ballistic accumulator (LBA, Brown and Heathcote 2008) to consumer attitudes and patient preferences and Berkowitsch et al. [2014], who applied decision field theory (DFT) to consumer choices for products such as computers, cameras and racing bicycles. In key comparisons against *traditional* choice models, DFT in particular has been found to outperform random utility and random regret based models (e.g. Berkowitsch et al. [2014], Hancock et al. [2018]). Other models from mathematical psychology are yet to be put to the test in such a rigorous manner.

DFT and the similarly popular (in mathematical psychology) multi-attribute linear ballistic accumulator (MLBA) [Trueblood et al., 2013a, 2014] differ from more traditional discrete choice models in one specific dimension. RUM and RRM models are characterised by their utility and regret functions respectively, which are used to calculate a single *value function* for each alternative, where comparison of this across alternatives then leads to probabilities of a given alternative being chosen. This value function is calculated once per choice situation. On the other hand, DFT and MLBA are members of a broad family of *accumulator* models, where the preference values for an alternative in a single choice context are not static but are updated over time. It is important to note that this is different from work looking at preferences evolving over a sequence of choices, such as models incorporating value learning [McNair et al., 2012], state dependence [Bruno et al., 2015] or dynamic discrete choice models [Liu and Cirillo, 2018]. *Accumulator* models are structures for internal preference accumulation at the level of every single choice, not models that accumulate evidence over a sequence of choices. The accumulation models thus capture the mental deliberation from the time a particular choice is faced (or stated choice scenario is presented) to the point where the choice is made. The preferences are reset after that, so that the accumulation effect is not carried over to the next choice task, i.e. the accumulation made for choice t does not affect choice $t+1$ although such extensions are possible too.

Under DFT, the decision maker updates his/her preference for given alternatives by repeated comparisons between them, considering one attribute at a time, where the attribute values of the alternatives in that situation remain constant across these comparisons. Under MLBA, a 'drift rate' is generated at the outset for each alternative allowing the preference values to update within a single choice context. Thus far, there has, to the best of our knowledge, not been any application of MLBA to transport data and only a few, mainly theoretical, applications of DFT. The way in which preferences evolve over time and their inherent ability to accommodate a range of what economists might call behavioural anomalies however make these models at first hand very appealing for studying travel behaviour.

The large and rich datasets typically found in transport have meant that computational limitations have until now limited the use of DFT in transport applications [Otter et al., 2008]. Our previous work on DFT has focused on methodological improvements that have made it possible to rigorously test DFT against typical choice models [Hancock et al., 2018]. This motivates us to investigate the suitability of MLBA in modelling travel behaviour as well, as it has been found to outperform DFT in applications in the mathematical psychology literature [Trueblood et al., 2014, Cohen et al., 2017, Turner et al., 2018][1].

Beyond simply comparing the two structures, we make a number of methodological improvements to both DFT and MLBA to facilitate their application to rich multi-alternative multi-attribute datasets. The key contribution relates to allowing analysts to use DFT with attributes that have opposite effects on choice probabilities, and where this directionality is not known a priori. Previously, DFT models included 'attention weights' which could be used to capture the relative importance of attributes. As these weights must be positive (and sum to one), a priori knowledge is required as to whether an attribute has a positive (e.g. comfort of journey) or negative (e.g. cost) impact on the likelihood of an alternative being chosen. This is particularly an issue for consumer attributes which some decision-makers may like and others dislike, such as the size of a car. We propose the use of attribute-specific scaling coefficients, meaning that such a priori knowledge is no longer required. We show that these coefficients can also be added to MLBA to capture the relative importance of different attributes, a feature not typically accounted for in standard MLBA implementations. Further improvements include the ability to capture the influence of socio-demographics and the presence of underlying preferences for specific alternatives, in a manner equivalent to alternative specific constants in typical discrete choice models. We also look in detail at identification issues for both models, with a number of empirical tests to help inform

[1]Furthermore, there has been increasing attention in transport on best-worst datasets [Giergiczny et al., 2013, Rose, 2014] and research in mathematical psychology has shown that the linear ballistic accumulator (a simpler form of the model, where each alternative has a mean drift rate simply equal to an alternative-specific constant), performs well for these datasets [Hawkins et al., 2014].

future applications.

In our empirical work, we offer what we believe to be the first in-depth simultaneous comparison of DFT and MLBA with typical discrete choice models, and also for the first time test both DFT and MLBA on a revealed preference dataset. We find that both models outperform typical RUM and RRM implementations for both estimation and out-of-sample prediction across our datasets, including in a large scale simulation experiment.

The remainder of this paper is organised as follows. In the next section, we first provide an overview of the two models in their current form before presenting our various methodological improvements. This is followed by our empirical work on stated choice and revealed preference data, before some further tests on simulated data. The final section summarises the findings and presents some directions for future research.

2. METHODOLOGY: CONTRASTING AND IMPROVING MODELS FROM MATHEMATICAL PSYCHOLOGY

In this section, we first provide an introduction to accumulator models. We then present state-of-the-art implementations before making methodological improvements for both DFT and MLBA.

2.1. Introduction to accumulator models

Since the introduction of the drift diffusion model [Ratcliff, 1978], many different variations of sequential sampling models (or accumulation models) have been developed by mathematical psychologists [Busemeyer and Townsend, 1992, Usher and McClelland, 2001, Krajbich et al., 2012]. The idea of a sequential sampling model is that preferences for alternatives update over time depending on what information is being considered. An individual may consider, for example, cost, before then considering travel time. They might make comparisons across alternatives sequentially or randomly. By contrast, mainstream choice models such as random utility or random regret models construct just a single one-off preference or utility value for each alternative given a set of attribute values and then use that value to calculate choice probabilities. Critically, accumulation models instead assume that these preferences change over the course of the deliberation process whilst the decision-maker is choosing an alternative (even if the attributes of the alternatives stay the same). As already highlighted in the introduction, this preference accumulation is internal and happens at the level of every single choice, i.e. it is not an accumulation over a sequence of choices. This thus allows us to contrast the models with typical discrete choice structures.

Accumulator models aim to 'understand the motivational and cognitive mechanisms that guide the deliberation process involved in decisions' [Busemeyer and Townsend, 1993]. These models have subsequently been shown to resemble neural activity. For example, Gold and Shadlen [2000] found that during a motion perception task, there was an accumulation of sensory evidence in the neural circuits of a monkey's brain, creating a behavioural response when the appropriate amount of information had been received. Furthermore, accumulator models have been demonstrated to predict contextual effects [Hotaling et al., 2010, Trueblood et al., 2014], capture risky choice behaviour [Busemeyer and Townsend, 1993, Stewart and Simpson, 2008] and can predict preference reversals [Diederich, 2003]. Additionally, dynamic models provide a naturalistic method for the modelling of decision making in dynamic choice settings [Holmes et al., 2016].

One popular model from mathematical psychology that can easily be compared to traditional choice models is decision field theory (DFT), first introduced by Busemeyer and Townsend [1992, 1993] and first operationalised in the context of travel behaviour by Hancock et al. [2018]. In a DFT model, preference values for the alternatives update stochastically over time. At each moment, a single attribute is compared across alternatives and a valence (momentary preference) is added to the preference value for each alternative. At some point, the decision-maker comes to a conclusion, either as one of the alternatives reaches some threshold (similar to satisficing [Kaufman, 1990, Schwartz et al., 2002, González-Valdés and Ortúzar, 2017]) or as an external cue forces the decision-maker to make a choice, in which case the decision-maker chooses the alternative with the highest preference value at that moment. As an example, the left panel in Figure 1 demonstrates that different alternatives may be chosen depending on which threshold is used. The first alternative to reach the internal threshold value is outperformed by another alternative if the decision is not made until the time threshold applies. Here, it should be noted that the value that evolves over comparison is a *preference value*, rather than a probability, where the latter is calculated from the expectation of the former, a point we will return to later. The horizontal axis is measured in *timesteps*, which relate to the number of comparisons between alternatives, each time using one attribute.

FIGURE 1 : An example decision process under both accumulation models

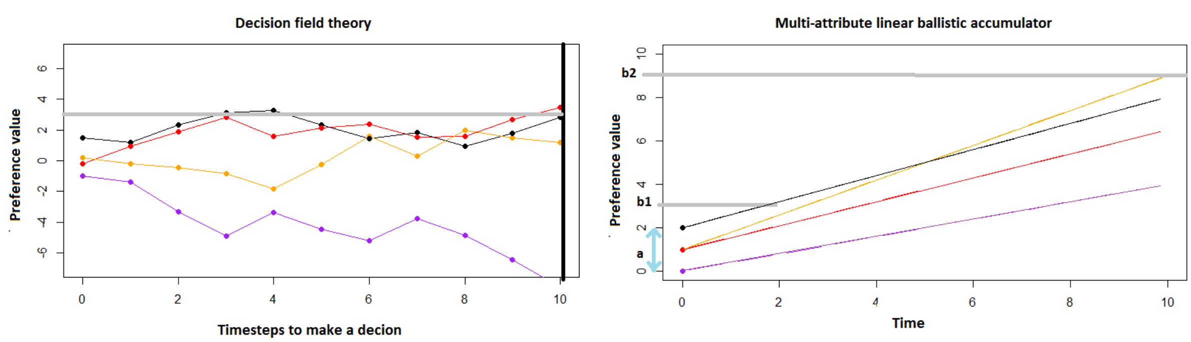

The linear ballistic accumulator model [Brown and Heathcote, 2008] and its multi-attribute version MLBA [Trueblood et al., 2013a, 2014] have a similar accumulation process for the preference of alternatives, but, in contrast with DFT, the updating is not stochastic. Instead, decision-makers start with some random amount of initial ‘evidence’ for each alternative, that then ‘drifts’ until one of the alternatives reaches a threshold. The preference values grow linearly at some drift rate dependent on the attributes of the alternative. Depending on the level of the threshold, different alternatives may be chosen. This is demonstrated in the right panel of Figure 1, in which the alternatives start with some random initial value, which we show as being in an interval a, and different alternatives are chosen if the threshold value is b_1 or b_2. The linear drift rates imply that, once the alternative with the largest drift rate value ‘gains the lead’, unlike in DFT, there is no way for another alternative to recover and be chosen. Whilst this would not be the case with a non-linear specification, the current model is specifically linear to allow for simple calculation of the probabilities of alternatives. Of course, different alternatives can be chosen depending on the length of the deliberation process. As with DFT, the value that evolves over time is a preference value, while the horizontal axis in Figure 1 now relates to actual time, given that no additional comparisons are made.

The mathematics underlying MLBA and DFT is vastly different. LBA was specifically designed such that it is 'simple' [Brown and Heathcote, 2008] and mathematically tractable, with MLBA subsequently developed such that it can also accurately capture and predict context effects. The simpler mathematical nature means that the probabilities of alternatives can easily be calculated from a combination of normal and uniform cumulative density functions (see Section 2.3 for a full description of MLBA).

It may be noted that there are numerous other accumulation models from mathematical psychology that seek to explain choice processes and predict choices. However, not all are currently suitable for transitioning into applied choice modelling. Given the complex nature of transport datasets, for example in terms of number of alternatives, even simple implementations will impose large computational costs. This is then further increased if analysts wish to add random heterogeneity in preferences, and models from mathematical psychology thus need to be efficient to run at a basic level if they are to compete. This means that models that do not have analytical solutions for calculating the probability of alternatives will likely not be suitable options. For example, the leaky competing accumulator model (LCA, Usher and McClelland 2001) would require two levels of simulation in order to incorporate random parameters. Requirements for computer intensive simulation are also issues for the associative accumulation model [Bhatia, 2013] or the attentional drift diffusion model [Krajbich et al., 2012].

2.2. Decision field theory (DFT)

In this section, we first look at the existing implementation of DFT before making a number of methodological improvements and finally turning to identification issues.

2.2.1. Existing implementation

2.2.1.1 *Theoretical model*

Decision field theory, as an accumulator model, has preference values for each alternative that update over time[2] (see left panel in Figure 1).

$$P_t = S \cdot P_{t-1} + V_t \tag{1}$$

The previous values, P_{t-1}, are multiplied by a feedback matrix, S, and a valence vector V_t is added. The feedback matrix has two parameters that control for the impact of attraction, similarity and compromise effects [Roe et al., 2001, Hotaling et al., 2010, Noguchi and Stewart, 2014], and is

[2]Note that whilst it is possible that separate choice tasks could be linked through parameters controlling for learning effects, all of the DFT models in this paper assume that all choice tasks are entirely independent of each other to make them comparable with the RUM and RRM models without state-dependence.

defined as:

$$S = I - \phi_2 \times \exp(-\phi_1 \times D^2) \quad (2)$$

where ϕ_1 is a sensitivity parameter, ϕ_2 is a memory parameter and D is the distance between the attributes of the alternatives. Whilst the relative importance of the different attributes can be taken into account with a psychological distance function [Hotaling et al., 2010] and work on new distance functions is possible [e.g. Berkowitsch et al., 2015], the Euclidean distance between the attributes can also be used for simplicity [Qin et al., 2013]. The sensitivity parameter, ϕ_1, affects how much the alternatives compete with each other. Values very close to zero results in the distance between the attributes of alternatives becoming less important, whereas higher values result in more competition between similar alternatives. The memory parameter (also known as the decay parameter) determines the relative importance of attributes considered towards the end of the decision process relative to those considered at the start. A value of one results in zeros on the diagonals of the feedback matrix, which results in the preference already accumulated becoming irrelevant. As this value tends towards zero, the importance of the already accumulated preference increases.

At each timestep, a DFT model assumes that the decision-maker compares a single attribute across all of the alternatives. This results in a random valence vector at time t, V_t, which can be calculated as:

$$V_t = C \cdot M \cdot W_t + \varepsilon_t \quad (3)$$

where C is a contrast matrix used to rescale the values such that they total zero, M is the matrix of attribute values and $W_t = [0..1..0]'$ with the k^{th} entry, i.e. $W_{t,k} = 1$ if and only if attribute x_k is the attribute being attended to by the decision-maker at timestep t. A DFT model thus estimates a weight, w_k, for the likelihood of attribute x_k being the single attribute attended to at a given timestep, where $\sum_k w_k = 1$. There is also a random error vector, $\varepsilon_t = [\varepsilon_1..\varepsilon_n]'$, with $\varepsilon_i \sim N(0,s)$, distributed identically and independently across alternatives, time and individuals. This allows for flexibility in the range of probability values that DFT predicts. This is in essence an error or noise parameter [Roe et al., 2001], for which higher values would be expected for more complex decision-making tasks [Hotaling et al., 2010].

2.2.1.2 *Calculating probabilities*

Under a DFT model, at the conclusion of the deliberation process, the alternative that is chosen is the one with the greatest preference value, regardless of whether the individual stopped deliberating due to a time threshold or due to the preference value for one of the alternatives reaching a preference threshold. Given that most choices do not have a strict time threshold, some applications of DFT calculate the probability for each alternative's preference value reaching a particular threshold first (see examples in Hotaling et al. [2010] and Turner et al. 2018). However, as this probability has no closed-form solution for more than two alternatives, we rely on Roe et al. [2001]'s method to calculate the probability for each alternative after t preference accumulation steps. This uses the expected value and the covariance of the preference values (ξ_t and Ω_t) and results in the stochastic variation being averaged out such that probabilities of the alternatives can be calculated without the requirement of computationally heavy simulation.

To calculate the expected value of the preference values, we must first expand Equation 1, which results in:

$$P_t = \sum_{r=0}^{t-1} S^r \cdot V_{t-r} + S^t \cdot P_0 \tag{4}$$

where $P_0 = [\delta_1, \delta_2, .., \delta_n]$, the initial preference vector. This is often assumed to be a vector of zeros [Busemeyer and Diederich, 2002] but can also be used to capture underlying preferences for different alternatives [Hancock et al., 2018].

The attribute weights w_k are stationary, therefore W_t can be considered a stationary stochastic process. This means that V_t is also a stationary stochastic process with mean $E[V_t]$ and a variance covariance matrix given by $Cov[V_t]$. We let ε_t vary according to a normal distribution with mean zero and variance σ_ε^2. Thus, given that $\mu = E[V_t]$, it can be calculated as $\mu = C \cdot M \cdot w$, where w is a vector containing the attribute attention weights, w_k, which corresponds to the probability of each of the attributes being considered. We also have $Cov[V_t] = \Phi = C \cdot M \cdot \Psi \cdot M' \cdot C' + \sigma_\varepsilon^2$, where $\Psi = Cov[W_t]$ (C and M are matrices of constants). We can then calculate the expected value and the covariance of P_t. With S being a constant, $E[P_t]$ reduces to:

$$E[P_t] = \xi_t = \sum_{r=0}^{t-1} S^r \cdot \mu + S^t \cdot P_0 \tag{5a}$$

$$= (I - S)^{-1}(I - S^t) \cdot \mu + S^t \cdot P_0 \tag{5b}$$

We can also now calculate the covariance of the preference values:

$$Cov[P_t] = \Omega_t = Cov\left[\sum_{r=0}^{t-1} S^r \cdot V_{t-r} + S^t \cdot P_0\right] \tag{6a}$$

$$= \sum_{r=0}^{t-1} \left[S^r \cdot \Phi \cdot S^{r'}\right] \tag{6b}$$

The resulting calculations are complex, but as shown in our earlier work [Hancock et al., 2018], we can further simplify $Cov[P_t]$ such that we can avoid the summation. We replace the feedback matrices with a matrix Z of size $J^2 \times J^2$ (where J is the number of alternatives) and reshape Φ (with entries $p_{i,j}$) into a column matrix, $\overline{\Phi}$:

$$Z = \begin{bmatrix} z_{1,1} & z_{1,2} & \cdots & z_{1,J^2} \\ z_{2,1} & z_{2,2} & \cdots & z_{2,J^2} \\ \vdots & \vdots & \ddots & \vdots \\ z_{J^2,1} & z_{J^2,2} & \cdots & z_{J^2,J^2} \end{bmatrix}, \overline{\Phi} = \begin{bmatrix} p_{1,1} \\ p_{2,1} \\ \vdots \\ p_{J,1} \\ p_{2,1} \\ \vdots \\ p_{J,J} \end{bmatrix} \tag{7}$$

As S is a symmetric matrix, we can then define Z by setting it as the Kronecker product of S with itself: $z_{i,j} = s_{(i \bmod J, j \bmod J)} \cdot s_{(\lceil i/J \rceil, \lceil j/J \rceil)}$. The calculation of the covariance of P_t now simplifies to:

$$Cov[P_t] = \Omega_t = \sum_{r=0}^{t-1} \left[S^r \cdot \Phi \cdot S^{r'} \right] \tag{8a}$$

$$= \sum_{r=0}^{t-1} \left[Z^r \cdot \overline{\Phi} \right] \tag{8b}$$

$$= (I - Z)^{-1} (I - Z^t) \overline{\Phi} \tag{8c}$$

These succinct forms for ξ_t and Ω_t mean that we can now calculate the probability with which each alternative is chosen. On the basis of the multivariate central limit theorem, P_t converges to the multivariate normal distribution [Roe et al., 2001]. This means that if a time threshold is reached, the probability of choosing alternative A from a set of n alternatives at time t is:

$$Prob\left[\max_{i \in n} P_t[i] = P_t[A]\right] = \int_{X>0} exp\left[-(X-\Gamma)'\Lambda^{-1}(X-\Gamma)/2\right] / (2\pi|\Lambda|^{0.5}) dX \tag{9}$$

with $X = [P_t[A] - P_t[B], ..., P_t[A] - P_t[n]]'$, $\Gamma = L\xi_t$, $\Lambda = L\Omega_t L'$ where

$$L = \begin{bmatrix} 1 & -1 & 0 & \dots & \dots & 0 \\ 1 & 0 & -1 & \ddots & & \vdots \\ 1 & \vdots & \ddots & \ddots & \ddots & \vdots \\ 1 & \vdots & & \ddots & -1 & 0 \\ 1 & 0 & \dots & \dots & 0 & -1 \end{bmatrix} \tag{10}$$

with L being a matrix constructed with a column vector of 1s and a negative identity matrix of size $n-1$ where n is the number of alternatives. The column vector of 1s is placed in the i^{th} column where i is the chosen option[3].

2.2.2. New developments

2.2.2.1 *Scaling of DFT*

In a typical linear additive RUM or RRM model, changing the units of a single attribute only affects the parameter for that attribute. For example, changing the unit of travel time from minutes to hours results in the corresponding marginal utility component being multiplied by 60, with no impact on other parameters.

DFT on the other hand is scale-variant [Busemeyer and Diederich, 2002, Trueblood et al., 2013a], where the requirements that the attribute weights sum to one ($\sum_k w_k = 1$) means that a change in the scale for one attribute will have an impact on the relative importance weights for all attributes. An illustration of this is given in Table 1. If we originally have weights of 0.6 and 0.4 for cost and time respectively, then the new importance for time will be multiplied by a value of 60 if we change from minutes to hours. These then need to be rescaled to ensure summation to 1, leading to new weights of 0.024 for cost and 0.976 for time. Whilst this adjustment can be easily made for only two attributes, estimation is simpler if this can be avoided.

[3] For further details and a more comprehensive description of decision field theory, please refer to Section 2 in Hancock et al. [2018].

TABLE 1 : Impact of changing the unit of time on attribute importance estimates

		Cost	Time
DFT coefficients	original weight	0.600	0.400
	new importance	0.600	24.000
	new weight	0.024	0.976

We define a new scaling method which translates attribute values into subjective values. This is achieved by multiplying the values by a set of attribute-specific scaling coefficients[4], β_{DFT}. This results in a different function for the random valence vector at time t, V_t, which can now be calculated as:

$$V_t = C \cdot \beta \cdot M \cdot W_t + \varepsilon_t \tag{11}$$

where M is the original attribute matrix, but with each attribute multiplied by its corresponding scaling value from the diagonal matrix β, which has the set of estimated scaling values, β_{DFT}, on the diagonal entries. With this specification, a decision-maker still attends to a given attribute at random in a given evaluation. In W_t, as before, one element is equal to 1 with all others 0. Furthermore, we also multiply the attributes by their corresponding scaling coefficients before they are used to estimate the distances between alternatives in the feedback matrix (D in Equation 2).

The addition of these scaling parameters results in three possible versions of DFT models:

1. A general model incorporating both attribute attribute-specific scaling coefficients and attribute weights.
2. A model with attribute-specific scaling coefficients only. The attribute weights, w_k, are not estimated, with the modeller simply fixing the weights equal to $1/K$ where K is the number of attributes.
3. A model with attribute weights only. The scaling coefficients are all fixed to a value of 1. This is of course equivalent to the original DFT specification by Roe et al. [2001].

The addition of attribute scaling coefficients to DFT results in a number of important benefits.

- First, the revised version of DFT is no longer scale-variant. Changing the unit of travel time from minutes to hours will now impact the estimate for the travel time scaling coefficient only. This means that for each marginal utility coefficient in a RUM model (or a marginal regret coefficient in a RRM model), there is a corresponding attribute scaling coefficient in the DFT model. This allows us to make comparisons across the different models in terms of relative importance of different attributes.

[4] We use the term β_{DFT} here as these values correspond to the marginal utility components, β, of RUM, but they cannot be used equivalently in, for example, value of travel time calculations.

- Second, the attributes are now adjusted accordingly for their relative importance before they enter the feedback matrix, meaning that we can calculate an appropriate psychological distance by simply taking Euclidean distances in the calculation of the feedback matrix. Consequently, we do not need a separate parameter w, as defined by Berkowitsch et al. [2015] to take the relative importance of the attributes into account.

- Third, we now separate out the effects of the frequency of considering individual attributes, through the weights, and the importance of different attributes in changing the preference values, through the scaling parameters. Conceptually, only the latter should be influenced by changes in units.

- Fourth, an even more important benefit of the proposed scaling approach relates to the possibility of attributes having opposite impacts on probabilities, i.e. some attributes being desirable and others being undesirable. In the traditional DFT model, an analyst needs to make a priori assumptions about this directionality, and failing to correct for the *sign* of the impact of attributes can have undesired consequences, as illustrated in Table 3 of Hancock et al. [2018]. With our new approach, we no longer require a priori knowledge or assumptions about whether an attribute has a positive or negative impact on the likelihood of an alternative being chosen, as the attribute scaling parameters can be estimated to be either positive or negative. This not only results in it being possible to take all attributes into account without any initial adjustments, but would also, in a random coefficients DFT model, allow for the possibility of different signs for a given parameter across different individuals.

- Finally, this new scaling method allows for the impact of attributes to be adjusted by incorporating socio-demographic interactions such as income effects or alternative specific coefficients for given attributes.

2.2.2.2 *Identification of parameters*

The literature on DFT (and other mathematical psychology models) often lacks crucial details in relation to model identification. As with any standard choice model, a DFT model requires a normalisation of location and scale. For DFT, this is a result of the multiplication of the expectation, ξ_t, and covariance, Ω_t, by L (see Equation 9). This results in only differences between alternatives, in terms of both expected values and covariances, mattering. In order to estimate the probability of alternatives under a DFT model, we require estimates for four 'process parameters'[5], which are exclusive to DFT and inform the process by which alternatives accumulate preference. These are ϕ_1 and ϕ_2, the sensitivity and memory parameters respectively, the number of deliberation timesteps[6], t, and the variance of the error term, σ_ε^2.

[5]Henceforth, if we refer to 'process parameters', we mean parameters which have no equivalent measure in a traditional model such as RUM or RRM.

[6]The choices that we investigate in this paper do not have a strictly imposed time threshold, and we do not impose one as a result of estimating this parameter, as we make no assumptions on the relationship between the number of deliberation steps and the real world time taken to make the choice.

Theoretical identification

To study identification, we write the expectation of the preference values after t deliberation timesteps as $\xi_t = a + b$ and the covariance of the preference values as $\Omega_t = c\overline{(d+e)}$. Thus:

- $a = (I-S)^{-1}(I-S^t)\cdot\mu$
- $b = S^t \cdot P_0$
- $c = (I-Z)^{-1}\,(I-Z^t)$
- $d = C\cdot\beta\cdot M\cdot\Psi\cdot M'\cdot\ldots.\cdot C'$
- $e = \sigma_\varepsilon^2$.

An overspecification occurs if we can have two sets of parameters (θ_1 and θ_2) such that exactly the same set of probabilities are generated for the full set of choice scenarios. This occurs if all means change by the same amount, or if the means and square root of the covariance matrix are scaled by the same factor. We thus consider the following three scenarios:

1. $a_1 + b_1 = \delta + a_2 + b_2$
2. $a_1 + b_1 = (a_2 + b_2)\cdot\gamma$
3. $c_1\overline{(d_1+e_1)} = (c_2\overline{(d_2+e_2)})\cdot\gamma^2$

An overspecification of location could exist if some value δ exists such that scenario 1 holds. This is a result of the specification of Equation 9, where ξ_t is multiplied by L, giving a vector of differences between the expected preference of the chosen alternative and each other alternative. Consequently, the addition of $delta$ to each preference value will results in no change in the value of $L\xi_t$.

To avoid this overspecification, the simplest solution is to fix one of the alternative specific constants in the initial preference matrix P_0. Under DFT, adding the same value to each alternative specific constant will result in the same increase in expected preference value for each alternative when $\phi_2 = 0$ (as this results in the feedback matrix becoming an identity matrix, which means there is no impact on P_0, see Equation 2.2.1.2).

An overspecification of scale can exist if some value γ exists such that scenarios 2 and 3 both hold. This is a result of the probabilities under DFT being calculated with the use of multivariate normal distributions (see Equation 9). Critically, these scenarios result in $\xi_{t,1} = \gamma\cdot\xi_{t,2}$ and $\Omega_{t,1} = \gamma^2\cdot\Omega_{t,2}$. Then if we have a set of parameters, θ_1 that results in probabilities, Pr_1:

$$Pr_1 = \int_{X_1>0} exp\left[-(X_1-\Gamma_1)'\Lambda_1^{-1}(X_1-\Gamma_1)/2\right]/(2\pi|\Lambda_1|^{0.5})dX, \tag{12}$$

with the corresponding parameter values θ_2 resulting in probabilities, Pr_2:

$$Pr_2 = \int_{X_2>0} exp\left[-(X_2-\Gamma_2)'\Lambda_2^{-1}(X_2-\Gamma_1)/2\right]/(2\pi|\Lambda_2|^{0.5})dX, \tag{13}$$

we can now simply observe that the substitutions implied by scenarios 2 and 3, $X_1 = \gamma \cdot X_2$, $\Gamma_1 = \gamma \cdot \Gamma_2$ and $\Lambda_1 = \gamma \cdot \Lambda_2$, result in $Pr_1 = Pr_2$, as the γ terms all drop out of the equation. Thus we obtain two distinct sets of parameters θ_1 and θ_2 that result in the same estimated probabilities of observing the full set of choices.

To avoid this overspecification, we need to understand the situations under which scenarios 2 and 3 hold. One example of when this is the case is when:

$$\theta_1 = [w, \beta_{DFT}, P_0, \phi_1, \phi_2, t, \sigma_\varepsilon^2] \tag{14}$$

and:

$$\theta_2 = [w, \gamma \cdot \beta_{DFT}, \gamma \cdot P_0, \phi_1/(\gamma^2), \phi_2, t, \gamma \cdot \sigma_\varepsilon^2], \tag{15}$$

with $\beta_{DFT} = [\beta_1, \beta_2, .., \beta_k]$, $w = [w_1, w_2, .., w_k]$ and $P_0 = [\delta_1, \delta_2, .., \delta_n]$. To avoid this overspecification, we need to fix a parameter such that $\gamma = 1$ and thus $\theta_1 = \theta_2$. In the absence of a priori knowledge on the directionality of the attributes or the directionality of the underlying preferences towards an alternative, the scaling parameters $\beta_1, \beta_2, .., \beta_k$ and the alternative specific constants should not be fixed. Furthermore, ϕ_2 could have an estimate of zero, resulting in ϕ_1 having no impact in equations 14 and 15. Thus, the safest option is to fix $\sigma_\varepsilon^2 = 1$.

It is easy to see a relationship between these two normalisations of location and scale and their corresponding normalisations in a RUM context.

It is also worth noting at this point that whilst in choice-only datasets, the attribute attention weights and attribute scaling parameters may appear to capture the same feature (the relative importance of an attribute), these parameters are not theoretically confounded, as they have differing impacts within the calculation of probabilities under DFT. The attribute attention weights do not impact the feedback matrix (whilst the attribute scaling parameters do, as detailed in Section 2.2.2.1). Additionally, the attribute scaling parameters do not impact Ψ (the covariance of the attention weights), which is used directly in the calculation of the covariance of the preference values (see Equation 2.2.1.2). This is instead calculated solely with the estimates for the attribute attention weights. As a result, only a model containing both will have the full flexibility to capture both effects. I wonder if you need a couple of equations in this para?

Empirical identification and restrictions

The process parameters in DFT have important behavioural roles. However, DFT models are routinely estimated on data where the only observed outcome is the choice itself, with little information about the process by which that choice was reached. If such process information was available, analysts could use it as additional indicators (i.e. additional dependent variables) in a joint estimation of process and outcome, and this would help inform the values of these parameters. In the absence of such data however, some of the parameters may become partially confounded.

We additionally impose various restrictions to aid empirical identification of parameters under a DFT model. These, together with the restrictions identified in the previous section, are detailed in Table 2.

Some of these constraints are necessary to avoid identification issues, while others simply avoid sign issues. The following list details why each restriction is required.

- (1, 2 and 3) There are two theoretical identifications (restriction 1 and either restriction 2 or 3 in Table 2) detailed in the previous section, applied to avoid a theoretical overspecification.
- (4 and 5) There is an empirical identification issue as a result of the decay parameter, ϕ_2. In case there is no impact for the decay parameter, we obtain a value of ϕ_2 equal to or close to zero. This results in the sensitivity parameter, ϕ_1, having no impact on the probabilities of alternatives, thus resulting in an overspecified model. As a consequence, we try two specifications of DFT for each dataset in our empirical applications: one with and one without estimated values for these two feedback parameters.
- (6 and 7) A DFT model assumes that a decision-maker attends one attribute at each preference updating step, which implies that the attribute attention weights must all be positive and sum to one. Estimation of these attention weights is thus aided through the use of logistic transformations and identification is then ensured by fixing one of the adjusted attribute attention weights $w_1^* = 0$.
- (8) An exponential transformation ensures that the sensitivity parameter is positive. This restriction results in alternatives that are more similar to each other competing more with each other than alternatives that are less similar, as assumed implicitly by DFT models.
- (9) The number of preference updating timesteps must exceed a value of one. This can be also be solved through the use of an exponential transformation.
- (10) The noise that is added on at each timestep to the valence (see Equation 3) is drawn from a normal distribution with mean 0 and variance σ_ε^2. Consequently, as $\sigma_\varepsilon^2 \geq 0$, we instead estimate the standard deviation, σ_ε.

2.3. The multi-attribute linear ballistic accumulator model (MLBA)

In this section, we first look at the existing implementation of MLBA before making a number of methodological improvements and finally turning to identification issues.

2.3.1. Existing implementation

2.3.1.1 *Theoretical model*

Under MLBA, each alternative has a value that grows linearly towards a threshold (see right panel in Figure 1). The chosen alternative in an MLBA model is the first alternative to pass a threshold value, χ. There are two components in this process; the start points and the drift rates.

TABLE 2 : Restrictions on DFT parameters

No.	Parameter	Description	Restrictions	Reason	Estimated parameter	Relation
1	asc_n	asc for alternative n	=0	Theoretical identification	n/a	n/a
2*	σ_ε^2	error	=1	Theoretical identification	n/a	n/a
3*	β_k	scale for attribute k	fixed	Theoretical identification	n/a	n/a
4**	ϕ_2	decay	=0	Empirical identification	n/a	n/a
5**	ϕ_1	sensitivity	fixed	Empirical identification	n/a	n/a
6	w_i	weights for attributes	$\sum_1^k w_k = 1$	DFT assumption	w_i^*	$w_i^* = exp(w_i)/\sum_1^k exp(w_k)$
7	w_1	first attribute weight	$w_1^* = 0$	DFT assumption	n/a	n/a
8**	ϕ_1	sensitivity	> 0	DFT assumption	ϕ_1^*	$\phi_1 = exp(\phi_1^*)$
9	t	timesteps	> 1	DFT assumption	t^*	$t = 1 + exp(t^*)$
10*	σ_ε^2	error	≥ 0	Mathematical	σ_ε	$\sigma_\varepsilon^2 = (\sigma_\varepsilon)^2$

*Note that only one of restrictions 2 and 3 should be used, and that restriction 10 should be applied if 2 is not used.
** If the estimate for ϕ_2 is equal to or close to zero, restrictions 4 and 5 may also be required to ensure standard errors can be estimated for the other parameters (with restriction 8 no longer required).

Start points for each of the alternatives are drawn separately from a uniform distribution $U[0,A]$ where A is estimated. For example, Figure 1 demonstrates what a decision might look like if the start points are drawn from a distribution $U[0,2]$. A different value A_j could be estimated for each alternative j, although it is common practice [Trueblood et al., 2014] to assume that all alternatives have starting values that are drawn using the same estimate A[7].

Trueblood et al. [2013a, 2014] demonstrate that there are different methods for specifying drift rates for an MLBA model such that they explain context effects. In this application, however, we choose to fit versions similar to the mainstream version of MLBA [Trueblood et al., 2014] as this outperforms the first version (described by Trueblood et al. [2013a]) for our choice datasets (see appendix A). Under MLBA, we define the drift rates for the different alternatives as independent draws from normal distributions (truncated below zero), where, for alternative j, we have the drift rate D_j given as:

$$D_j \sim TN\left(d_j, s_j\right) \tag{16}$$

with mean drift rate d_j and standard deviation s_j. Typically, the standard deviation is set to be the same value for all alternatives, i.e. $s_j = s, \forall j$, but a different value could be estimated for each drift rate [Trueblood et al., 2013b].

In the current version of MLBA, mean drift rates follow the specification used by Trueblood et al. [2014]:

$$d_j = v_j + I_0 \tag{17}$$

where I_0 is a positive constant (which can be specified such that all drift rates have a positive mean) and v_j is a value function, similar to random regret minimisation in that it compares an alternative j against all other alternatives i across each attribute x. Specifically, with K attributes, we have that:

$$v_j = \sum_{j \neq i} \sum_{k=1}^{K} \left(w_{x_{k,i,j}} \cdot (x_{k,j} - x_{k,i})\right). \tag{18}$$

In this notation, $x_{k,i}$ is the objective value for the k^{th} attribute for alternative i, and $w_{x_{k,i,j}}$ is a weight for attribute x_k and alternative pairing i and j, which relates to the similarity between them[8]. In particular, the similarity is assumed to be an exponential decaying function of distance, with:

$$w_{x_{k,i,j}} = exp(-\lambda \cdot |x_{k,i} - x_{k,j}|) \tag{19}$$

Two different values of λ are used depending on whether the difference between $x_{k,i}$ and $x_{k,j}$ is positive or negative:

$$\lambda = \begin{cases} \lambda_1, & \text{if } x_{k,j} \geq x_{k,i}. \\ \lambda_2, & \text{if } x_{k,j} < x_{k,i}. \end{cases} \tag{20}$$

This feature can capture differences between the subjective similarity between A and B and the subjective similarity between B and A, which may not be equal [Tversky, 1977], with gains and

[7]Thus we too assume that the same value A is used for each alternative, in both the theoretical identification (2.3.2.3) and the applications of MLBA in this paper (3).

[8]Note that Equation 18 is equivalent to Equation 3 in Trueblood et al. [2014], but for multiple alternatives and multiple attributes.

losses regularly having been shown to be treated differently in a transport context [Hess et al., 2008, Masiero and Hensher, 2010, Stathopoulos and Hess, 2012].

In the original MLBA work [Trueblood et al., 2014], the model additionally translates attribute values into 'subjective values'. In their example, Trueblood et al. [2014] had two similar attributes: testimony strength of eyewitness P and testimony strength of eyewitness Q. Thus to translate objective values into subjective values, a parameter was introduced such that an 'indifference curve' could be calculated to avoid issues of extremeness aversion [Chernev, 2004], where, for example, values of 50-50 might be preferred to 70-30. However, under typical travel choice tasks, attributes are not as closely related. Thus to translate objective attribute values into subjective values in this case, we instead require some measure to translate the values appropriately such that the relative importance of the attributes is accounted for. Consequently, we use attribute importance parameters instead of m, the additional parameter in the original specification of MLBA [Trueblood et al., 2014]. We discuss this further in Section 2.3.2.2.

Calculating probabilities

If we have values for the drift rates of the alternatives and for the start point and threshold (A and χ respectively), we can calculate the probability of each alternative's accumulator being the first to finish, i.e. for its value function to exceed the threshold χ before any others do [Brown and Heathcote, 2008].

Given that the starting evidence is drawn from a uniform distribution $U[0,A]$, the amount of evidence that needs to be accumulated for an alternative to reach the threshold χ is U[χ-A,χ] (assuming $\chi > A$). Given an alternative's drift rate distribution, D_j, the cumulative distribution function for the time taken for the accumulator associated with alternative j is given by:

$$F_j(t) = Prob\left(\frac{U[\chi - A, \chi]}{D_j} < t\right) \tag{21}$$

Brown and Heathcote [2008] demonstrate that for a mean drift rate following a normal distribution[9], this reduces to:

$$\begin{aligned} F_j(t) = 1 + \frac{\chi - A - t\cdot d_j}{A}\cdot\Phi\left(\frac{\chi - A - t\cdot d_j}{t\cdot s}\right) - \frac{\chi - t\cdot d_j}{A}\cdot\Phi\left(\frac{\chi - t\cdot d_j}{t\cdot s}\right) \\ + \frac{t\cdot s}{A}\cdot\phi\left(\frac{\chi - A - t\cdot d_j}{t\cdot s}\right) - \frac{t\cdot s}{A}\cdot\phi\left(\frac{\chi - t\cdot d_j}{t\cdot s}\right) \end{aligned} \tag{22}$$

where ϕ and Φ are the standardised normal distribution's density and cumulative density functions, respectively. The associated probability density function is then:

$$\begin{aligned} f_j(t) = \frac{1}{A}\Bigg[-d_j\cdot\Phi\left(\frac{\chi - A - t\cdot d_j}{t\cdot s}\right) + d_j\cdot\Phi\left(\frac{\chi - t\cdot d_j}{t\cdot s}\right) \\ + s\cdot\phi\left(\frac{\chi - A - t\cdot d_j}{t\cdot s}\right) - s\cdot\phi\left(\frac{\chi - t\cdot d_j}{t\cdot s}\right)\Bigg] \end{aligned} \tag{23}$$

[9] We follow the first adjustment made by Heathcote and Love [2012] to translate this for truncated normals.

To then calculate the probability of a given alternative j being chosen[10], we need to calculate the probability density function of alternative j reaching the threshold χ before all other alternatives $i \neq j$:

$$Prob(j) = \int_0^\infty PDF_j(t)dt = \int_0^\infty f_j(t) \prod_{i \neq j} (1 - F_i(t))dt \tag{24}$$

2.3.2. New developments

2.3.2.1 *Incorporating baseline preferences in MLBA*

A key feature of discrete choice models belonging to the RUM family is the concept of alternative specific constants that capture baseline preferences for specific alternatives. Hancock et al. [2018] discusses in detail how this can be implemented in a DFT model. Here, we extend this to a MLBA model too.

In particular, we rewrite Equation 17 as

$$d_j = max(0, \delta_j + v_j + I_0), \tag{25}$$

where δ_j is an additional alternative specific estimated constant capturing a baseline preference for alternative j. The final adjustment we make is to ensure that each mean drift rate has a minimum value of zero (as oppose to fixing I_0 such that this is the case)[11].

2.3.2.2 *Incorporating attribute specific weights in MLBA*

An additional limitation of the original implementation of MLBA is in the treatment of the different attributes. Firstly, this applies in terms of directionality, noting that λ_1 is used for a positive difference between $x_{k,i}$ and $x_{k,j}$ independently of whether gains in attribute x_k increase or decrease the attractiveness of an alternative. This limitation is analogous to the issue with using weight parameters in DFT and would require an analyst to a priori change the sign on undesirable attributes. Secondly, the actual impact of differences between alternatives in a given attribute x_k is constant across attributes. Whilst one possibility is to use different valuation and weighting functions [Cohen et al., 2017], Trueblood et al. [2014] suggest that attribute biases can be dealt with by including attribute-specific 'bias parameters', β_k (an approach analogous to the attribute-specific scaling coefficients that we defined for DFT) in Equation 19, which becomes:

$$w_{x_{k,i,j}} = exp(-\lambda \cdot \beta_k \cdot |x_{k,i} - x_{k,j}|) \tag{26}$$

However, we can relax the limitations of attribute bias and directionality simultaneously by also making an adjustment to the value function (Equation 18), which now takes the same form

[10] For full derivations of equations 22, 23 and 24, refer to appendix A of Brown and Heathcote (2008).

[11] The use of truncated normals results in it being possible that adding the same constant to each drift rate can result in some more deterministic choices (when the mean drift rates are negative, for example) as well as some less deterministic choices. Adjusting the mean drift rates such that are at least zero avoids this and aids the estimation of MLBA. For more details of this, please see appendix B

as that of the original specification with the exception that we add in attribute-specific scaling coefficients, β_k. This results in the value function from Equation 18 being redefined as:

$$v_j = \sum_{j \neq i} \sum_{k=1}^{K} (w_{x_{k,i,j}} \cdot \beta_k \cdot (x_{k,j} - x_{k,i})). \tag{27}$$

As with the scaling applied to DFT, this change allows us to also make inferences about the relative importance of different attributes in MLBA, as well as incorporating interactions with socio-demographics at the level of individual attributes.

2.3.2.3 *Identification of parameters*

For the estimation of the probability of choosing alternatives in a MLBA model, we require estimates for K attribute scaling parameters, n alternative specific drift rate constants and estimates for six process parameters (A and χ, the start and threshold parameters respectively, a drift rate constant, I_0, a parameter for the standard deviation of the drift rates, s, and similarity parameters λ_1 and λ_2)[12]. MLBA can explicitly estimate the probability of an alternative being chosen as well as giving an output relating to the time the participant took to make the choice. As we use choice-only data, this results in the requirement for a number of restrictions, both theoretical and empirical, to aid and improve estimation.

Theoretical identification

For choice-only data (i.e. where no additional process information is available), the start and threshold parameters A and χ are perfectly confounded. This is a consequence of the fact that multiplying A and χ by scale factor f results in no change in the probabilities with which each alternative is chosen. Instead, this simply changes the time that alternative j finishes in from t_j to $f \cdot t_j$. As all alternatives are impacted in the same way, the probabilities of each alternative being chosen are not impacted. This can be demonstrated by multiplying A, χ and t by the factor f in the cumulative distribution function for the time taken by an accumulator (Equation 2.3.1.1). The factor f drops out, resulting in identical cumulative distribution functions being given by the set of parameters $\theta_1 = [A, \chi, t]$ and $\theta_2 = [A \cdot f, \chi \cdot f, t \cdot f]$. Thus, when integrated over $t = 0 \rightarrow \infty$ (see Equation 24), the resulting probabilities are equal. This means that we need to fix either A or χ. In the applications in this paper, we choose (in line with previous applications, see Trueblood et al. 2014) to fix $A = 1$.

Additionally, it is possible that the choice probabilities remain exactly the same if all mean drift rates d_j and the standard deviation s are multiplied by the same factor, g. This is possible, for example, with the parameter sets θ_1 and θ_2:

$$\theta_1 = [\beta_{MLBA}, \delta, A, \chi, s, I_0, \lambda_1, \lambda_2] \tag{28}$$

[12]Different values for these process parameters could be used for different alternatives. However, in the applications in this paper and for the purposes of identification in the sections below, we assume that the same value is used for each alternative with the exception of the drift rate constants.

and

$$\theta_2 = [\beta_{MLBA} \cdot g, \delta \cdot g, A, \chi, s, I_0, \lambda_1 \cdot (1/g), \lambda_2 \cdot (1/g)], \quad (29)$$

with $\beta_{MLBA} = [\beta_1, \beta_2, .., \beta_k]$ and $\delta = [\delta_1, \delta_2, .., \delta_n]$. As with the previous example for changing start and threshold parameters, this effect simply changes the time that alternative j finishes in. Again, all alternatives are impacted in the same way, and thus the probabilities remain the same (as we are integrating over $t = 0 \rightarrow \infty$), with the only change being that each alternative j now takes $\frac{t_j}{g}$ rather than t_j time to finish. Thus we must fix either s or one of the drift rates d_j to ensure identification. As before, we follow Trueblood et al. [2014] in choosing to fix $s = 1$.

Furthermore, as a contrast to RUM models, where the same differences between alternative specific constants δ_j (see Equation 25) results in the same probabilities, each mean drift rate can have a separately identified constant, as the greater the rates, the less deterministic the choice is. This is a result of each alternative taking less time to finish, meaning there is less time for the alternatives to 'overtake' each other. Consequently, the draw for the starting point of an alternative has a greater influence on the probability of that alternative finishing first. Note that if we additionally estimate I_0 (which is desirable as this allows us to determine whether the baseline preferences for the alternatives differ significantly from each other), then one of the constants d_j must be fixed to ensure identification.

Empirical identification and restrictions

We additionally impose various restrictions to aid empirical identification of parameters under a MLBA model. These, together with the restrictions identified in the previous section, are detailed in Table 3. The following list details why each restriction is required.

- (1) As detailed in the section on including baseline preference parameters for the different alternatives (see Section 2.3.2.1), one of these alternative specific constants must be fixed to avoid a theoretical overspecification.
- (2 and 3) Further theoretical restrictions as discussed in Section 2.3.2.3.
- (4 and 5) Exponential transformations ensure that the lambda parameters are positive. This is required if the similarity between alternatives is to be a decaying function of distance [Trueblood et al., 2014].
- (6) A further restriction is required on the threshold, χ, which must be specified such that it is at least the same value as the start parameter (to avoid the possibility that more than one alternative reaches the threshold before any deliberation has taken place).

2.3.3. Variations between MLBA and econometric choice models

Crucially, the above probability for an alternative being chosen (Equation 24) holds specifically for MLBA models with truncated normal drift rate distributions. Terry et al. [2015] demonstrate

TABLE 3 : Restrictions on the parameters within MLBA

No	Parameter	Description	Restrictions	Reason	Estimated parameter	Relation
1	asc_n	asc for alternative n	=0	Theoretical identification	n/a	n/a
2	A	start	=1	Theoretical identification	n/a	n/a
3	s	drift rate standard deviation	=1	Theoretical identification	n/a	n/a
4	λ_1	similarity	≥ 0	MLBA assumption	λ_1^*	$\lambda_1 = exp(\lambda_1^*)$
5	λ_2	similarity	≥ 0	MLBA assumption	λ_2^*	$\lambda_2 = exp(\lambda_2^*)$
6	χ	threshold	$\geq$ A	MLBA assumption	χ^*	$\chi = A * (1 + exp(\chi^*))$

that the assumption of Fréchet-distributed drift rates with start rate parameter $A = 0$ results in probabilities:

$$Prob(j) = \frac{exp(d_j)}{\sum_i^J exp(d_i)}, \tag{30}$$

where J is the number of alternatives and d_j is still the mean drift rate for alternative j. Consequently, should MLBA's value function (Equation 27) be replaced such that it is linear (possible if, for example, we set $w_{x_{k,i,j}} = 1$), we obtain new values:

$$v_j = \sum_{j \neq i} \sum_{k=1}^{K} (\beta_k \cdot (x_{k,j} - x_{k,i})), \tag{31}$$

which, when also combined with the assumption of Fréchet-distributed drift rates with start rate parameter $A = 0$ results in a standard multinomial logit (MNL) model (written in terms of utility differences)[13]. Similarly, if the value function becomes:

$$v_j = -\sum_{j \neq i} \sum_{k=1}^{K} ln(1 + exp(\beta_k \cdot (x_{k,i} - x_{k,j}))), \tag{32}$$

the MLBA model becomes equivalent to a random regret minimisation (RRM) model [Chorus, 2010].

The MLBA model that we test in the applications in this paper, however, has three key differences. Firstly, we assume truncated normal drift rate distributions. Secondly, we assume that $A \neq 0$. Thirdly, we use weights in the value functions (Equation 27) for the translation from objective to subjective values.

In Section 3.6, we test whether it is the 'error structure' (i.e. whether we use truncated normal drift rate distributions or Fréchet-distributed drift rates with start rate parameter $A = 0$) or the value function that drives differences between MNL, RRM and MLBA. This test is possible as both MNL and RRM use Equation 30 to generate probabilities of different alternatives. Consequently, we can define six different combinations of error structure and value function:

TABLE 4 : Variations of model structures

		'Error Structure'	
		Eq.30	Eq.24
'Value Function'	Eq.31	MNL	Specification 1
	Eq.27	Specification 2	MLBA
	Eq.32	RRM	Specification 3

This gives us MNL, RRM and MLBA and three additional specifications:

[13]Note that the models are only perfectly equivalent if the choice set size is identical across all choice tasks in the dataset.

- Specification 1: which would correspond to a MLBA model with linear attributes differences except that we exponentiate each of the calculated drift rates to ensure that the mean drifts are positive.
- Specification 2: a 'context-dependent' logit model.
- Specification 3: a MLBA model which still treats positive and negative attribute differences differently, but at the cost of no additional parameters.

3. EMPIRICAL APPLICATIONS ON REVEALED AND STATED CHOICE DATA

In this section, we present empirical results using DFT and MLBA on three different datasets, two from stated choice (SC) surveys and one from a revealed preference (RP) survey, where the latter is the first DFT/MLBA application to RP data. We provide a detailed investigation as to the empirical identification of DFT and MLBA. This is crucial, as in the context of choice-only data, some of the process parameters will be confounded, and it has not yet been established what normalisation should be applied. We also compare the estimation results to typical MNL and RRM models. We finally present an empirical comparison between the different existing specifications of DFT and our proposed new scaling approach.

The estimation of DFT and MLBA remains a non-trivial computational task even with our methodological developments, and efficient implementation as well as good starting values are essential. In all of our applications, we use the R packages maxLik [Henningsen and Toomet, 2011] and *Apollo* [Hess and Palma, 2019] for estimation of the likelihood function and the RCPP package together with the Armadillo C++ linear algebra library for fast calculation of the matrices required for finding the probability under which each alternative is chosen under a DFT model [Eddelbuettel et al., 2011, Sanderson and Curtin, 2016]. Additionally, we use an initial parameter search algorithm based on the heuristic for non-linear global optimisation developed by Bierlaire et al. [2010] in an attempt to reduce the risk of convergence to poor local optima as well as an excessively long estimation process.

3.1. First stated choice survey

Our first dataset is a subset of the data from the Danish value of time study [Fosgerau, 2006]. This dataset comes from a typical stated choice survey, where 545 participants faced a total of 4,214 choices between them. The choices were for car drivers and specifically the choice between two different routes, described only by travel cost and travel time, where one route is cheaper, but the other is faster. The aim of such a setup is to understand trade-offs between time and money, leading to estimates of the value of travel time (VTT). While very simplistic in nature, this type of dataset is a useful first step in moving from the abstract settings in mathematical psychology towards more complex choices in a transport setting. In all models, we only focussed on the time and cost attributes after earlier results confirmed there was no left-right bias that would require the inclusion of alternative specific constants.

Table 5 shows the results for the first SC dataset. Where appropriate, we used the constraints from Table 3 but then report the actual transformed estimates in Table 5, along with the transformed standard errors, obtained using the Delta method [cf. Daly et al., 2012].

We first have two MNL models, one using a purely linear specification while the second additionally estimates parameters for the logarithm of time and cost. This latter model offers a significant improvement in fit over the first model, and all four coefficients remain negative, where the significant estimates for the log-time (β_{LTT}) and log-fare (β_{LF}) parameters indicate non-linear sensitivities.

Whilst a number of different parameters within DFT could be fixed to solve the theoretical overspecification issue identified in Section 2.2.2.2, we follow restriction 2 from Table 2 such that a priori information about the other attributes is not required. We then initially trial two different DFT models to test the impact of removing the effect of the feedback matrix (model 2 compared to model 1).

As there are only two alternatives in the Danish dataset, ϕ_1 is of little meaning, given that it is a parameter for the level of competition between alternatives dependent on the distance between them. Given that there is only one distance calculated per choice, a significant estimate for ϕ_1 here means that a pair of more similar alternatives simply reduces the overall preference for both alternatives, resulting in a less deterministic choice. Additionally, ϕ_2, the memory parameter, has little meaning when the sequence of attribute attendance is not known. It however also contributes to the level of competition between alternatives as a value of $\phi_2 = 0$ results in the value of ϕ_1 having no impact (see Equation 2). For this dataset, the use of an identity matrix in place of the feedback matrix (model 2) results in an insubstantial loss of model fit. Additionally, we find that if we fix $\sigma_\varepsilon = 0$ (meaning that all of the variation in the DFT model comes from the random attribute attendance), there is again no substantial loss of model fit. Note that with $\sigma_\varepsilon = 0$, we require a further normalisation. Given that we have negative coefficients for the scaling parameters, we follow restriction 3 from Table 5 by fixing a beta to -1. As a consequence of the the noise parameter having an insignificant impact on model fit, fixing it to a value of 1 (as configured in models 1 and 2) results in insignificant parameter estimates for the β-coefficients. By appropriately fixing $\sigma_\varepsilon = 0$ and also fixing the first β, we recover significant parameter estimates elsewhere.

Similar to DFT, MLBA has many parameters that have little interpretable output when an analyst only has access to the choice data and no additional psychometric or process data. For example, a decision-maker could make a choice quickly because there is a small difference between the start and threshold parameters or because they have a higher deviation in the drift rates. Consequently, if we only have choice data and no information about the process in which the choice was made, then some of the MLBA parameters become confounded (see Section 2.3.2.3). As with DFT, we initially test MLBA using a full specification, which implies only fixing the start parameter A and the drift rate standard deviation s to values of 1.

In model 2, we fix the threshold parameter χ, which is a common approach in mathematical psychology [Trueblood et al., 2014, Cohen et al., 2017, Cataldo and Cohen, 2018] (fixing it to a value of 2 as done in the original MLBA paper Trueblood et al. 2014), but find that this is not appropriate in this case, leading to a substantial loss of fit. On the other hand, our initial estimate

TABLE 5 : Estimation results and identifications tests on the first SC dataset

Model		MNL	MNL	DFT			MLBA		
	Version	1	2	1	2	**3**	1	2	**3**
	Free Pars.	3	5	6	4	**3**	7	6	**6**
	Log-likelihood	-2,301.25	-2,211.69	-2,015.54	-2,016.24	**-2,016.28**	-2,005.68	-2,036.36	**-2,005.92**
	BIC	4,627.54	4,465.11	4,081.15	4,065.87	**4,057.60**	4,069.79	4,122.79	**4,061.93**
β_{TT}	est.	-0.1938	-0.1590	-0.3572	-1.3077	-1.0000	-2.2715	-3.0842	-3.2257
	r. t-rat.	-13.53	-8.85	-1.11	-1.05	**fixed**	-8.10	-34.91	-3.17
β_F	est.	-2.4079	-1.7637	-5.6778	-20.9080	-16.0515	-35.9239	-45.8770	-50.9630
	r. t-rat.	-13.51	-9.20	-1.08	-1.03	-25.89	-8.22	-29.39	-3.25
β_{LTT}	est.		-1.03						
	r. t-rat.		-2.70						
β_{LF}	est.		-1.90						
	r. t-rat.		-5.39						
δ_1	est.	0.0264	0.0326	0.2337	1.4729	1.1068	0.3562	0.4397	0.5163
	r. t-rat.	1.08	1.29	1.43	0.96	2.51	2.39	1.86	3.05
ϕ_1	est.			0.5195	0.0000	0.0000			
	r. t-rat.			1.64	**fixed**	**fixed**			
ϕ_2	est.			-0.3388	0.0000	0.0000			
	r. t-rat.			-9.77	**fixed**	**fixed**			
σ_ε	est.			1.0000	1.0000	0.0000			
	r. t-rat.			**fixed**	**fixed**	**fixed**			
t	est.			38.2704	6.6058	6.4712			
	r. t-rat. (*vs* 1)			2.14	13.61	15.62			
χ	est.						1.1929	2.0000	1.1612
	r. t-rat. (*vs* 1)						1.88	**fixed**	7.05
I_0	est.						1.4611	39.4990	1.4726
	r. t-rat.						1.68	12.60	3.59
λ_1	est.						0.0009	0.0025	0.0000
	r. t-rat.						1.56	33.74	**fixed**
λ_2	est.						0.2539	0.0369	0.1983
	r. t-rat.						1.90	8.86	5.27

for λ_1 is so close to zero that fixing it does not lead to any significant loss of fit (model 3). A value of $\lambda_1 = 0$ results in weights, $w_{x_{k,i,j}} = 1$, thus resulting in a simplified calculation of the mean drift rates with linear contributions from positive differences of attributes $x_{k,j} - x_{k,i}$. This means that an MLBA model which does not find a similarity effect will result in λ parameters approaching zero. This issue resulted in previous applications of MLBA resorting to different valuation and weighting functions [Cohen et al., 2017]. In this case, however, only one λ approaches zero.

In terms of model performance using BIC, we see that DFT and MLBA both outperform MNL. The difference in fit between DFT and MLBA is much smaller than between these two models and MNL, with a slight advantage for DFT in terms of BIC, but for MLBA using log-likelihood.

3.2. Second stated choice survey

The second stated choice dataset we consider has a total of 368 participants, each completing 10 choice tasks resulting in 3,680 choices. The participants are all public transport commuters living in the UK. Each task involves the choice between an invariant reference trip and two hypothetical alternatives, where each of the three alternatives is described by travel time, cost, rate of crowded trips, rate of delays (both out of 10 trips), the average length of delays (entered into models both as the average extent of delays, RA, and as the expected delay, RB, by multiplying the length of delays by the rate of delays) and the provision of a delay information service (none used as the base, with parameters for a charged, ICH, and free, IFR, service). Following earlier results by Hess and Stathopoulos [2013], we applied a log-transform to the fare attribute (described as LF).

Table 6 shows the results for the second SC dataset. In the presence of three alternatives, we can now include a RRM model alongside MNL, where we see fairly similar performance for these two models, with a slight advantage for MNL. All parameters have the expected sign in these models, and we are also able to include two alternative specific constants (ASCs), which results in improvements in log-likelihood of 46 and 61 units respectively for DFT and MLBA.

For DFT, we follow the same specification tests as on our first SC dataset. However, this time constraining the feedback matrix to be an identity matrix (as in model 2) now clearly leads to a significant drop in model fit[14]. This is a direct result of having more than two alternatives, meaning that the feedback matrix is needed for capturing the different similarities between alternatives.

For MLBA, we again show that constraining $\chi = 2$ is not appropriate, leading to a loss of fit for model 2. We are however able to constrain it to $\chi = 1$ and in addition can constrain $I_0 = 0$ (model 3) without a significant loss of fit. Higher values are observed for λ_2 compared to λ_1, meaning that a greater importance weight is given to positive attribute differences $x_{k,j} \geq x_{k,i}$ compared to negative ones, $x_{k,j} < x_{k,i}$ in the estimation of the mean drift rates.

In terms of model performance, we see that DFT and MLBA again both outperform MNL (and also RRM), where, with the present data, DFT offers better performance than MLBA, potentially as it is better able to deal with the differential competition between the three alternatives than MLBA (in contrast to the earlier binary dataset).

3.3. RP data

Whilst both DFT and MLBA have been used extensively on experimental data and have been shown to accurately explain choices in stated preference surveys, as far as we are aware, neither model has been fitted to revealed preference (RP) data. In this section, we first fit MNL, RRM, DFT and MLBA models to our full RP dataset. We then provide elasticities as well as additionally

[14] Note that in this case, a model with $\sigma_\varepsilon = 0$ as demonstrated to be effective for the first dataset, has a log-likelihood of -3,306.61, thus this time does result in worse fit than a model with $\sigma_\varepsilon = 1$.

TABLE 6 : Estimation results and identifications tests on the second SC dataset

Model		MNL	RRM	DFT		MLBA		
Version		1	1	**1**	2	1	2	**3**
Free Pars.		10	10	**13**	11	14	13	**13**
Log-likelihood		-3,360.43	-3,363.91	**-3,299.82**	-3,327.28	-3,321.90	-3,331.50	**-3,322.26**
BIC		6,802.97	6,809.92	**6,706.41**	6,744.88	6,758.75	6,769.73	**6,743.04**
β_{TT}	est.	-0.0471	-0.0320	-0.3749	-0.1321	-0.0626	-0.0300	-0.0753
	r. t-rat.	-9.50	-9.58	-3.32	-5.20	-7.88	-2.02	-8.59
β_{LF}	est.	-5.9990	-4.1090	-53.9560	-17.5868	-11.4724	-6.3619	-14.0066
	r. t-rat.	-18.87	-17.66	-3.31	-5.69	-6.73	-1.67	-13.57
β_{CR}	est.	-0.2230	-0.1212	-1.7221	-0.6618	-0.2878	-0.1560	-0.3600
	r. t-rat.	-8.58	-5.82	-3.22	-4.45	-9.03	-2.16	-9.44
β_{RA}	est.	-0.1870	-0.0441	-1.2248	-0.5488	-0.1634	-0.0740	-0.1864
	r. t-rat.	-5.96	-2.71	-2.81	-2.98	-2.47	-3.01	-1.29
β_{RE}	est.	-0.0619	-0.1457	-0.7332	-0.1750	-0.1312	-0.0519	-0.1678
	r. t-rat.	-2.64	-8.59	-2.17	-1.15	-2.08	-0.74	-1.10
β_{RB}	est.	-0.0293	-0.0186	-0.1109	-0.0683	-0.0187	-0.0110	-0.0220
	r. t-rat.	-3.25	-3.06	-1.64	-1.77	-7.07	-1.08	-2.82
β_{ICH}	est.	-0.0910	-0.0510	0.0002	-0.1825	-0.0518	-0.0211	-0.0669
	r. t-rat.	-1.13	-0.95	0.03	-0.98	-3.07	-1.16	-1.11
β_{IFR}	est.	0.3305	0.2179	1.8357	0.7773	0.2801	0.1319	0.3180
	r. t-rat.	4.95	4.85	3.25	4.18	8.27	2.57	4.29
δ_1	est.	0.3902	-0.2730	2.2548	2.0355	1.1070	0.5354	1.3043
	r. t-rat.	5.85	-4.17	4.88	7.82	9.48	2.46	12.08
δ_2	est.	0.1633	-0.1656	1.0733	0.6109	0.3518	0.1658	0.3543
	r. t-rat.	3.30	-3.38	3.02	3.30	3.29	1.05	2.74
ϕ_1	est.			0.0003	0.0000			
	r. t-rat.			1.32	**fixed**			
ϕ_2	est.			-0.5560	0.0000			
	r. t-rat.			-4.92	**fixed**			
σ_ε	est.			1.0000	1.0000			
	r. t-rat.			**fixed**	**fixed**			
t	est.			5.1846	7.5332			
	r. t-rat. (*vs* 1)			8.33	6.37			
χ	est.					1.0087	2.0000	1.0000
	r. t-rat. (*vs* 1)					3.33	**fixed**	**fixed**
I_0	est.					0.5747	0.8215	0.0000
	r. t-rat.					0.60	1.08	**fixed**
λ_1	est.					0.1128	0.3386	0.0959
	r. t-rat.					1.57	3.13	7.84
λ_2	est.					0.8777	4.3636	1.1646
	r. t-rat.					7.15	2.99	8.02

testing out-of-sample prediction for all four models.

Our RP data comes from the national UK value of travel time study [Arup, ITS Leeds and Accent, 2015]. Questionnaires were completed by 2,646 individuals travelling by train from Birmingham, Stoke or Peterborough to London. After extensive data cleaning (see page 164 of Arup, ITS Leeds and Accent 2015), 725 observations were left, with either one or two observations for each of the 578 individuals[15]. For every decision recorded, the available alternatives are one or two of Chiltern railways, Northern rail and Midlands railways as well as one of Virgin Trains and East Coast. Travel time, travel cost and headway were used to describe the alternatives.

We run a basic MNL model with a specification based on the model used by Arup, ITS Leeds and Accent [2015], with four different travel time coefficients for different groups. Individuals are first segmented by travel purpose (employees' business, commute (TT_C in Table 7) or 'other non-work' (TT_O)). Individuals on employees' business where further segmented into those who were very sure (TT_{EB1}) and those who were quite sure (TT_{EB2}) about the attributes of the unchosen alternatives. Two further attribute parameters are estimated (travel cost, TC, and headway, HW). For all three attributes, log values are used [Arup, ITS Leeds and Accent, 2015]. Additionally, Arup, ITS Leeds and Accent [2015] use three alternative specific constants for train services run by Chiltern railways (ASC_C), Midlands railways (ASC_M) and Northern rail (ASC_N). Finally, two parameters are incorporated to capture income effects. Travel time coefficients (β_{TT_n}) are calculated for each individual n:

$$\beta_{TT_n} = \beta_{TT_{i,n}} \cdot \left({RI_n}^{\lambda_{inc}} \cdot (1 - z_{miss,n}) + \lambda_{miss} \cdot z_{miss,n} \right) \tag{33}$$

where $\beta_{TT_{i,n}}$ is a travel time coefficient depending on the individual's trip purpose, RI_n is the relative income of the individual, λ_{inc} is an income elasticity on the time sensitivity and λ_{miss} is a multiplier on the time sensitivity used only if the individual did not provide their income in the questionnaire (in which case the dummy variable $z_{miss,n} = 1$). Table 7 provides model estimates for these parameters under MNL, RRM, DFT and MLBA.

For DFT, we again test two versions. With 118 out of 725 observations having three alternatives available and the rest having only two alternatives available, it is unsurprising that, in line with the results from the first SC dataset, the effect of the feedback parameters being removed (DFT model 2 relative to DFT model 1) has little impact on the log-likelihood.

For MLBA, we see that fixing one of the similarity parameters, λ_2, to infinity (which results in the corresponding weight, $w_{x_{k,i,j}} = 0$, when $x_{k,j} < x_{k,i}$) has no impact on model fit (model 2 compared to model 1). This implies that there is no asymmetry between positive and negative differences and that the model captures the choices just as well by only considering positive differences. Additionally fixing both I_0 and χ results in an insignificant impact on model fit with a lower BIC value obtained for model 3 compared to model 2.

With this data, we observe that MLBA obtains a lower BIC value than DFT, with both DFT

[15] Note that this means we have data that is panel data for some individuals, but not others. This does not impact the results of the models as the only impact this has is in the calculation of standard errors.

TABLE 7 : Results, estimates and robust t-ratios from MNL, RRM, DFT and MLBA models on the RP dataset

Model		MNL	RRM	DFT		MLBA		
Version		1	1	1	**2**	1	2	**3**
Free Pars.		11	11	14	**12**	15	14	**12**
Log-likelihood		-370.05	-373.31	-362.53	**-363.31**	-351.13	-351.13	**-352.07**
BIC		812.54	819.07	817.26	**805.66**	801.06	794.47	**781.50**
TT_C	est.	-4.4541	-1.1490	-5.4662	-5.2201	-26.2436	-26.2735	-25.9481
	rob. t-rat.	-3.88	-3.70	-1.43	-1.96	-2.61	-2.60	-2.39
TT_O	est.	-2.0021	-0.5272	-2.4161	-2.4187	-4.6805	-4.6898	-4.3771
	rob. t-rat.	-2.46	-2.47	-1.44	-2.00	-3.36	-3.18	-5.76
TT_{EB1}	est.	-3.7769	-0.9341	-3.7178	-3.6950	-8.2882	-8.3010	-7.8395
	rob. t-rat.	-4.63	-4.66	-1.73	-2.47	-5.48	-5.29	-5.79
TT_{EB2}	est.	-5.7016	-1.4342	-6.3211	-6.1886	-11.7130	-11.7293	-11.2309
	rob. t-rat.	-7.10	-6.85	-1.65	-2.28	-10.61	-7.18	-9.16
TC	est.	-2.2127	-0.6152	-1.9211	-1.8505	-4.8725	-4.8766	-4.8152
	rob. t-rat.	-8.52	-9.27	-2.52	-2.99	-5.62	-4.65	-9.62
HW	est.	-0.1267	-0.0532	-0.1087	-0.1037	-0.0962	-0.0966	-0.1058
	rob. t-rat.	-0.64	-0.75	-0.75	-0.73	-0.67	-0.37	-1.09
δ_C	est.	0.7549	0.6813	2.5936	3.2931	1.5675	1.5697	1.4996
	rob. t-rat.	2.75	2.52	1.61	3.24	5.08	4.98	4.85
δ_M	est.	-0.4882	-0.5251	-0.5357	-0.4285	-0.5397	-0.5393	-0.5910
	rob. t-rat.	-1.86	-1.96	-0.45	-0.41	-2.01	-1.68	-3.03
δ_N	est.	-0.4879	-0.5164	-0.3405	-0.3706	-0.4461	-0.4442	-0.5138
	rob. t-rat.	-1.65	-1.68	-0.28	-0.31	-3.46	-3.07	-1.81
λ_{inc}	est.	0.4563	0.4690	0.5594	0.5411	0.5695	0.5692	0.5849
	rob. t-rat.	4.43	4.46	4.50	4.48	5.13	5.12	4.88
λ_{miss}	est.	0.4844	0.4939	0.8040	0.7948	0.6073	0.6082	0.5831
	rob. t-rat.	1.13	1.18	1.05	1.09	0.65	0.64	0.76
ϕ_1	est.			1.4686	0.0000			
	r. t-rat.			0.47	**fixed**			
ϕ_2	est.			-0.0867	0.0000			
	r. t-rat.			-1.75	**fixed**			
σ_ε	est.			1.0000	1.0000			
	r. t-rat.			**fixed**	**fixed**			
t	est.			8.1692	8.1585			
	r. t-rat. (*vs* 1)			2.55	3.032322			
χ	est.					2.0797	2.0781	2.0000
	r. t-rat. (*vs* 1)					5.87	1.57	**fixed**
I_0	est.					-0.1389	-0.1414	0.0000
	r. t-rat.					-3.73	-1.97	**fixed**
λ_1	est.					0.1033	0.1033	0.1046
	r. t-rat.					12.50	10.43	13.03
λ_2	est.					269.4277	Inf	Inf
	r. t-rat.					0.46	**fixed**	**fixed**

and MLBA outperforming MNL and RRM, thus demonstrating that they work well for RP data as well as SC data.

With a view to not just focussing on model fit, Table 8 contrasts the cost and time elasticities on the RP data for the four models. We see that the elasticities for MNL and RRM are quite similar to each other. MLBA obtains visibly higher time and cost elasticities than MNL and RRM. For DFT, the cost and time elasticities are between MNL/RRM and MLBA. These results again show that DFT and MLBA offer more significant departures from standard models than for example RRM.

TABLE 8 : Cost and time elasticities on RP data

	elasticities	
	cost	time
MNL	-0.537	-0.933
RRM	-0.530	-0.901
DFT	-0.604	-1.017
MLBA	-0.670	-1.278

We finally test all four models for their ability to make out-of-sample predictions. For each of the five data subsets, we take choices corresponding to a random 80% of the individuals in the data to be used for estimation, with the remaining 20% used for validation. We fit each model to each estimation subset and then calculate log-likelihoods for the remaining 20% of the data using the parameter estimates obtained for the first 80%. Table 9 gives the log-likelihoods of the estimation and validation subsets of the data under each model.

TABLE 9 : Out-of-sample estimation and holdout log-likelihoods for the RP data

	MNL (11 pars)		RRM (11 pars)	
	estimated	forecast	estimated	forecast
Dataset 1	-302.88	-68.92	-306.05	-69.27
Dataset 2	-298.59	-72.76	-301.04	-73.59
Dataset 3	-296.70	-75.08	-299.31	-75.76
Dataset 4	-302.29	-68.18	-304.81	-68.84
Dataset 5	-296.64	-75.74	-299.41	-76.28
	DFT (12 pars)		**MLBA (12 pars)**	
	estimated	forecast	estimated	forecast
Dataset 1	-296.90	-67.80	-282.54	-70.30
Dataset 2	-293.80	-70.64	-282.32	-69.93
Dataset 3	-293.12	-71.75	-281.39	-71.36
Dataset 4	-295.41	-68.29	-286.19	-65.94
Dataset 5	-293.23	-72.90	-282.06	-71.76

We see that DFT and MLBA outperform MNL and RRM across all five subsamples in both estimation and performance on the holdout sample except for DFT in holdout sample 4 and MLBA in holdout sample 1. MNL outperforms RRM in estimation and holdout across all samples, while MLBA typically outperforms DFT. Overall, these findings confirm the results on the full sample.

3.4. Comparison of results

To summarise the results, Table 10 shows the BIC for the final recommended specification for each model type on each dataset. We see that DFT and MLBA consistently offer better performance than MNL and RRM. While DFT marginally outperforms MLBA on the Danish SC data, the differences are more substantial on the remaining two datasets, with DFT performing best on the UK SC data and MLBA best on the RP data.

TABLE 10 : Model fit (BIC) comparison across models and datasets

	MNL	RRM	DFT	MLBA
Danish SC	4,457.58	-	4,057.60	4,061.93
UK SC	6,802.97	6,809.92	6,706.41	6,743.04
RP	812.54	819.07	805.66	781.50

An additional benefit of the new scaling method we use for DFT is that it allows us to more directly compare parameter estimates across different models, notwithstanding the different meaning of the parameters. This is possible as a result of the new specifications of both MLBA and DFT having attribute-specific scaling coefficients, which have a role analogous to marginal utility coefficients in RUM models. Although these scaling coefficients cannot be directly translated into measures such as the value of travel time, we can calculate 'relative importance of travel time with respect to travel cost'. In Table 11, we set the calculated MNL ratios of time and cost parameters to a base rate of 1 (with the rates being based on the MNL value for commuters in the RP dataset). Consequently we can compare whether DFT and MLBA assign more or less importance to travel time with respect to travel cost.

TABLE 11 : The relative importance of travel time compared to cost parameter coefficients across different models in comparison to MNL

		MNL	RRM	DFT	MLBA
SP	Danish	1.000	1.000	0.774	0.786
	UK	1.000	0.992	0.885	0.685
RP	Commuters	1.000	0.928	1.401	2.677
	Other Non-Work	0.449	0.426	0.649	0.452
	Employees' Business 1	0.848	0.754	0.992	0.809
	Employees' Business 2	1.280	1.158	1.661	1.159

Across the SP datasets, it appears that MNL tends to assign higher importance to travel time with respect to travel cost relative to DFT and MLBA. The opposite is the case for the RP datasets. RRM always estimates lower ratios than MNL, while DFT has some similar values to MLBA, with key exceptions being the UK data and commuters in the RP data, for which DFT is more similar to MNL.

3.5. DFT model specifications: alternative scaling methods or weights

In this section, we compare our new method (see Section 2.2.2.1) to scaling methods that have been used in previous DFT applications, where we do this for both SP datasets. As DFT is scale-variant (see Section 2.2.2.1), a failure to appropriately adjust the attribute values can result in inferior model fit (see Table 5 of Hancock et al. [2018]). Crucially, all previous methods rely on a priori knowledge of the directionality of the attributes, whereas the new method proposed does not. The different models tested are listed below.

1. Unity-based normalisation, as used by Berkowitsch et al. [2014], where we rescale the attribute values to a range between 0 and 1. For undesirable attribute x, with some value x_k, we define a new attribute value $x'_k = 1 - \frac{x_k - min(x)}{max(x) - min(x)}$, with desirable attribute values being set as $x'_k = \frac{x_k - min(x)}{max(x) - min(x)}$.
2. No scaling method other than taking the negative value for all 'negative' attributes (as DFT can only capture 'positive' effects of attributes as the relative importance weights must be positive[16] - see Section 4.3.1 of Hancock et al. [2018] for an illustration of the results of failing to do this for DFT models)
3. Standard score normalisation, as previously found to be effective for DFT (see results in Hancock et al. [2018]), where for undesirable attribute x, with value x_k, we define a new attribute value $x'_k = -\frac{x_k - mean(x)}{sd(x)}$. For desirable attributes, $x'_k = \frac{x_k - mean(x)}{sd(x)}$.
4. Minimum rescaling (dividing each attribute by the smallest value for that attribute across the choice set), as previously shown to be effective for a previous version of MLBA [Trueblood et al., 2013a]
5. Maximum rescaling (dividing each attribute by the largest value for that attribute across the choice set), as previously shown to be effective for a previous version of MLBA [Trueblood et al., 2013a]
6. Our new method detailed in Section 2.2.2.1, which removes the scale-variant nature of DFT.
7. The new method with both attribute scaling coefficients and attribute importance weights estimated.

For both datasets, it appears that our new method (model 6) has the best model fit if either the scales or the weights are fixed parameters. This result holds regardless of whether we include DFT's feedback matrix. Scale 6 appears to better capture the impact of the feedback matrix for the UK data, resulting in an improvement by more than 40 log-likelihood units, whereas this improvement is much smaller with the other scaling methods corresponding to models 1-5. On the other hand, with the Danish data, the feedback matrix is needed for some of the other scalings to obtain model fits more in line with the new scaling.

[16]Note that here we adjust the attributes accordingly so that our new scaling method does not have an unfair advantage for attributes which have a positive sign.

TABLE 12 : The log-likelihood (LL) values obtained from models for the two stated choice datasets, with different types of scaling for DFT

Dataset			Danish				UK			
Model			with feedback		without feedback		with feedback		without feedback	
No.	wt est.	scale est.	free pars.	LL	free pars.	LL	free pars.	LL	free pars.	LL
1	yes	no	5	-2,020.24	3	-2,020.24	11	-3,404.09	9	-3,405.60
2	yes	no	5	-2,034.24	3	-2,040.94	11	-3,395.50	9	-3,400.88
3	yes	no	5	-2,021.86	3	-2,021.86	11	-3,390.31	9	-3,400.41
4	yes	no	5	-2,112.62	3	-2,112.62	11	-3,419.43	9	-3,420.48
5	yes	no	5	-2,139.50	3	-2,146.38	11	-3,442.14	9	-3,443.22
6	yes	no	5	**-2,017.23**	3	**-2,018.73**	11	**-3,346.23**	9	**-3,387.38**
7	yes	yes	6	-2,014.21	4	-2,014.76	18	-3,338.37	16	-3,382.86

For both datasets, we observe that improvements in model fit can be obtained by estimating both weights and scales. Whilst this improvement is from just one extra parameter for the Danish dataset, the gain in fit comes at a heavy cost for the UK data, which has 7 extra parameters for a relatively insubstantial gain in model fit. This implies that for basic comparisons of DFT against alternative models, analysts may be better off fixing either the weights or the scales. Finally, the results here additionally demonstrate that scales corresponding to models 4 and 5 offer relatively poor performance for DFT, which could in part explain why these scaling methods resulted in MLBA outperforming DFT previously [Trueblood et al., 2013a].

3.6. MLBA model specifications

Table 13 demonstrates the results of applying the different specifications of MLBA models described in Section 2.3.3.

TABLE 13 : Log-likelihoods of models inbetween MNL, RRM and MLBA

Model	Error Structure	Value Function	first SC	second SC	RP
MNL	Eq.30	Eq.31	-2,301.25	-3,360.43	-370.05
Specification 2	Eq.30	Eq.27	-2,105.36	-3,326.43	-358.44
RRM	Eq.30	Eq.32	-2,301.25	-3,363.91	-373.31
Specification 1	Eq.24	Eq.31	-2,279.55	-3,373.57	-395.61
MLBA	Eq.24	Eq.27	-2,005.68	-3,321.87	-351.13
Specification 3	Eq.24	Eq.32	-2,009.78	-3,343.71	-403.42

In all cases, standard MLBA has the best model fit. However, the good performance of specification 2 implies that most of the improvement in model fit for MLBA over MNL and RRM comes from its alternative value function as opposed to having a probability based on drift rates. This is perhaps unsurprising given that Terry et al. [2015] found little differences in fit between LBA models with different drift rate distributions.These results are also displayed in Figure 2, which implies that it is only for the first SC dataset that MLBA has a substantial gain over the alternative models.

4. SIMULATED DATA EXPERIMENTS

The work in Section 3 has provided initial insights about the potential benefits of DFT and MLBA compared to more traditional structures. Of course, these results are dataset specific and the advantages might be a result of the true (and unobserved) data generation process. In this section, we provide some further evidence based on simulated data, where we have a number of aims. In particular, we test the impacts of considering choices generated by different models, compare the ability of the different accumulator models at capturing various complexities in the data, and finally consider parameter recoverability.

FIGURE 2 : Relative model fit of alternative structures in comparison to MLBA across the three datasets.

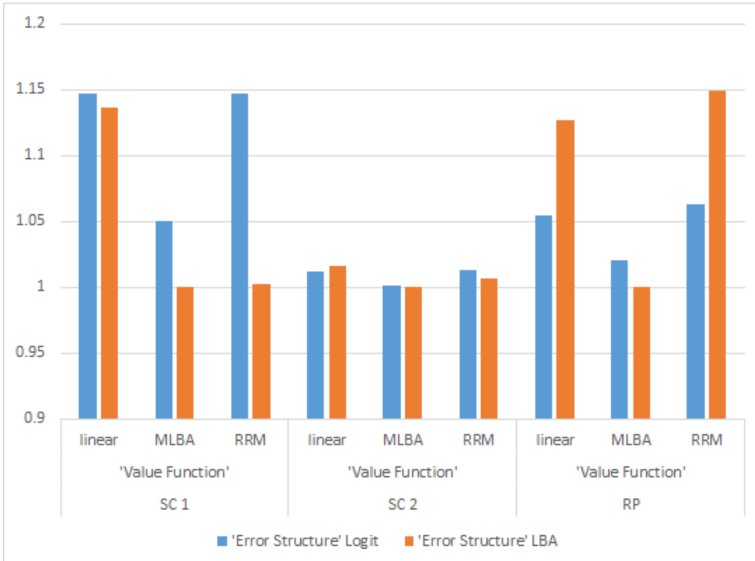

4.1. Generation of simulated data

We use an efficient design to generate 5,000 mode choice observations where each choice task has four alternatives (car, air, rail and high-speed rail), each described by travel cost (TC) and travel time (TT). Additionally, all alternatives other than car have an access time (AT) attribute.

We then generate choices four times using a MNL model, a RRM model, a DFT model and an MLBA model. The aim of this exercise is to see how robust each of the models is to the case where the data stems from a different model.

For our MNL model, we define the utility a respondent n obtains from alternative j in choice task t as:

$$U_{jnt} = ASC_j + ASC_{F_j} \cdot z_{F,n} + \beta_{TT} \cdot \alpha_{TT_j} \cdot TT_{jnt} + \beta_{TC} \cdot TC_{jnt} \cdot \alpha_{IE,n} + \beta_{AT} \cdot AT_{jnt} + \varepsilon_{jnt} \tag{34}$$

where ASC_j and ASC_{F_j} are alternative specific constants, with the latter capturing the difference between male and female participants through the use of an appropriate dummy term, $z_{F,n}$, which takes a value of 1 if individual n is female. TT_{jnt} is the travel time, TC_{jnt} is the travel cost and AT_{jnt} is the access time, all for alternative j in choice situation t for respondent n. There are coefficients for travel cost, access time and mode-specific coefficients for travel time, which are defined as $\beta_{TT} \cdot \alpha_{TT_j}$. A general value β_{TT} is estimated, with appropriate adjustments applied by multiplying by α_{TT_j} for mode j (for identification purposes we fix this coefficient for cars, $\alpha_{TT_{car}} = 1$). We additionally have an income effect, $\alpha_{IE,n}$, which is defined as $\alpha_{IE,n} = (\frac{income_n}{2500})^{\alpha_I}$, where $income_n$ is the income for individual n and α_I is an estimated income elasticity.

These additional coefficients are simple to add in for psychological choice models too following our modifications. For the DFT simulated dataset, we incorporate underlying preferences by setting $P_{0_{jnt}} = ASC_j + ASC_{F_j} \cdot z_{F,n}$, with this having been effective previously (see results in Hancock et al. [2018]). The alternative specific travel time coefficients can be included in DFT and MLBA by multiplication of the attribute values, as we use our new scaling method (see Section 2.2.2.1) which means that these coefficients will have an equivalent impact on the attributes in DFT and MLBA as they would in a RUM model. Finally, in the MLBA models, we incorporate alternative specific constants (δ_j) by adding them to the mean drift rates as in Equation 25. All of the values used for the parameters to generate probabilities for each alternative are given in Table 15.

4.2. Results for simulated data

We next test the performance of the different models across the four datasets, i.e. seeing also how well each model performs on data generated with a different model, thus giving an indication of robustness to the underlying data generation process. The log-likelihood and BIC values obtained from these models are displayed in Table 14.

TABLE 14 : The log-likelihood and BIC values obtained from models for the simulated datasets

Model	free pars.	dataset							
		MNL		RRM		DFT		MLBA	
		LL	BIC	LL	BIC	LL	BIC	LL	BIC
MNL	13	**-4,842.60**	**9,795.92**	-4,773.88	9,658.48	-4,907.49	9,925.70	-4,991.23	10,093.18
RRM	13	-4,853.38	9,817.48	**-4,727.57**	**9,565.86**	-4,923.43	9,957.58	-5,007.51	10,125.74
DFT	16	-4,847.94	9,832.16	-4,751.80	9,639.88	**-4,853.03**	**9,842.33**	-4,934.57	10,005.42
MLBA	17	-4,845.31	9,835.42	-4,741.07	9,626.92	-4,874.60	9,894.00	**-4,930.17**	**10,005.12**

The main difference between the RRM and MNL datasets compared to the MLBA and DFT datasets is that there are parameters for competition between psychologically similar alternatives in the MLBA and DFT models. It appears that MNL and RRM cannot capture this effect and thus have worse model fits for these datasets[17]. Crucially, DFT and MLBA also outperform RRM for the MNL dataset and outperform MNL for the RRM dataset.

These results are highlighted by the comparison of each of the four models with that of the model type used for data generation, demonstrated in Figure 3. This shows that DFT and MLBA show much smaller differences in fit compared to the model consistent with the data generating process (DGP) and suggests that they are more robust to potential misspecification. MNL and RRM, by contrast, perform poorly on the datasets generated by DFT and MLBA, with MNL also having poor fit on the RRM data.

4.3. Recovery of parameters from simulated datasets

We next consider how well the different models recover the parameter values that were used to generate the simulated datasets for the same model. Table 15 gives the parameters used in simulating the data (labelled as 'setup') as well as the parameters produced in estimation, and the

[17]Note that this is also suggested by the fact that the removal of DFT's feedback matrix results in a loss of 10.28 log-likelihood units for the DFT dataset, but only 0.56 units for the MNL dataset.

FIGURE 3 : Log-likelihood of estimated models compared to model consistent with data generating process (DGP)

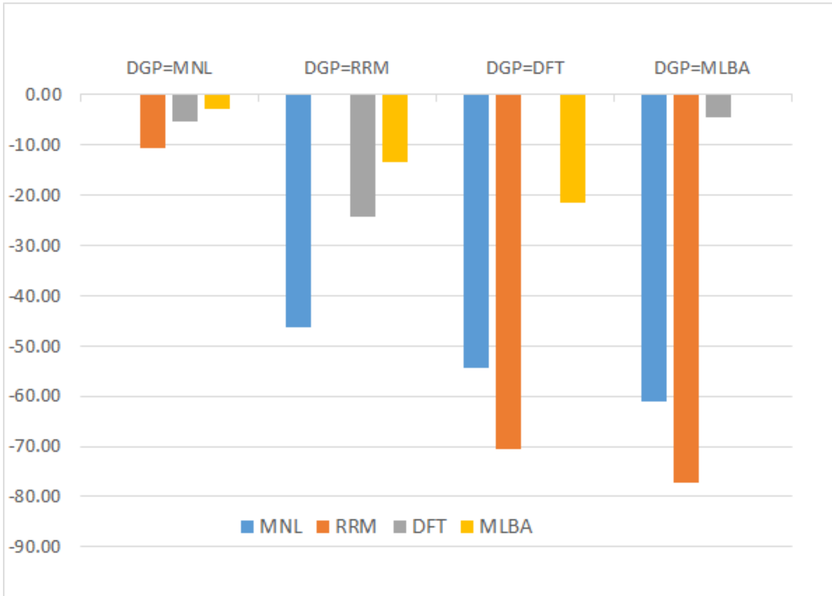

difference between those two. As each model is tested against a dataset generated by the same model, we can test the stability of the parameters. Using our new scaling method allows us to use similar parameter setup values across models, with the exception that parameters are adjusted such that the data generation process has similar amounts of noise across all datasets no matter which model is used to generate the choices.

All four models appear to accurately recover the three β-coefficients associated with the explanatory variables. These appear to be more recoverable than the alternative specific constants. All four models, however, additionally perform well at recovering the attribute-specific travel time coefficients. Most importantly for DFT and MLBA, the process parameters are fairly well recovered too.

5. CONCLUSIONS

In this paper, we have considered two alternate accumulator choice models developed in mathematical psychology and compared them against models typically used in choice modelling. The models in question are decision field theory (DFT), a model where preferences for alternatives stochastically update over time, and the multi-attribute linear ballistic accumulator (MLBA), where the preferences for alternatives *race* towards a threshold.

We first made a number of methodological developments to improve the suitability of the models for studying travel behaviour and other non-laboratory based choices. For DFT, we implemented a new scaling method on the attributes, which results in a number of benefits such as the modeller not having to know the sign of the impact of the attributes before running the model. This has an immediate benefit for the UK dataset, for which one attribute (whether the delay informa-

TABLE 15 : Parameter values used to generate datasets and estimates for full models for their respective datasets

	MNL			RRM			DFT			MLBA		
Parameter	Setup	Estimate	Bias	Setup	Estimate	Bias	Setup	Estimate	Bias	Setup	Estimate	Bias
β_{TT}	-0.0050	-0.0044	-12%	-0.0030	-0.0029	-3%	-0.0050	-0.0049	-2%	-0.0030	-0.0037	23%
β_{TC}	-0.0280	-0.0279	0%	-0.0160	-0.0162	1%	-0.0280	-0.0304	9%	-0.0160	-0.0169	6%
β_{AT}	-0.0060	-0.0053	-12%	-0.0040	-0.0046	15%	-0.0060	-0.0057	-5%	-0.0040	-0.0039	-3%
δ_{car}	-0.5000	-0.8238	65%	-0.5000	-0.6965	39%	-0.5000	-1.4179	184%	-0.5000	0.68986	-238%
δ_{air}	-1.5000	-1.8053	20%	-1.5000	-1.8363	22%	-1.5000	-2.7631	84%	-1.5000	-1.5054	0%
δ_{rail}	-1.0000	-0.9036	-10%	-1.0000	-1.0067	1%	-1.0000	-1.8977	90%	-1.0000	0.62732	-163%
$\delta_{car_{fem}}$	-0.5000	-0.4752	-5%	-0.5000	-0.4020	-20%	-0.5000	-0.6257	25%	-0.5000	-0.6254	25%
$\delta_{air_{fem}}$	0.5000	0.6952	39%	0.5000	0.6822	36%	0.5000	0.8528	71%	0.5000	0.57662	15%
$\delta_{rail_{fem}}$	1.0000	1.1188	12%	1.0000	1.1855	19%	1.0000	-0.1264	-113%	1.0000	1.21153	21%
$\beta_{TT_{air}}$	1.2500	1.1041	-12%	1.2500	1.7576	41%	1.2500	1.1109	-11%	1.2500	1.07889	-14%
$\beta_{TT_{rail}}$	2.0000	2.3845	19%	2.0000	2.2579	13%	2.0000	1.9182	-4%	2.0000	2.24317	12%
$\beta_{TT_{hsr}}$	1.5000	1.7723	18%	1.5000	1.3528	-10%	1.5000	1.8306	22%	1.5000	1.17701	-22%
α_I	-0.5000	-0.5106	2%	-0.5000	-0.4962	-1%	-0.5000	-0.3585	-28%	-0.5000	-0.4553	-9%
ϕ_1							0.0500	0.03646	-27%			
ϕ_2							0.1000	0.138	38%			
σ_ε							1.4142	1.4142	**fixed**			
t							10.0000	8.83029	-12%			
A										1.0000	1.0000	**fixed**
χ										2.0000	2.009	0%
s										2.0000	2.0000	**fixed**
I_0										10.0000	11.0373	10%
λ_1										0.1000	0.06936	-31%
λ_2										0.2000	0.17905	-10%

tion service is free) is a desirable attribute. A comparison with other available scaling approaches in Section 3.5 also highlights the benefits of this approach.

We also considered the impacts of including parameters to capture underlying preferences towards specific alternatives in MLBA and DFT. Results from our UK dataset suggest that MLBA and DFT make substantial gains when these parameters are included and can consequently capture status quo biases. We have, however, only considered one method for incorporating preferences in these models. Whilst we add parameters to the drift rate in MLBA, alternative specifications would allow for an adjustment of the starting point A or the threshold χ, such that alternatives have different values for these parameters. It is easily possible that some alternatives may not require as much evidence to be chosen (for example, a commuter's usual route to work), meaning that an MLBA model including alternative specific thresholds may work well. This could be investigated in future research, with, for example, work on accumulator models with collapsing thresholds already popular [Bowman et al., 2012, Hawkins et al., 2015, Evans et al., 2019]. We additionally only test MLBA with truncated normal drift rate distributions and simplified Fréchet-distributed drift rates (with start rate parameter $A = 0$), with further operationalisations possible with alternative distributions [Terry et al., 2015]. However, results from comparisons of models that have features of MLBA and features of MNL/RRM models suggest that the main improvement in model fit for MLBA is from the alternative value function. This implies that further tests of a utility model with MLBA value functions are required.

The operationalisation of the two models in this paper provides promising results, and paves the way for the incorporation of data on the processes of decision-making in these models, such as eye-tracking information, response times and EEG data.

We also consider in detail the relative importance of different parameters of our models. Whereas additionally fixing the threshold parameter for MLBA does not have a significant impact for our simulated datasets, it does have an impact for our SP data. The opposite is true for the drift rate constant, I_0, which is important for our simulated datasets but is less important for our SP data. It is possible that the importance of these parameters varies according to how deterministic the data is and further work could test datasets with specified variations in the level of noise. This could help an analyst determine which parameters are important for MLBA for complex choice data. For DFT, it appears that our new method for the scaling of attributes significantly improves the impact of the feedback matrix parameters. It appears that the feedback matrix is not relevant for choices where there are only two alternatives. However, regardless of whether the feedback matrix has an impact or not, DFT outperforms MNL and RRM for our SP and RP datasets.

We test the models extensively using simulated data, where the findings suggest that DFT and MLBA may be less sensitive to model misspecification (i.e. if the estimated model differs substantially from that used for data generation) than the corresponding RUM and RRM models. Crucially, both DFT and MLBA outperform MNL and RRM across the two SP datasets and the RP dataset, including in out of sample validation for the latter, which is to the best of our knowledge the first use of both DFT and MLBA on RP data. The good model fits for both DFT and MLBA for our second stated survey dataset suggest that if there is competition between psychologically similar alternatives (when there are two alternatives that have attributes that are more similar than those of a third alternative), a move towards a choice model with psychological foundations becomes more appealing.

Moving away from RUM has obvious pitfalls, especially in terms of the use of models for welfare analysis [see e.g. Hess et al., 2018]. The evidence in this paper suggests that if an analyst is willing to accept these pitfalls, then moving further away from RUM than for example with a RRM model, may be beneficial, and models from mathematical psychology provide an interesting avenue for such work. Of course, more research is needed in terms of additional comparisons, including on larger datasets with more alternatives and attributes. Also, whilst we have considered DFT and MLBA, future research should also consider models from mathematical psychology that do not have closed-form likelihood functions. A large number of models from mathematical psychology such as the drift diffusion model [Wiecki et al., 2013], the leaky competing accumulator [Usher and McClelland, 2001] and the feed-forward inhibition model [Turner et al., 2016] can be estimated using hierarchical Bayesian estimation combined with probability density approximation [Turner and Sederberg, 2014]. This means that there is large scope for further comparisons between psychological and mainstream choice models using hierarchical Bayesian estimation, a method already popular in traditional choice modelling for mixed logit models [Train, 2001, Burda et al., 2008, Dumont et al., 2015, Akinc and Vandebroek, 2018].

Additionally, the linear ballistic accumulator [Brown and Heathcote, 2008], a simplified version of MLBA for alternatives without multiple attributes, has been demonstrated to work well with dynamic datasets where the drift rates change over time [Holmes et al., 2016]. A similar concept could be applied to both DFT and MLBA, for which changing attributes could easily be incorporated. Thus DFT and MLBA may work well with dynamic revealed preference datasets such as the lane merging decisions made by drivers, where typical choice models may not do so well due to their static nature. Complex datasets such as these, as well as datasets with additional process or psychometric measures, would also be useful for further testing the functionality and usefulness of the process parameters within both DFT and MLBA. Additionally, given that in Hancock et al. [2018], we demonstrate that DFT can efficiently incorporate random parameters, it is possible that similar adjustments could also be made for MLBA. All of these potential extensions of DFT and MLBA, combined with the results in this paper, demonstrate that accumulator models such as DFT and MLBA are attractive alternative approaches to random utility models, particularly when it comes to forecasting. It therefore appears that these models, as well as others, may hold significant promise in improving the behavioural realism in choice models, in both transport and beyond.

ACKNOWLEDGEMENTS

The authors acknowledge the financial support by the European Research Council through the consolidator grant 615596-DECISIONS. The data used for the revealed preference analysis comes from a study conducted for the UK Department for Transport. The results reported, findings inferred, and views expressed are the responsibility of the authors, and not a statement of the Department's policy or guidance. We would also like to thank Jennifer Trueblood and two anonymous reviewers for their comments on a previous version of this paper.

REFERENCES

Akinc, D. and Vandebroek, M. (2018). Bayesian estimation of mixed logit models: Selecting an appropriate prior for the covariance matrix. *Journal of choice modelling*, 29:133–151.

Arup, ITS Leeds and Accent (2015). Provision of market research for value of time savings and reliability. Phase 2 report to the Department for Transport.

Berkowitsch, N. A., Scheibehenne, B., and Rieskamp, J. (2014). Rigorously testing multialternative decision field theory against random utility models. *Journal of Experimental Psychology: General*, 143(3):1331.

Berkowitsch, N. A., Scheibehenne, B., Rieskamp, J., and Matthäus, M. (2015). A generalized distance function for preferential choices. *British Journal of Mathematical and Statistical Psychology*, 68(2):310–325.

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, 120(3):522.

Bierlaire, M., Thémans, M., and Zufferey, N. (2010). A heuristic for nonlinear global optimization. *INFORMS Journal on Computing*, 22(1):59–70.

Bowman, N. E., Kording, K. P., and Gottfried, J. A. (2012). Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, 75(5):916–927.

Brown, S. D. and Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3):153–178.

Bruno, H. A., Cebollada, J., and Chintagunta, P. K. (2015). Brand preference and state-dependence with intra-household heterogeneity.

Burda, M., Harding, M., and Hausman, J. (2008). A bayesian mixed logit–probit model for multinomial choice. *Journal of Econometrics*, 147(2):232–246.

Busemeyer, J. R. and Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Sciences*, 43(3):345–370.

Busemeyer, J. R. and Townsend, J. T. (1992). Fundamental derivations from decision field theory. *Mathematical Social Sciences*, 23(3):255–282.

Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432.

Cataldo, A. M. and Cohen, A. L. (2018). Reversing the similarity effect: The effect of presentation format. *Cognition*, 175:141–156.

Chernev, A. (2004). Extremeness aversion and attribute-balance effects in choice. *Journal of consumer research*, 31(2):249–263.

Chorus, C. G. (2010). A new model of random regret minimization. *EJTIR, 10 (2), 2010*.

Chorus, C. G., Arentze, T. A., and Timmermans, H. J. (2008). A random regret-minimization model of travel choice. *Transportation Research Part B: Methodological*, 42(1):1–18.

Cohen, A. L., Kang, N., and Leise, T. L. (2017). Multi-attribute, multi-alternative models of choice: Choice, reaction time, and process tracing. *Cognitive Psychology*, 98:45–72.

Daly, A., Hess, S., and de Jong, G. (2012). Calculating errors for measures derived from choice modelling estimates. *Transportation Research Part B*, 46(2):333–341.

Diederich, A. (2003). Mdft account of decision making under time pressure. *Psychonomic Bulletin & Review*, 10(1):157–166.

Dumont, J., Keller, J., and Carpenter, C. (2015). Rsghb: functions for hierarchical bayesian estimation: a flexible approach. *R package version*, 1(2).

Eddelbuettel, D., François, R., Allaire, J., Ushey, K., Kou, Q., Russel, N., Chambers, J., and Bates, D. (2011). Rcpp: Seamless r and c++ integration. *Journal of Statistical Software*, 40(8):1–18.

Evans, N. J., Hawkins, G. E., and Brown, S. D. (2019). The role of passing time in decision-making. *Journal of experimental psychology: learning, memory, and cognition*.

Fosgerau, M. (2006). Investigating the distribution of the value of travel time savings. *Transportation Research Part B: Methodological*, 40(8):688–707.

Giergiczny, M., Hess, S., Dekker, T., and Chintakayala, P. K. (2013). Testing the consistency (or lack thereof) between choices in best-worst surveys. In *Third International Choice Modelling Conference, Sydney*.

Gold, J. I. and Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404(6776):390–394.

González-Valdés, F. and Ortúzar, J. d. D. (2017). The stochastic satisficing model: A bounded rationality discrete choice model. *Journal of Choice Modelling*, 27:74–87.

Hancock, T. O., Hess, S., and Choudhury, C. F. (2018). Decision field theory: Improvements to current methodology and comparisons with standard choice modelling techniques. *Transportation Research Part B: Methodological*, 107:18–40.

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., and Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35(6):2476–2484.

Hawkins, G. E., Marley, A., Heathcote, A., Flynn, T. N., Louviere, J. J., and Brown, S. D. (2014). Integrating cognitive process and descriptive models of attitudes and preferences. *Cognitive Science*, 38(4):701–735.

Heathcote, A. and Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in psychology*, 3:292.

Henningsen, A. and Toomet, O. (2011). maxlik: A package for maximum likelihood estimation in R. *Computational Statistics*, 26(3):443–458.

Hess, S., Daly, A., and Batley, R. (2018). Revisiting consistency with random utility maximisation: theory and implications for practical work. *Theory and Decision*, 84(2):181–204.

Hess, S. and Palma, D. (2019). Apollo: a flexible, powerful and customisable freeware package for choice model estimation and application, www.apollochoicemodelling.com.

Hess, S., Rose, J. M., and Hensher, D. A. (2008). Asymmetric preference formation in willingness to pay estimates in discrete choice models. *Transportation Research Part E: Logistics and Transportation Review*, 44(5):847–863.

Hess, S. and Stathopoulos, A. (2013). A mixed random utility - random regret model linking the choice of decision rule to latent character traits. *Journal of Choice Modelling*, 9:27–38.

Holmes, W. R., Trueblood, J. S., and Heathcote, A. (2016). A new framework for modeling decisions about changing information: The piecewise linear ballistic accumulator model. *Cognitive Psychology*, 85:1–29.

Hotaling, J. M., Busemeyer, J. R., and Li, J. (2010). Theoretical developments in decision field theory: comment on Tsetsos, Usher, and Chater (2010). *Psychological Review*, 117(4):1294–1298.

Kaufman, B. E. (1990). A new theory of satisficing. *Journal of Behavioral Economics*, 19(1):35–51.

Krajbich, I., Lu, D., Camerer, C., and Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in psychology*, 3:193.

Liu, Y. and Cirillo, C. (2018). A generalized dynamic discrete choice model for green vehicle adoption. *Transportation Research Part A: Policy and Practice*, 114:288–302.

Masiero, L. and Hensher, D. A. (2010). Analyzing loss aversion and diminishing sensitivity in a freight transport stated choice experiment. *Transportation Research Part A: Policy and Practice*, 44(5):349–358.

McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behaviour. In Frontiers in Econometrics, ed. P. Zarembka. (New York: Academic Press).

McNair, B. J., Hensher, D. A., and Bennett, J. (2012). Modelling heterogeneity in response behaviour towards a sequence of discrete choice questions: a probabilistic decision process model. *Environmental and Resource Economics*, 51(4):599–616.

Noguchi, T. and Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition*, 132(1):44–56.

Otter, T., Johnson, J., Rieskamp, J., Allenby, G. M., Brazell, J. D., Diederich, A., Hutchinson, J. W., MacEachern, S., Ruan, S., and Townsend, J. (2008). Sequential sampling models of choice: Some recent advances. *Marketing letters*, 19(3-4):255–267.

Qin, H., Guan, H., and Wu, Y.-J. (2013). Analysis of park-and-ride decision behavior based on decision field theory. *Transportation Research Part F: Traffic Psychology and Behaviour*, 18:199–212.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59.

Roe, R. M., Busemeyer, J. R., and Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108(2):370.

Rose, J. M. (2014). Interpreting discrete choice models based on best-worst data: a matter of framing. *The Transportation Research Board (TRB) 93rd Annual Meeting*.

Sanderson, C. and Curtin, R. (2016). Armadillo: a template-based c++ library for linear algebra. *Journal of Open Source Software*.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., and Lehman, D. R. (2002). Maximizing versus satisficing: happiness is a matter of choice. *Journal of Personality and Social Psychology*, 83(5):1178.

Stathopoulos, A. and Hess, S. (2012). Revisiting reference point formation, gains–losses asymmetry and non-linear sensitivities with an emphasis on attribute specific treatment. *Transportation Research Part A: Policy and Practice*, 46(10):1673–1689.

Stewart, N. and Simpson, K. (2008). A decision-by-sampling account of decision under risk. *The probabilistic mind: Prospects for Bayesian cognitive science*, pages 261–276.

Swait, J. (2001). A non-compensatory choice model incorporating attribute cutoffs. *Transportation Research Part B: Methodological*, 35(10):903–928.

Terry, A., Marley, A., Barnwal, A., Wagenmakers, E.-J., Heathcote, A., and Brown, S. D. (2015). Generalising the drift rate distribution for linear ballistic accumulators. *Journal of Mathematical Psychology*, 68:49–58.

Train, K. (2001). A comparison of hierarchical bayes and maximum simulated likelihood for mixed logit. *University of California, Berkeley*, pages 1–13.

Trueblood, J., Brown, S., and Heathcote, A. (2013a). The multi-attribute linear ballistic accumulator model of decision-making. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 35.

Trueblood, J. S., Brown, S. D., and Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological review*, 121(2):179–205.

Trueblood, J. S., Brown, S. D., Heathcote, A., and Busemeyer, J. R. (2013b). Not just for consumers context effects are fundamental to decision making. *Psychological science*, 24(6):901–908.

Turner, B. M., Schley, D. R., Muller, C., and Tsetsos, K. (2018). Competing theories of multialternative, multiattribute preferential choice. *Psychological review*, 125(3):329.

Turner, B. M. and Sederberg, P. B. (2014). A generalized, likelihood-free method for posterior estimation. *Psychonomic bulletin & review*, 21(2):227–250.

Turner, B. M., Sederberg, P. B., and McClelland, J. L. (2016). Bayesian analysis of simulation-based models. *Journal of Mathematical Psychology*, 72:191–199.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4):281.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4):327.

Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232.

Usher, M. and McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108(3):550.

Wiecki, T. V., Sofer, I., and Frank, M. J. (2013). HDDM: hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in neuroinformatics*, 7:14.

A. APPENDIX: ALTERNATIVE VERSIONS OF MLBA

Whilst we use the mainstream version of MLBA [Trueblood et al., 2014] in this paper, it should be noted that the original version of MLBA [Trueblood et al., 2013a] has also not been tested on large-scale consumer choice data. Whilst this version of MLBA, here denoted '$MLBA_0$' uses the same start, threshold and standard deviation for its drift rates, it differs in the specification for the value of the mean drift rate:

$$d_j = \frac{10}{1 + exp(-\gamma \cdot v_j)} \tag{35}$$

where v_j is a valence function and γ is a logistic parameter. Small values of the logistic parameter γ would result in $exp(-\gamma \cdot v_j) \to 1$, meaning that the valences, v_j, are less influential and the probabilities of the alternatives become more similar, resulting in a less deterministic choice. The valences are similar to a decision field theory model's valences with the exception that they attempt to additionally capture the comparison process achieved by DFT's feedback matrix. Thus we have

$$V = C \cdot M \cdot W \tag{36}$$

where W is a vector comprising of a set of attribute importance weights that sum to 1, M is the attribute matrix and C is a $n \times n$ comparison matrix (n being the number of alternatives) with diagonal entries of 1 and off-diagonal elements:

$$C_{i,j \neq i} = \frac{exp(-\phi \cdot Dist_{i,j}) - 1}{n-1}. \tag{37}$$

Finally, ϕ is a sensitivity parameter such that high values result in the distance between the attributes of the alternatives becoming insignificant. Low values allow for more similar alternatives to compete more with each other relative to less similar alternatives.

Results from applying the previous version of MLBA to both of the SP datasets and the RP dataset are given in Table 16 below.

TABLE 16 : Comparison of different versions of MLBA

Dataset	$MLBA_0$	MLBA	Difference
Danish	-2,189.78	-2,005.92	-183.86
UK	-3,394.36	-3,324.42	-69.94
RP	-375.24	-352.07	-23.17

From these results, it appears that the old version of MLBA has far inferior fits compared to that of the mainstream MLBA. Consequently, it would appear that modellers should focus on the mainstream version of MLBA.

B. APPENDIX: A MINIMUM MEAN DRIFT RATE

Originally, the mean drift rate in MLBA was specified as:

$$d_j = v_j + I_0 \tag{38}$$

As well as choosing to add an alternative specific component here, we also choose to truncate the mean drift rate such that it is always positive:

$$d_j = max(0, \delta_j + v_j + I_0), \tag{39}$$

To consider the impact of not restricting mean drift rates to be positive, we consider the probabilities generated by drift rates i and j where these have values of $-5 < i < 5$ and $i < j < i + 10$. Figure 4 gives the probability of choosing alternative i given the use of these drift rate values as well as values of $A = 1$, $\chi = 2$ and $s = 1$.

FIGURE 4 : The probability of picking alternative i as the mean drift rate values change

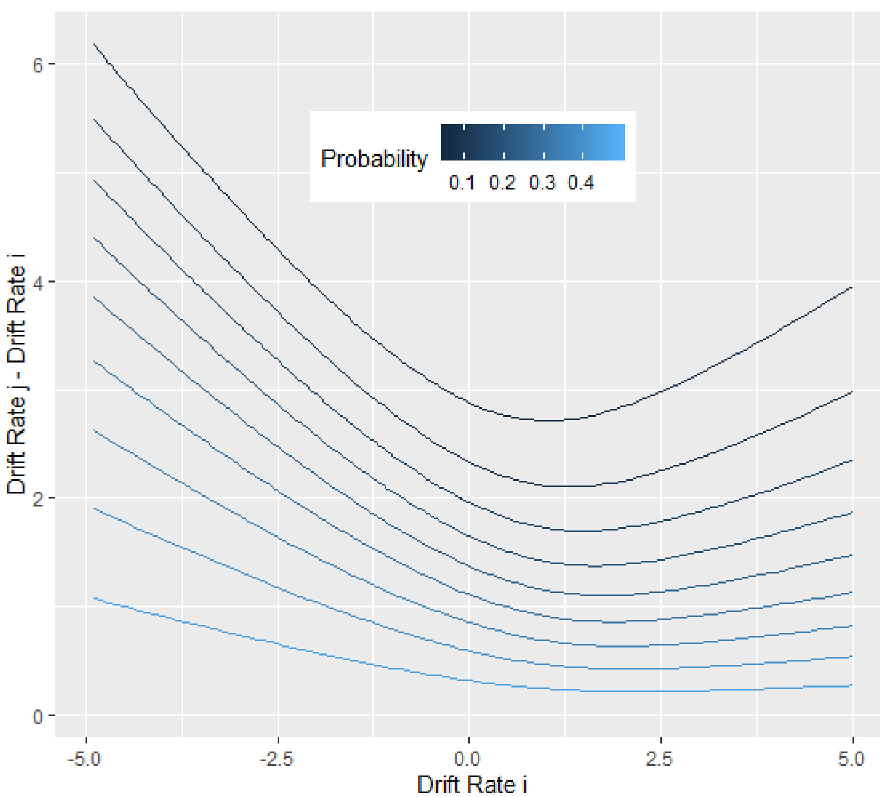

In this figure, the contours demonstrate the points at which probabilities remain the same. If these lines were horizontal, it would indicate that for MLBA, only differences in drift rates matter (much as only differences in utility matter in RUM). However, the presence of curved lines indicates that at some points adding I_0 will increase the probability of alternative i, whereas it will decrease the probability at other points. This can cause convergence issues for MLBA. We find that fixing the mean drift rates to have a minimum value of zero reduces the impact of this issue and improves estimation of MLBA models.