A CAREFUL RESPONDENT OR AN UNCERTAIN RESPONSE: DISENTANGLING CONFOUNDING SOURCES OF INCREASED DELIBERATION TIME USING DECISION FIELD THEORY

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
tstoh@leeds.ac.uk

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

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

ABSTRACT

Decision field theory (DFT), although popular in mathematical psychology, has only recently been used in choice modelling for consumer and travel choices. A key difference that DFT has from standard choice models is that it has preference values for each alternative that update over time. This results in a different probability of picking each alternative depending on how long a decision-maker considers their alternatives. However, the computational complexities of DFT have resulted in failures to utilise its dynamic nature. Recent advances in the underlying computational methods for DFT have allowed for the analytical calculation of the probability of alternatives at any time point. Consequently, the number of preference accumulation steps can be estimated as a function of choice response time. We demonstrate that the model fit for DFT models can be improved by considering response times across route, accommodation and conservation programme choice contexts. We also explore the confounding nature of choice response time, with a key assumption within DFT and other accumulation models being that preference grows over time, contradicting a well-known result that a longer response time often indicates a less certain and hence less deterministic choice from a decision-maker. In line with DFT and preference accumulation, we find that across all three datasets, a longer mean response time indicates that a decision-maker is more deterministic. However, within a decision-maker, our models suggest that fast decisions are typically more deterministic, demonstrating that a longer response time indicates a less certain decision. Furthermore, results from multinomial logit (MNL) models suggest that DFT's time parameters performs similarly to a MNL's scale parameter. This suggests that without further consideration of measures such as choice certainty, DFT may struggle to truly capture the process of preference updating, as the time parameters are not necessarily correlated with response time and instead simply capture how deterministic a choice is.

1. INTRODUCTION

Decision field theory (DFT), first developed in the 1990s (Busemeyer and Townsend, 1992, 1993) is a dynamic, stochastic choice model in which the preference for each alternative updates over time, changing the probability with which each alternative is chosen. These preference values update at each accumulation step as the decision-maker considers the different attributes of the alternatives. The decision-maker then comes to a conclusion either when the preference value for an alternative reaches some satisfactory internal threshold value (equivalent to satisficing (Kaufman, 1990; Schwartz et al., 2002), where a participant chooses one of the alternatives if it is 'good enough') or when the decision-maker reaches some external threshold (such as running out of time). For example, a voter may not have a firm preference for a candidate but be forced to make a decision on election day. At this point the alternative with the highest preference value is chosen.

An analyst needs to decide whether to use DFT with internal or external thresholds. DFT was initially used as a model for understanding risky choice decisions, with internal thresholds used if the decision-maker chose when to stop deliberating, and external thresholds used if a time restriction was imposed. Both versions were developed for two alternatives with two attributes but DFT with an external threshold has since been expanded to allow for multiple attributes and multiple alternatives, and renamed multialternative decision field theory (MDFT, Roe et al. 2001). Decision field theory in various forms has been used widely across the mathematical psychology literature, having been used to model a variety of choices including monetary gambles (Schall, 2003), decision-making in sport (Raab and Johnson, 2004), likely crime suspects (Trueblood et al., 2014) and consumer decisions (Noguchi and Stewart, 2014). However, it has only recently been compared to models developed in econometric choice modelling (Berkowitsch et al., 2014). This is in part due to the complexity of calculating the probability of alternatives being chosen under a DFT model. In particular, for a DFT model with internal thresholds, simulation is often required as there is no closed-form solution for the probability for which each alternative is chosen if there are more than two alternatives. Simulation can also be used for MDFT models with an external threshold (Turner et al., 2018). Many applications thus far have chosen to avoid this computationally-intensive procedure by using MDFT and fixing the number of preference accumulation steps to a high value (Tsetsos et al., 2010; Berkowitsch et al., 2014; Trueblood et al., 2014; Cohen et al., 2017). Alternatively, fixing the number of preference accumulation steps to infinity results in a simpler closed-form analytical solution for the probability of alternatives (Berkowitsch et al., 2014). However, crucially, both these approaches lose MDFT's dynamic nature. Thus, there has been a requirement to improve the computational methods behind MDFT such that the number of preference accumulation steps has a better behavioural underpinning and does not need to be set to an arbitrary value. Hancock et al. (2018) proposes computational developments meaning that the probability with which each alternative is chosen can in fact be calculated simply after any number of preference accumulation steps in a MDFT model. The important question that remains is what these preference accumulation steps actually represent mathematically in a choice process. The work in this paper considers how this feature is linked to choice response time.

There have been a number of studies considering choice response time in choice modelling, often with different foci and aims. For example, consistent differences in response times are found depending on the size of the difference between travel times and costs in a choice task compared to those from a reference trip (Börjesson and Fosgerau, 2015), with increased response times for larger travel time differences and decreased response times for larger differences in costs. There have also been suggestions that choice response times reflect how much cognitive effort a participant uses (Rose and Black, 2006), in which case longer response times would suggest more deterministic behaviour. Whilst Qin et al. (2013) demonstrated that the probability of each alternative at each timepoint under a MDFT model could be matched with the proportion with which each alternative was chosen under different time restrictions, this current paper demonstrates for the first time how the decision-maker's response time can be naturally incorporated into a DFT model such that the probability of alternatives being chosen across multiple choices is impacted. A key issue with such an approach is in accounting for the different reasons a choice might take more time. Whilst a complete model would control for such features, factors such as measurement errors (i.e. whether a recorded response time accurately reflects when the decision-maker came to a conclusion), levels of concentration of the decision-maker and choice certainty are all difficult to account for. Whilst the work in this paper does not control for such considerations, we demonstrate how an analyst might begin to restore the dynamic nature of DFT, through the inclusion of response time.

Thus far, it has been typical for analysts to use decision field theory with an internal threshold if the decision-maker is free to choose when to conclude deliberating on their choice and decision field theory with an external threshold if a time limit is imposed. However, we argue that the choice of which version to use is not so simple, with complications for both versions often not accounted for. For example, a strict external threshold does not account for the different speeds at which the different decision-makers may accumulate evidence, nor any differences in method. A certain number of iterations may result in weak evidence in favour of an alternative for some individuals, but far greater evidence for that same alternative for others. Additionally, internal thresholds cannot be precisely measured and are likely to vary vastly both across and within decision-makers. Furthermore, a decision-maker may not choose an alternative because it satisfices, but because they do not wish to deliberate on a choice any longer. This could either be viewed (analytically) as a decision-maker reaching a self-imposed time threshold, or an alternative reaching a lowered preference evidence threshold.

However, regardless of whether the decision-maker stopped due to an internal or external threshold, we know that under DFT model assumptions, the alternative chosen has the highest preference value at the moment when the decision-maker concludes the deliberation process. Whilst we cannot know how many iterations of preference value updating have occurred, it is possible to estimate the number of iterations as a function of the choice response time, which can be recorded. As both versions of DFT have a parameter related to the number of iterations of preference updating, we can thus restore DFT to being a properly dynamic model in which the time taken to make a decision impacts the probability of each alternative being chosen. This paper considers state-of-the-art implementations of both versions of DFT in detail. We demonstrate a number of methods for incorporating the response time in the models, thus beginning the process of restoring DFT to being a properly dynamic model in which the amount of time taken to make a choice influences the probability with which each alternative is chosen. We then provide a number of empirical applications, demonstrating the increased flexibility of having free rather than fixed time parameters, as well as detailing the results of the different methods for including response times in the DFT models. These different models as well as comparisons with multinomial logit models allow us to investigate the nature and precise function of the time parameters in DFT models.

The remainder of this paper is organised as follows. First, we present the methodology behind decision field theory with internal and external thresholds, demonstrating how the probability with which each alternative is chosen can be calculated. Next, we demonstrate how response time can be incorporated into the DFT models. We then give an outline of the latent class and multinomial logit models that are used in this paper. We next detail our empirical applications, where we test the impact of including response time for models on three very different stated choice datasets. Finally, we finish with some conclusions and present directions for future research.

2. METHODOLOGY

In this section we first describe how the probability with which each alternative is chosen can be calculated in decision field theory models, both with external (time) thresholds and internal (evidence) thresholds. We then demonstrate how reparameterising the time parameter results in a natural method for choice response times to be incorporated into the model. Finally, we briefly introduce the latent class and multinomial logit models that are used in the empirical applications of this paper.

2.1. Decision field theory with an external threshold (MDFT)

2.1.1. Basic theory

Multialternative decision field theory is a dynamic model, meaning that preferences for alternatives accumulate and de-accumulate over time. After each preference accumulation step, a new set of preference values can be calculated as follows:

$$P_\tau = S \cdot P_{\tau-1} + V_\tau. \tag{1}$$

The previous values are stored in a column vector, $P_{\tau-1}$, which corresponds to the preference values after $\tau - 1$ preference accumulation steps[1]. These are multiplied by a feedback matrix, S, after which a valence vector V_τ is added. The feedback matrix has two parameters that allow for the attraction, similarity and compromise effects to occur (Roe et al., 2001) and is defined as:

$$S = I - \phi_2 \times \exp(-\phi_1 \times D^2) \tag{2}$$

where ϕ_1 is a sensitivity parameter and ϕ_2 is a memory parameter. D is the distance between the attributes of the alternatives and can include a factor to control for the importance of the different attributes (Hotaling et al., 2010). The Euclidean distance between the attributes can also be used to avoid an additional parameter (Qin et al., 2013). The memory parameter, ϕ_2, allows for preferences to naturally increase or decrease over time. If this value is less than one, the set of preference values

[1]Note that Equation 1 here is equivalent to Equation 1 in Berkowitsch et al. (2014) and is obtained by setting $h = 1$ in Equation 2 in Busemeyer and Diederich (2002). Thus a preference accumulation step here is defined as a discrete step h of size 1.

stabilise for a large number of accumulation steps (Roe et al., 2001; Berkowitsch et al., 2014).The sensitivity parameter, ϕ_1, affects how much the alternatives compete with each other, and allows for more similar alternatives (in terms of attribute levels) to deduct higher amounts of preference from each other. This is in effect a similar concept to a nested logit model (Williams, 1977; Daly and Zachary, 1978; McFadden, 1978), although tests have not yet been carried out to test whether MDFT models capture nesting effects as effectively.

At each preference accumulation step, a decision field theory model assumes that the decision-maker compares a single attribute across all of the alternatives. This results in a random valence vector for step τ, V_τ, which can be calculated as:

$$V_\tau = C \cdot M \cdot W_\tau + \varepsilon_\tau \tag{3}$$

where C is a contrast matrix used to rescale the values such that they total zero, M is the matrix of attribute values (which is assumed to be fully available to the decision-maker over the course of the choice process), W_τ is a vector indicating which attribute is referred to in step τ and ε_τ is a random error vector. The weighting vector $W_\tau = [0..1..0]'$ with entry $W_{\tau,k} = 1$ if and only if attribute x_k is the attribute being attended to by the decision-maker at preference accumulation step τ. This means that a MDFT model estimates $n-1$ weights (where n is the number of attributes), where each weight, w_k corresponds to the proportion of time (preference accumulation steps) the decision-maker considers attribute k and $\sum_k w_k = 1$. These proportions, w_k, can thus be used to partition a unit interval, meaning that to simulate the process, draws from a standard uniform distribution $X \sim U(0,1)$ can be used to randomly select which attribute a decision-maker attends to at each time step (whilst simultaneously accounting for the importance of each attribute through the likelihood with which each attribute k is sampled, which depends on w_k). There is also a random error vector, ε_τ added on to allow for noise and variation in the probability values that MDFT predicts. $\varepsilon_\tau = [\varepsilon_{\tau_1}, .., \varepsilon_{\tau_n}]'$ with each ε_{τ_i} independently drawn from a normal distribution with mean zero and variance σ_ε^2, with higher values of σ_ε^2 expected for more complex decision-making tasks (Hotaling et al., 2010). Whilst different variances could be estimated for the different alternatives, we estimate a single variance for all alternatives in this paper.

2.1.2. Estimation of the probability with which alternatives are chosen in MDFT

The probability of choice of each alternative after τ preference accumulation timesteps can be calculated once the expected value and the covariance of the preference values (ξ_τ and Ω_τ) are known (Roe et al., 2001). Thus if we assume that the decision-maker concludes their decision-making process after τ preference updating steps, we can calculate the probability of each alternative being chosen. To calculate the expected value of the preference values, we first expand Equation 1, which results in:

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

where P_0 is the initial preference vector. The attribute weights w_k are stationary, therefore W_τ can be considered a stationary stochastic process. This results in V_τ also being a stationary stochastic process with mean $E[V_\tau]$ and a variance covariance matrix given by $Cov[V_\tau]$. We follow Roe et al. (2001) in letting ε_τ vary according to a normal distribution with mean zero and variance

σ_ε^2. Thus, if we have $\mu = E[V_\tau]$, it can be calculated as $\mu = C \cdot M \cdot w_m$, where w_m is a vector containing the probabilities of each of the attributes being considered. We also have $Cov[V_\tau] = \Phi = C \cdot M \cdot \Psi \cdot M' \cdot C' + \varepsilon$, where $\Psi = Cov[W_\tau]$ (C and M are matrices of constants). We can then calculate the expected value and the expected covariance of P_τ. With S being a constant, $E[P_\tau]$ reduces to:

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

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

We can also now calculate the covariance of the preference values[2]. By setting Z as the Kronecker product of S with itself and using Φ reshaped as a column matrix, $\overline{\Phi}$, the following reduction occurs (see Hancock et al. (2018)):

$$Cov[P_\tau] = \Omega_t = \sum_{k=0}^{\tau-1} \left[S^k \cdot \Phi \cdot S^{k'} \right] \tag{6a}$$

$$= \sum_{k=0}^{\tau-1} \left[Z^k \cdot \overline{\Phi} \right] \tag{6b}$$

$$= (I-Z)^{-1} (I-Z^\tau) \overline{\Phi} \tag{6c}$$

This simplified form for Ω_τ, together with ξ_τ, means that we can now calculate the probabilities of the alternatives after any number of preference accumulation steps τ. On the basis of the multivariate central limit theorem, P_τ converges to the multivariate normal distribution (Roe et al., 2001). This allows us to approximate the probability of the chosen alternative, which is the alternative with the greatest preference value at the conclusion of the deliberation process. Thus the probability of choosing alternative A from a set of n alternatives after τ steps is:

$$Prob[\max_{i \in n} P_\tau[i] = P_\tau[A]] = \int_{X>0} exp\left[-(X-\Gamma)'\Lambda^{-1}(X-\Gamma)/2\right] / (2\pi|\Lambda|^{0.5}) dX \tag{7}$$

with $X = [P_\tau[A] - P_\tau[B], \ldots, P_\tau[A] - P_\tau[n]]'$, $\Gamma = L\xi_\tau$, $\Lambda = L\Omega_\tau L'$ where

$$L = \begin{bmatrix} 1 & -1 & 0 & \ldots & \ldots & 0 \\ 1 & 0 & -1 & \ddots & & \vdots \\ 1 & \vdots & \ddots & \ddots & \ddots & \vdots \\ 1 & \vdots & & \ddots & -1 & 0 \\ 1 & 0 & \ldots & \ldots & 0 & -1 \end{bmatrix} \tag{8}$$

with L a matrix comprised of 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.

[2]Previously, a simplification was often made by assuming that $\tau \to \infty$ (Berkowitsch et al., 2014). However, this approach is not satisfying as it loses the dynamic nature of the accumulation of preferences, a key feature of the model. If the response time taken by the decision-maker is to be used, then we must further simplify $Cov[P_\tau]$ such that we avoid requiring a summation whilst simultaneously keeping a parameter for the number of preference accumulation steps.

2.2. Decision field theory with an internal threshold

2.2.1. Basic theory

For our decision field theory models using an internal threshold (which we henceforth call DFT-I), we consider the original specification of DFT by Busemeyer and Townsend (1992, 1993), which considers two alternatives, A_R and A_L. This differs from MDFT in a few small but distinct ways. Similarly to MDFT, the new preference values are a function of previous values and a number of other parameters. As DFT-I is only analytically solvable for two alternatives, we here only consider binary choice scenarios[3]. This means that we can consider just a single value, $P_\tau = Pref_\tau(A_R) - Pref_\tau(A_L)$. Alternative A_R is chosen when P_τ reaches some internal threshold θ, whereas A_L is chosen if P_τ reaches $-\theta$. The underlying key assumptions are still the same in that at each step, a single attribute k is attended to with probability w_k. This means that there is still a mean valence input, denoted δ, dependent on the valences v_R and v_L, which are weighted sums of the respective attributes for A_R and A_L (hence relative importance weights, w_k are still required for each attribute). The preference difference between the two alternatives then updates according to:

$$P_t = [1-(s+c)\cdot h]\cdot P(t-h) + [\delta \cdot h + \varepsilon(t)], \tag{9}$$

where s is a growth-decay parameter, which, similarly to ϕ_2 for MDFT, effects whether recently considered attributes or initially considered attributes have more impact. The parameter c is a goal-gradient parameter, which is used to explain the effect that 'avoidance-avoidance' decisions (choosing between two negatives) take longer than 'approach-approach' decisions (choosing between two positives) and is specified precisely in Equation 13. δ is the mean valence input, h is a time unit and $\varepsilon(t)$ is the error input, equivalently drawn from a normal distribution with mean zero, and variance $h\cdot\sigma^2$. Crucially, as a contrast to MDFT, the error variance is not assumed to be uncorrelated with the mean valence input. Instead, the variance of the error is defined directly as the input variance (Busemeyer and Townsend, 1993):

$$\sigma^2 = Var[V_R - V_L] = \sigma_R^2 + \sigma_L^2 - 2\cdot\sigma_{RL}, \tag{10}$$

with σ_R^2 the variance of the valence for choice alternative A_R, σ_L^2 the variance of the valence for choice alternative A_L, and σ_{RL} the covariance between the two.

We next detail the precise specification for the various parameters discussed above. First, to control for underlying biases towards an alternative, the initial value, P_0 is set to some value z, which is typically set as an increasing function of the mean valence input δ and the internal threshold θ. We follow Busemeyer and Townsend (1993)'s definition[4]:

[3] For the full equations for DFT using an internal threshold for more than two alternatives, readers should refer to the Appendix of Busemeyer and Townsend (1993). Further work should also consider searching for analytical solutions for more than two alternatives, as this is currently a major limitation for implementations of DFT-I.

[4] Note that in general a different function may be used, but we follow Busemeyer and Townsend (1993) in using the hyperbolic tangent function as it works well for predictions in Table 8 of their paper.

$$z = tanh(z^* \cdot \delta) \cdot \theta, \tag{11}$$

where z^* is a parameter to be estimated. Next, the internal threshold, θ, is defined to control speed-accuracy trade-offs (Busemeyer and Townsend, 1993). The threshold is therefore assumed to increase over time, and thus is set as an increasing function of some time parameter, τ_{DFTI}:

$$\theta = f(\tau_{DFTI}) \cdot \sigma, \tag{12}$$

where σ is the standard deviation of the error, as defined in Equation 10. The goal gradient parameter, c is then defined based on an approach gradient, a, and an avoidance gradient, b:

$$c = b \cdot (v_{R-} + v_{L-}) + a \cdot (v_{R+} + v_{L+}), \tag{13}$$

where the mean valence inputs v_R and v_L are split into a weighted sum of 'positive' attributes that are desirable, v_{R+} and v_{L+} (i.e. the average gain for alternative A_R and A_L respectively) and the weighted sum of 'negative' attributes that are undesirable, v_{R-} and v_{L-} (i.e. the average losses). Finally, the mean input valence is defined based on the approach gradient, avoidance gradient, internal threshold and expected gains and losses for each alternative:

$$\delta = (v_{R+} + v_{L+})(1 - a \cdot \theta) + (v_{R-} + v_{L-})(1 - b \cdot \theta), \tag{14}$$

2.2.2. Estimation of the probability of alternatives in DFT with an internal threshold

Whilst the probabilities with which each alternative is chosen under a DFT-I model can be calculated using Markov chain methods (Busemeyer and Diederich, 2002), we instead choose to solve the original integral given by Busemeyer and Townsend (1993), as this is simple to use for the choice scenarios described in this paper[5]. Busemeyer and Townsend (1993) demonstrate that by assuming a continuous time process ($h \to 0$), the probability of choosing alternative A_R over A_L is:

$$Pr(A_R, A_L) = \frac{S(z)}{S(\theta)}, \tag{15}$$

where the function $S(x)$ is the integral:

$$S(x) = \int_{-\theta}^{x} exp\left[\frac{(c+s) \cdot y^2 - 2 \cdot \delta \cdot y}{\sigma^2}\right] dy, \tag{16}$$

with the initial preference value, $P_0 = z$, θ the internal threshold, c the goal-gradient parameter, s the growth-decay parameter, δ the mean input valence and σ^2, the variance of this input, all as defined and specified in Section 2.2.1. To solve this integral, we use the substitution:

[5]Note that Markov chain methods are more practical to use for more complex scenarios (Bhattacharya and Waymire, 1990; Busemeyer and Townsend, 1993).

$$u = \frac{(c+s)\cdot y - \delta}{\sigma \cdot \sqrt{(c+s)}}. \tag{17}$$

In order to simplify the integral, we square both sides of Equation 17 and rearrange to get:

$$\frac{(c+s)\cdot y^2 - 2\cdot\delta\cdot y}{\sigma^2} = u^2 - \frac{\delta^2}{(c+s)\cdot\sigma^2}. \tag{18}$$

This, together with the derivative:

$$\frac{du}{dy} = \frac{\sqrt{c+s}}{\sigma}, \tag{19}$$

can be used to rearrange the integral in Equation 16 in terms of u and du:

$$S(x) = \int_{ll}^{ul} exp\left[u^2 - \frac{\delta^2}{(c+s)\cdot\sigma^2}\right]\cdot\frac{\sigma}{\sqrt{c+s}}\,du \tag{20a}$$

$$= exp\left[-\frac{\delta^2}{(c+s)\cdot\sigma^2}\right]\cdot\frac{\sigma}{\sqrt{c+s}}\cdot\int_{ll}^{ul} exp[u^2]\,du, \tag{20b}$$

where the limits of integration are $ll = \frac{-\theta\cdot(c+s)-\delta}{\sigma\cdot\sqrt{(c+s)}}$ and $ul = \frac{x\cdot(c+s)-\delta}{\sigma\cdot\sqrt{(c+s)}}$. Finally, we note that we require the imaginary error function ($erfi(x)$, Abramowitz and Stegun 1965), which is defined:

$$erfi(x) = \frac{2}{\sqrt{\pi}}\cdot\int_0^x exp[u^2]\,du. \tag{21}$$

Following some rearrangement, this results in a solution to Equation 15:

$$Pr(A_R, A_L) = \frac{S(z)}{S(\theta)} \tag{22a}$$

$$= \frac{erfi\left(\frac{z\cdot(c+s)-\delta}{\sigma\cdot\sqrt{(c+s)}}\right) - erfi\left(\frac{-\theta\cdot(c+s)-\delta}{\sigma\cdot\sqrt{(c+s)}}\right)}{erfi\left(\frac{\theta\cdot(c+s)-\delta}{\sigma\cdot\sqrt{(c+s)}}\right) - erfi\left(\frac{-\theta\cdot(c+s)-\delta}{\sigma\cdot\sqrt{(c+s)}}\right)}. \tag{22b}$$

2.3. The impact of time in DFT

Under a decision field theory model, the preferences of alternatives change as the decision-maker deliberates on the attributes of the alternatives. Consequently, the probability of each alternative being chosen changes depending on how long the decision-maker takes to make their choice. Under MDFT, the probabilities of each alternative being chosen are impacted by the number of preference accumulation steps, τ, (which we henceforth refer to as τ_{MDFT}), whilst under DFT with an evidence threshold, the probabilities depend on the threshold, θ, which in turn depends on the time parameter, τ_{DFTI}. For a typical simple setting in transport choice modelling, a decision-maker might need to complete choice tasks where there are two route alternatives, one which is cheaper and slower and another which is more expensive and faster. Figure 1 demonstrates how the probabilities that each alternative is chosen in such a choice (where the second alternative is 1 minute faster but 1 Swedish Krona more expensive) might change with an increase in decision time. Behaviourally, the assumption here is that a longer response time results in more comparisons of the alternatives, which results in more time for the preferred attributes to have an impact on the preferences.

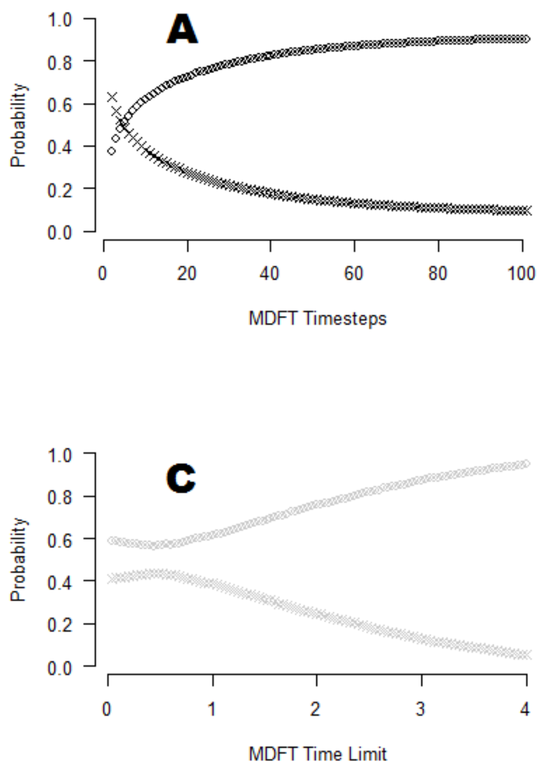

		Time (minutes)	Cost (SEK)
Alt 1	O	11	11
Alt 2	X	10	12
weight		0.4	0.6

FIGURE 1 : The probability of choosing the two different alternatives as the amount of time taken considering the choices increases.

Graph A shows a choice under MDFT and how the probability of choosing each alternative changes as the number of preference accumulation steps comparing the alternatives increases. Graph B shows the equivalent choice scenario under a DFT-I model, with the probability of alternatives changing with the time parameter (which impacts the value of the internal threshold depending on Equation 12). Graph C shows a MDFT model with an evidence threshold. This is created by using an MDFT model as described by Equations 1 - 3, but rather than calculating expected values and covariances, we run 10,000 simulations of the preferences evolving over time, concluding the choice when the difference in the preference values reaches some 'internal threshold', where the threshold is set as an increasing function of response time (as done for DFT-I). Under each version of DFT, the probability of choosing alternative 1 increases as the number of iterations of preference updating increase.

For MDFT, this figure clearly demonstrates that the higher the number of preference accumulation steps, the more deterministic the decision is[6]. Given additional comparisons, the fact that the decision-maker is more likely to consider the cost (which has a weight of 0.6 compared to that of travel time, which is 0.4), increases the probability of choosing alternative 1, which is cheaper. Additionally, DFT-I also becomes more deterministic, unsurprisingly, given that 'choice probability becomes more extreme as the threshold criterion increases' (Busemeyer and Townsend, 1993). Consequently, increasing the time parameter also increases the probability of alternative 1.

There are visible differences in the models with DFT-I producing S-shaped curves and MDFT producing a more logarithmic-shaped curve. However, results from simulations of MDFT with an internal threshold produces probabilities more similar to DFT-I, suggesting that the difference between MDFT and DFT-I is more likely due to the difference in when the choice is made (time or evidence threshold) rather than the difference in model structure (in terms of the different parameters the models have). Crucially, independently of the variation of DFT model used, an increased response time results in more time for evidence to accumulate, and consequently a more deterministic choice. Initially, the impact of the attribute weights is minimal, such that the chosen alternative depends on the initial preference matrix (which in Figure 1 is in favour of alternative 2) and which attribute is considered first[7]. As the number of preference accumulation steps increase, the higher weight for cost begins to have an impact, with the cheaper alternative gradually becoming more likely to be chosen. For MDFT, we set $\phi_2 = 0.2$ (for the purpose of this illustration), resulting in the preference values stabilising for a large number of preference accumulation steps (Berkowitsch et al., 2014), and hence the probabilities with which each alternative is chosen stabilise also.

We now look into incorporating choice response time into the models. In general, the functional form for the number of preference accumulation steps, $\tau_{MDFT_{np}}$, for individual n in a particular choice task p in a MDFT model could include multiple parameters:

$$\tau_{MDFT_{np}} = f(R_{np}) + g(z_n) + \varepsilon_{np} \tag{23}$$

where f is a function of the choice response time, R_{np}, in a particular choice task, $g(z_n)$ is a function of the characteristics of an individual, n, and ε_{np} is an error term, to be estimated[8]. For example, $g(z_n)$ could capture the fact that some individuals may not process information as quickly as others, hence they may require a larger number of comparisons before coming to the same conclusions. Whilst previous applications of decision field theory have simply estimated or fixed the number of preference accumulation steps, this paper considers choice tasks where choice response time is recorded. We subsequently consider three distinct methods for incorporating response time into the parameter for the number of preference accumulation steps for MDFT and the time parameter of DFT-I.

[6]Whilst this may not be the case under certain parameter values of MDFT and very specific attribute values, we choose typical attribute values and parameter estimates under MDFT in this scenario.

[7]Note that for the applications in this paper, decision-makers are presented with new alternatives in each trial, therefore there is little argument for the existence of initial preferences.

[8]In all of the datasets used in this paper, a single choice response time is observed for each choice task. Thus the deliberation and action components in making the choice cannot be disentangled. Thus this error term will incorporate 'reaction times' for the decision-makers to physically select their chosen alternative.

For our first set of models (T1), we wish to ensure that the number of preference accumulation steps is positive and that there is at least one step. Thus we define the number of preference accumulation steps as:

$$\tau_{MDFT_{np}} = 1 + e^{(t_0 + t_1 * R_{np})} \tag{24}$$

where R_{np} is the choice response time in a particular choice task as before, t_0 is an error term and t_1 is a parameter to be estimated. A positive value for t_1 thus indicates that the number of preference deliberation steps increases with increased response time whereas a negative estimate indicates a decrease in the number of steps for an increased response time[9].

Our second set of models (T2) additionally attempt to utilise differences in response times both across and within individuals. We do this by additionally considering a term t_2 for capturing the impact of a participant's mean response time. T2 therefore has the specification:

$$\tau_{MDFT_{np}} = 1 + e^{(t_0 + t_1 * RSD_{np} + t_2 * log(RM_n))} \tag{25}$$

where RM_n is the mean response time for individual n and RSD_{np} is the number of standard deviations the response time for task p is away from RM_n. Again, we use exponentials and add one to ensure that there is at least one step.

Finally, we also try running individual-specific models[10]. Here, for each choice task, we set the number of preference accumulation steps as:

$$\tau_{MDFT_{np}} = 1 + e^{(t_0 + t_1 * RSD_{np})} \tag{26}$$

where RSD_{np} is the number of standard deviations away the response time is from the individual's mean response time.

For DFT-I[11], we set the time parameter using the same functions as above but without adding 1 (thus $1 + \tau_{MDFT_{np}} = \tau_{MDFT_{np}}$), as there is no limit on how low the threshold can be (as long as it is positive for one alternative, and negative for the other).

2.4. Latent class models

As it is possible that there exist large differences between individuals, we also use latent class models (Kamakura and Russell, 1989) in this paper (as only one of our three datasets contains enough observations per individual to run individual models). These models allow for differences in sensitivities to be captured, with each different class capturing a different set of taste coefficients, with

[9] Note that in estimation, the number of preference accumulation steps does not have to be a discrete value, as reducing Equation 6a to 6c allows continuous values to be used for τ.

[10] Whilst many datasets in large-scale choice modelling applications only have a few observations per individual, one of our datasets has enough such that all observations included in a model can be from just one individual.

[11] Note that we do not run individual models using DFT-I, as the only dataset with enough observations per individual also has three alternatives, and we do not consider simulation of DFT probabilities in the empirical applications of this paper.

possibly even different models being used in the different classes (Hess et al., 2012). More often, the same model is used in the different classes, with S different copies of the model (where S is the number of classes) with a different set of parameter estimates β_S estimated for each class. Either way, we can then denote $P_{ni^*t}(\beta_s)$ as the probability of the chosen alternative, i^*, by individual n in choice task t under class s, where β_s is the set of parameters for class s. Allowing for S different classes results in the likelihood of the observed set of choices for individual n is:

$$L_n(\beta,\pi) = \sum_{s=1}^{S} \pi_{ns} \left(\prod_{t=1}^{T_n} P_{ni^*t}(\beta_S) \right) \tag{27}$$

where π_{ns} is the estimated share given to model s for participant n (summing to 1 for each n), and T_n is the set of choice tasks faced by the individual.

2.5. Multinomial Logit models

In the empirical applications of this paper, we also wish to compare DFT to static choice models (models in which preferences for alternatives do not change over time for a single set of alternative attributes), for which we use the multinomial logit (MNL) model (McFadden, 1974). This allows us to test the benefits generated by using a model that additionally attempts to capture the choice deliberation process.

Despite the fact that MNL models are static, there are various methods in which response time could be incorporated. Under the assumption of a decision-maker accumulating evidence in favour of alternative, a longer response time would result in a more deterministic choice. As the scale parameter within MNL (which is often normalised to 1) directly controls how deterministic a choice is, a scale parameter set as a function of response time would allow for the possibility of longer response times resulting in more deterministic choices. To illustrate how this would work, we first define our multinomial logit models. If we assume that an individual, n, has a utility, U^*, for alternative j, then:

$$U^*_{nj} = (\beta^*)' x_{nj} + \varepsilon^*_{nj} \tag{28}$$

where $(\beta^*)'$ is a set of parameters, x_{nj} is vector of observed variables relating to alternative j and ε^*_{nj} is the unobserved portion of utility. By assuming a type I extreme value distribution with variance $\sigma^2 \times (\pi^2/6)$, (McFadden, 1974) demonstrates that we can calculate the probability of alternative i being chosen as:

$$P_{ni} = \frac{e^{(\beta^*/\sigma)' x_{ni}}}{\sum_j e^{(\beta^*/\sigma)' x_{nj}}}, \tag{29}$$

with σ the scale parameter. Whilst this scale parameter cannot be identified alone from choice data, we can incorporate choice response time into our MNL models using the scale parameter, setting it as we did for the parameter for the number of preference accumulation steps in MDFT and the time parameter in DFT-I. Whilst the scale parameter does not need to be greater than one, it does need to be positive. Additionally we want the impact of t_1 and t_2 to have the same effect on MNL compared to DFT (higher estimates resulting in a more deterministic choice), so we set:

$$\sigma_{np} = \frac{1}{e^{(t_0 + t_1 * RSD_{np} + t_2 * log(RM_n))}}, \quad (30)$$

with RSD_{np} and RM_n defined as before.

3. EMPIRICAL APPLICATIONS

We will now demonstrate how the response time taken for choosing an alternative can be used in models applied to three very different datasets. All three datasets come from stated choice surveys, where participants consider several sets of alternatives amongst which they have to state their preference. Each choice task comprises of a set of hypothetical alternatives with differing attribute levels. All three questionnaires were completed on a computer, thus response times were recorded automatically when the participant selected an alternative.

3.1. Datasets

3.1.1. Route choice

The route choice dataset tested in this paper comes from a study on choice response time patterns in an online stated choice experiment (Börjesson and Fosgerau, 2015). In each choice task, respondents have two alternative car routes described by travel time and travel cost (in Swedish Krona). We discard choices with a recorded response time of 0 seconds and those with a response time of more than 60 seconds, as we assume that the respondent was either not attempting to respond to the choice seriously or was interrupted. Additionally, we omit choices made by respondents who have less than six (out of eight) choice tasks remaining after the above censor. This leaves us with 15,546 choice tasks completed by 2,358 respondents.

3.1.2. Conservation choice

The conservation dataset tested in this paper comes from a study exploring tree planting preferences in a stated choice survey (Mahieu et al., 2016). 146 participants completed 16 stated choice tasks where they were asked which of two conservation programmes they preferred. Each of the programmes are described by four attributes: country (Senegal or Peru), provision of online information (Yes/No), type of programme (restorative or preservative) and cost (2,5,10,15 EUR). Country and type of programme are both found to have an insignificant impact on the choice and are therefore omitted in this study. In all tasks, the participant also have a third ‘status quo’ alternative where they could choose not to invest in either of the presented programmes. We again exclude choices with a recorded response time of less than one second or choices with a time of more than a minute. This leaves us with 2,334 (out of 2,336) choice tasks.

3.1.3. Accommodation choice

Our accommodation choice dataset has a total of 32 participants each completing 45 choice tasks. It comes from Cohen et al. (2017)'s paper in which a version of MDFT is fitted with a fixed number of preference accumulation steps, $\tau_{MDFT} = 500$. In each choice task, decision-makers had three accommodation alternatives described by ease of transportation, size, condition and kitchen facilities (each on a scale of 1-5). Figure 2 gives an illustration of an example choice scenario. Trials were omitted if the decision-maker did not consider more than two out of twelve of the information panels as Cohen et al. (2017) deemed these to not be meaningful decisions[12]. Consequently, this leaves a total of 1,430 decisions.

FIGURE 2 : An example choice scenario from Cohen et al. (2017)

3.2. Basic models

We initially do not consider choice response time, simply using basic MDFT and DFT-I models to test the importance of the parameter for the number of preference accumulation steps and time parameter respectively. Given that the parameter for the number of preference accumulation steps is often fixed to a high value, these models effectively test how much flexibility is gained by freeing this parameter. We run five versions of MDFT on each dataset and five versions of DFT-I on just the route choice dataset, as this is the only dataset with only two alternatives.

In the first four versions, we fix the number of preference accumulation steps to 1, 10, 100 and 1,000, respectively for MDFT. For DFT-I, we fix the time parameter to 1, 2, 3, and 4. These values have no link with 'real time' in that the response time for each choice is not included, thus the different values simply test the parameter impact on the model. In the final version of each model for each dataset, we simply estimate the number of steps/time parameter to be some constant, $\tau_{MDFT} = \tau_{DFTI} = t_0$ (we are not yet including choice response time).

[12]Note that this was possible as the original experiment used eye-tracking equipment. See Cohen et al. (2017) for details.

To estimate the probability of alternatives under a decision field theory model with a time threshold, we require estimates for $n-1$ weight parameters (where n is the number of attributes) and estimates for four process parameters (ϕ_1 and ϕ_2, the sensitivity and memory parameters respectively, the constant for the number of preference accumulation steps, t_0, and the standard deviation of the error term, σ_ε).

We additionally consider new MDFT parameters equivalent to alternative specific constants in MNL (see Table 4 in Hancock et al. (2018)). These parameters are used to capture underlying preferences for alternatives and are particularly useful if there is a status-quo bias or an unfavourable alternative (such as not picking either conservation programme in our conservation choice dataset). For the work in this paper, we use additional attribute weights for time spent considering a specific alternative, j. For example, dummy attributes of $(0,0,1)$ could be used to represent 'additional time spent considering alternative 3 not captured by the attribute values'. These additional dummy attributes are omitted for the housing dataset, for which they are not found to be significant. We choose to use additional attribute weights rather than parameters in the initial preference matrix as we wish to test the importance and relative impact of the number of preference accumulation steps, which would be affected by an initial preference matrix.

For DFT-I, which is tested only on the route choice dataset, we also estimate two attribute weight parameters, with a third dummy attribute of $(0,-1)$ used to represent additional time spent considering the negatives of alternative 2. This is equivalent to $(1,0)$ but we use a negative as this results in all three attributes being negative, meaning that we do not need to estimate an approach gradient, a, and thus only need to estimate b, the avoidance gradient. However, we set $b=0$ and the growth-decay rate $s=0$ as this does not significantly impact model fit. Additionally, we set the bias term $z=0$ as, similarly to MDFT, we wish to avoid impacting the time parameter. This results in only two free parameters (the attribute weights) for DFT-I models with a fixed time parameter.

For all models, the free parameters are estimated by maximisation of the likelihood function of the observed choices. We use the R packages maxLik (Henningsen and Toomet, 2011) and cmcRcode (CMC, 2017) for estimation of the likelihood function. We use an RCPP package together with the Armadillo C++ linear algebra library for calculation of the matrices required for finding the probability of alternatives in the MDFT models (Eddelbuettel et al., 2011; Sanderson and Curtin, 2016). Finally, we use a initial parameter search algorithm based on Bierlaire et al. (2010)'s heuristic for non-linear global optimisation to try minimise the risk of our model finding poor local optima. The results of the basic DFT models are shown in Table 1.

For all three datasets, DFT achieves a better model fit if the parameter for the number of preference accumulation steps/time parameter is free. It is notable that MDFT models for both route and accommodation choice with a free parameter for the number of preference accumulation steps have vastly better fit than models with the number of steps fixed to 1,000, as often previously configured. Additionally, the estimate for the number of preference accumulation steps varies considerably across the datasets, further demonstrating the value of having a freely estimated parameter for the number of preference accumulation steps. DFT-I similarly produces results that are vastly worse if the time parameter is inappropriately fixed.

TABLE 1 : The number of free parameters (f.p), log-likelihoods (LL) and estimate/fixed value for the time parameter (τ_{DFTI}) or number of preference accumulation steps (τ_{MDFT}) for the basic DFT models

	DFT Evidence Threshold			**DFT Time Threshold**		
	Route			Route		
	f.p.	LL	τ_{DFTI}	f.p.	LL	τ_{MDFT}
restricted	2	-7,256.77	*1*	5	-7,441.19	*1*
restricted	2	-6,932.33	*2*	5	-6,904.70	*10*
restricted	2	-7,203.34	*3*	5	-6,915.92	*100*
restricted	2	-7,827.38	*4*	5	-6,915.92	*1,000*
τ free	3	**-6,930.75**	1.92	6	**-6,883.18**	5.83
	DFT Time Threshold			**DFT Time Threshold**		
	Accommodation			Conservation		
	f.p.	LL	τ_{MDFT}	f.p	LL	τ_{MDFT}
restricted	6	-1,429.20	*1*	6	-2,097.82	*1*
restricted	6	-1,329.75	*10*	6	-1,959.26	*10*
restricted	6	-1,320.69	*100*	6	-1,961.03	*100*
restricted	6	-1,330.40	*1,000*	6	-1,961.03	*1,000*
τ free	7	**-1,320.59**	290.54	7	**-1,959.17**	8.96

3.3. Models with choice response time

We next set the number of preference accumulation steps, τ_{MDFT} and the time paramater, τ_{DFTI} as functions of choice response time, using Equations 24 and 25 for models T1 and T2, respectively. For all three datasets, there is a statistically significant improvement in model fit by incorporating choice response time (see Table 2).

Whilst incorporating response time alone has little impact except in our conservation dataset (cf. models T1), where in general a faster decision is less deterministic, the real gain of using response time in our MDFT models is seen in models T2. This results in a slower mean response time indicating a more deterministic response ($t_2 > 0$, see Table 2), a finding that is consistent across all three datasets (for both MDFT and DFT-I). Additionally, an individual making a decision faster than their average response time is more deterministic ($t_1 < 0$) for that choice in some cases and less deterministic ($t_1 > 0$) in others. Consequently, by including both the decision maker's response time in comparison to others and to themselves, MDFT model fit is consistently improved across all three datasets. However, DFT-I has positive significant estimates for both t_1 and t_2, suggesting that a longer response time relative to either an individual's mean response time (cf. t_1) or a longer mean response time (cf. t_2) result in more deterministic choices.

These results are visualised through Figure 3, which shows the estimates for the number of preference accumulation steps and the time parameter for each choice, depending on the decision-

TABLE 2 : Results from including choice response time in DFT, with log-likelihoods and estimates and robust t-ratios for our time parameters.

Dataset/Model			T0	T1	T2
Route (Evidence Threshold)		Log-likelihood	-6,930.75	-6,927.20	-6,924.73
	t_0	estimate	0.65	0.60	0.27
		rob. t-ratio	24.81	18.71	1.56
	t_1	estimate	-	3.4E-03	0.02
		rob. t-ratio	-	2.82	1.99
	t_2	estimate	-	-	0.15
		rob. t-ratio	-	-	2.31
Route (Time Threshold)		Log-likelihood	-6,883.18	-6,882.37	-6,874.37
	t_0	estimate	1.58	1.50	0.16
		rob. t-ratio	12.04	9.85	0.36
	t_1	estimate	-	0.01	-0.02
		rob. t-ratio	-	1.35	-0.58
	t_2	estimate	-	-	0.54
		rob. t-ratio	-	-	3.22
Accommodation (Time Threshold)		Log-likelihood	-1,320.59	-1,320.38	-1,307.36
	t_0	estimate	5.67	6.25	4.40
		rob. t-ratio	2.55	4.40	4.85
	t_1	estimate	-	-0.01	-0.44
		rob. t-ratio	-	-0.45	-2.75
	t_2	estimate	-	-	0.72
		rob. t-ratio	-	-	3.53
Conservation (Time Threshold)		Log-likelihood	-1,959.20	-1,919.20	-1,917.96
	t_0	estimate	2.07	-4.73	-11.18
		rob. t-ratio	6.53	-1.50	-1.82
	t_1	estimate	-	2.08	1.05
		rob. t-ratio	-	2.42	1.62
	t_2	estimate	-	-	8.51
		rob. t-ratio	-	-	1.98

maker and their response times for each individual choice.

Whilst the estimated number of preference accumulation steps is extremely high (and implausible from a psychological perspective) for the conservation dataset, it is worth noting that the estimate for the memory parameter is $\phi_2 = 0.043$, meaning that the probabilities with which each alternative is chosen converges after a large number of preference accumulation steps. Consequently, if all decisions where $\tau_{MDFT} > 1000$ instead have 1000 steps, there is no change in model

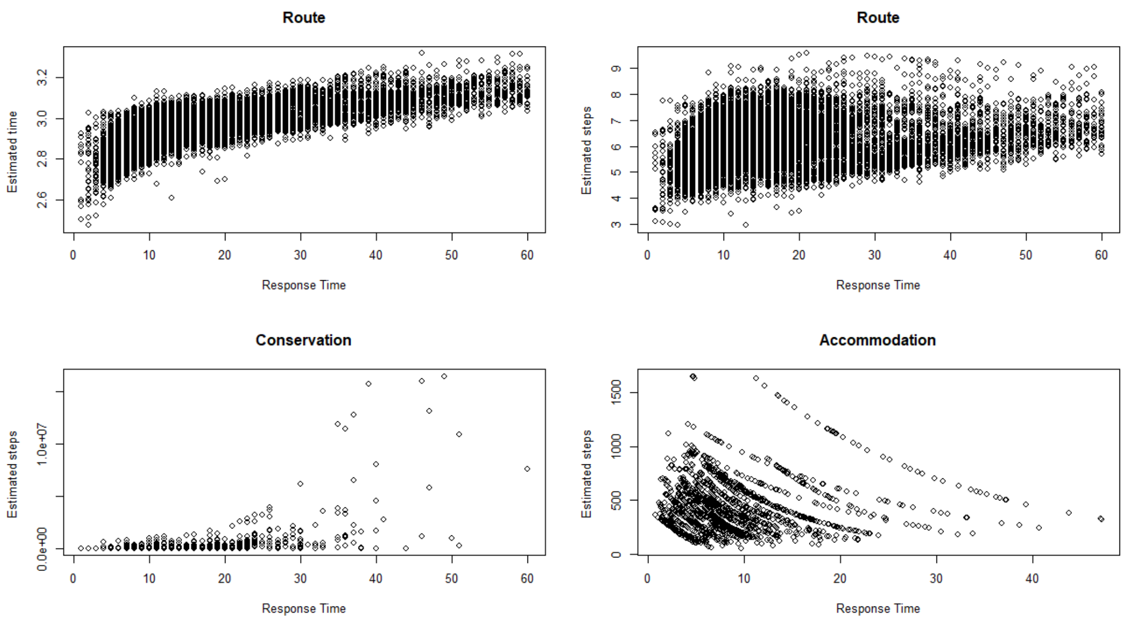

FIGURE 3 : The estimates for the number of preference accumulation steps or the time parameter depending on the response times

fit. This does however suggest that there is a large difference in how deterministic the different decision-makers are, as small changes for very fast responses have a significant impact.

3.4. Individual response time models

As we have up to 45 choice tasks per individual in our accommodation choice dataset, we can also run a MDFT model for each individual. This allows us to directly explore whether including response time aids model fit for explaining the choices made by each individual decision-maker. The results of these models are displayed in Figure 4. In line with the overall model (T2 in Table 2), we see that most estimates for t_1 are negative, resulting in decision-makers typically being more deterministic if they make a faster response. Whilst in many cases the estimate for t_1 is insignificant, in 10 out of 32 cases the log-likelihood for the model fit improves by more than 5% (with 8 of these 10 having a negative estimate for t_1 and the other 2 having a positive estimate). In 4 cases (all of which have negative estimates for t_1) the improvement is between 10 and 15%.

3.5. Latent class models

Given that some individuals might be more deterministic for faster responses and some might be more deterministic for slower responses, we also try latent class models to test whether separating the participants into two classes helps improve our models. Whilst we could use a function of parameters such that the class allocation is dependent on characteristics of an individual, in this case we are interested in whether there are different types of respondents based on their response time. Thus, in our first set of latent class models, we do not use response time and simply test the

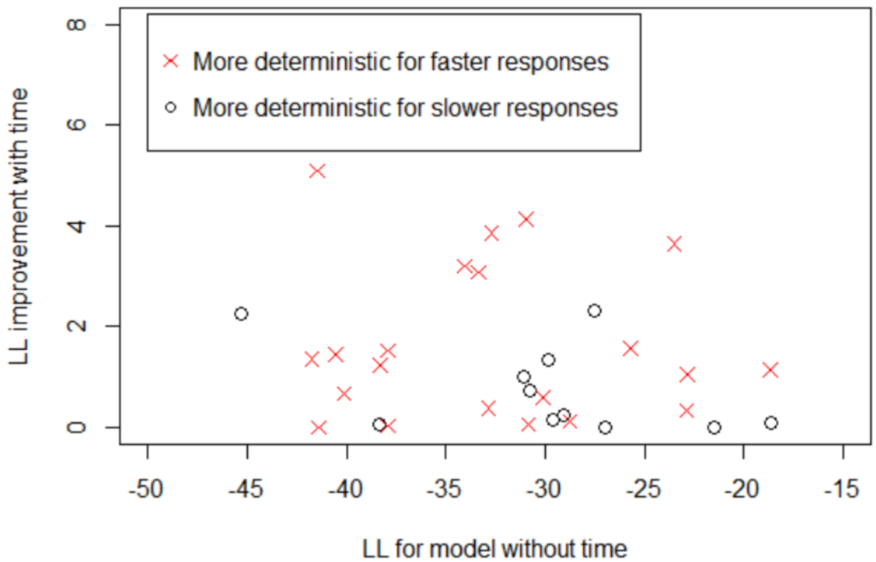

FIGURE 4 : The log-likelihoods for individual MDFT models compared to the improvement in fit by including response times

impact of an MDFT model with two different values for the number of preference accumulation steps and a DFT model with two different values for the time parameter (denoted τ_1 and τ_2 in Table 3 below). This means that the models only have two additional parameters, one for a second estimate for the number of preference accumulation steps, and one for the class share. The rest of the parameter estimates, including the attribute weights, are held constant across the two classes. The results are shown in Table 3.

TABLE 3 : Results for basic latent class DFT models for the three datasets

		classes	f.p.	LL	BIC	τ_1	class 1 share	τ_2
DFT-I	Route	1	3	-6930.75	13,890.45	1.92	100%	
		2	5	-6745.06	13,538.37	6.37E-04	9%	2.23
MDFT	Route	1	6	-6,883.18	13,824.27	5.83	100%	
		2	8	-6,742.23	13,561.67	1.00	20%	10.63
	Accommodation	1	7	-1,320.59	2,692.03	290.54	100%	
		2	9	-1,310.15	2,685.68	789.00	35%	4,704.68
	Conservation	1	7	-1,959.20	3,972.68	8.96	100%	
		2	9	-1,830.76	3,731.31	1.00	31%	19.28

Notably, the use of latent classes results in better fit for all three datasets, with a lower Bayesian Information Criteria (BIC) value for all datasets. For the route and conservation datasets, it appears that a subset of the decision-makers (9% for DFT-I or 20% for MDFT and 31% for MDFT, respectively for route and conservation) are much more random in their choices, with an estimate for the number of preference accumulation steps of 1 and a time parameter of 0 being found. This vastly improves model fit for both these datasets. This suggests that these individuals have very different sensitivities to the attributes in comparison to the majority of individuals, as the estimates suggest that random choices (where there is little/no time for evidence to accumulate) better represent their choices[13].

Given these results, it appears that our latent class models should also incorporate a separate set of attribute weights for each class (version 2 in Table 4). We can additionally also look at the impact of having two different classes for the DFT models whilst also including the response times taken by the decision-makers to make their choices. For these models (version 3 in Table 4), we use Equation 25 to set the number of preference accumulation steps/the time parameter, with different values for t_0, t_1 and t_2 in the two different classes. The results of latent class models 2 and 3 are given in Table 4.

Crucially, there is a vast improvement for all models in comparison to the first set of latent class models, demonstrating that there is a large amount of taste heterogeneity across all three datasets going beyond just the τ_1/τ_2 split. For the route choice dataset, it appears that some individuals are more sensitive to time (wt_1), whilst others are more sensitive to cost (wt_2). The key differences are the cost for the accommodation (wt_2) and the bias against picking neither alternative (wt_4) for conservation datasets respectively. Whilst the log-likelihoods for latent class models with response time are similar to those of latent class models without response time, there is still a significant improvement for models across all three datasets (though not in terms of BIC for route and conservation). The parameter estimates for the response time coefficients in both classes are shown in Table 4 also, with many positive and negative coefficients. Figure 5 shows this effect clearly, with, for example, the time parameter increasing with response time for class 1 for DFT-I, but the time parameter decreasing with response time for class 2. For the conservation dataset it appears that there is one group (class 2) where response time does not influence randomness, with the other (class 1) producing less deterministic choices with increasing response time. Both classes produce less deterministic choices with increasing response time for the accommodation dataset.

3.6. Reappraisal of results using multinomial logit models

As our results demonstrate that DFT time parameters can be successfully parameterised as a function of response time, this leads us to consider whether alternative models can similarly include response time information. The most obvious example of a model to test is the multinomial logit (MNL) model (McFadden, 1974). Whilst this is a static model, the response time can be incorporated easily through the use of Equation 30.

We now test five different MNL models for each dataset, with each one being equivalent to a DFT models tested already in this paper.

1. A basic MNL model without response time ($\sigma = 1$ in Equation 30).
2. A MNL model with response time captured as defined in Equation 30.
3. A MNL model with 2 latent classes, with only the scale parameter different and without response time

[13] Note that the wide range in the number of preference accumulation steps for the model incorporating response time for the conservation dataset also implies this.

TABLE 4 : The results of versions 2 and 3 of latent class models for both variations of DFT

Model		DFT Evidence Threshold		DFT Time Threshold					
Dataset		Route				Accommodation		Conservation	
Version		2	3	2	3	2	3	2	3
response time		no	yes	no	yes	no	yes	no	yes
free parameters		7	11	13	17	15	19	15	19
Log-likelihood		-6,017.46	-6,007.09	-5,983.38	-5,976.97	-1,255.22	-1,238.64	-1,745.76	-1,734.77
BIC		-12,102.49	-12,120.34	-12,092.22	-12,118.02	-2,619.42	-2,615.31	-3,521.53	-3,616.89
	class 1 share	0.62	0.62	0.64	0.63	0.59	0.69	0.75	0.73
Class 1 Par. Ests.	t_0	0.82	1.07	1.94	2.45	7.71	5.66	2.43	1.87
	t_1	-	-0.02	-	0.03	-	-0.34	-	-0.38
	t_2	-	-0.02	-	-0.19	-	0.93	-	0.36
	wt_1	0.25	0.25	0.26	0.25	0.11	0.13	0.04	0.03
	wt_2	0.62	0.63	0.64	0.65	0.47	0.42	0.10	0.08
	wt_3	0.13	0.12	0.10	0.10	0.14	0.17	0.11	0.13
	wt_4	-	-	-	-	0.28	0.29	0.75	0.76
	class 2 share	0.38	0.38	0.36	0.37	0.41	0.31	0.25	0.27
Class 2 Par. Ests.	t_0	1.10	0.66	7.28	-2.43	5.64	6.90	1.03	-1.21
	t_1	-	3.00E-04	-	-0.02	-	-0.82	-	-0.01
	t_2	-	0.03	-	2.11	-	-0.20	-	0.93
	wt_1	0.51	0.50	0.49	0.48	0.20	0.20	0.00	0.00
	wt_2	0.29	0.28	0.26	0.26	0.38	0.33	0.22	0.20
	wt_3	0.20	0.22	0.25	0.26	0.30	0.35	0.39	0.37
	wt_4	-	-	-	-	0.12	0.12	0.39	0.43

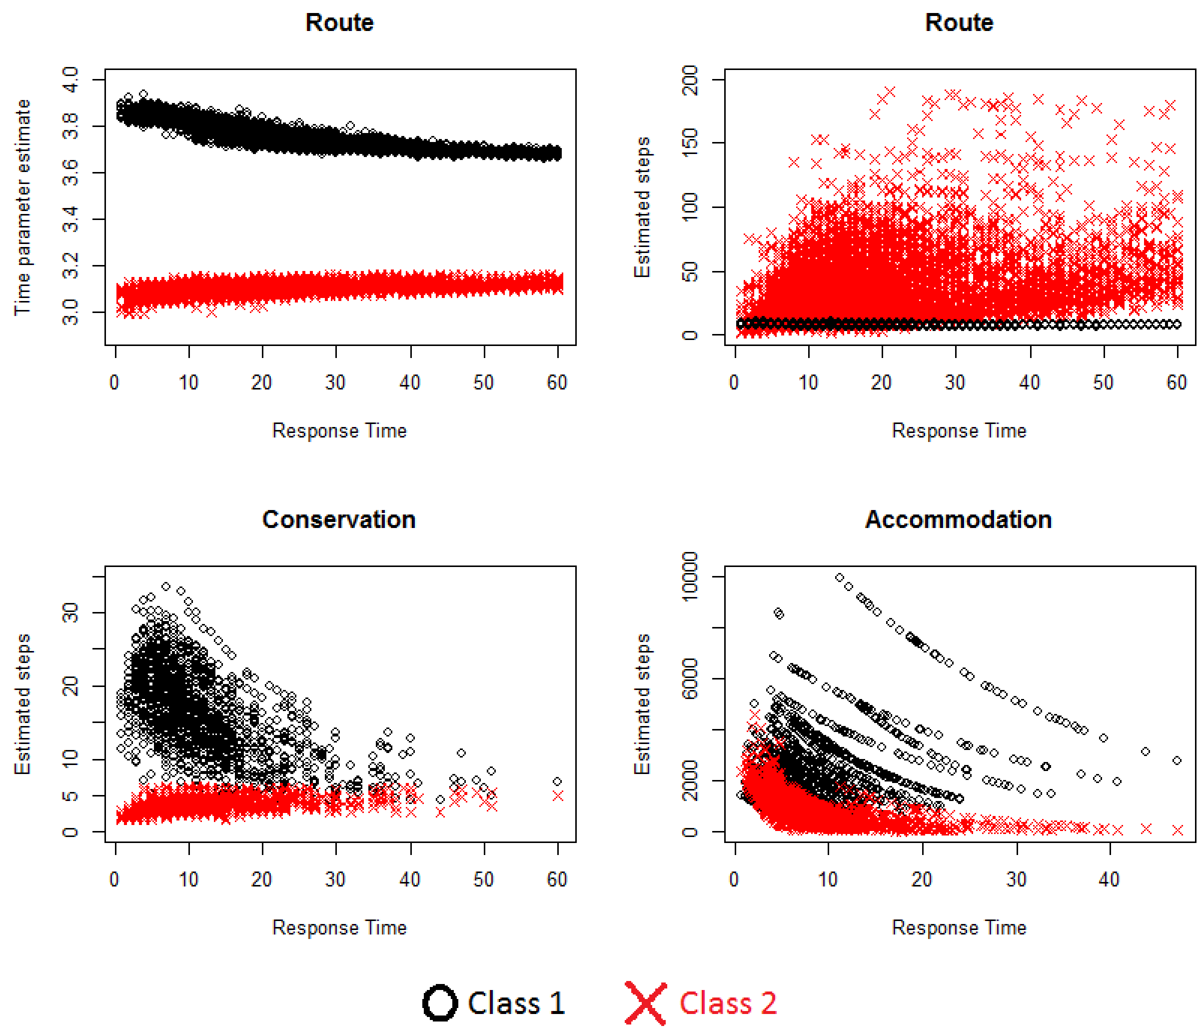

FIGURE 5 : The estimates for the number of preference accumulation steps and the time parameter depending on the response times for the latent class models.

4. A MNL model with 2 latent classes, with different scale and weight parameters, but without response time

5. A MNL model with 2 latent classes, each with a complete set of separate parameters, including separate response time parameters t_1 and t_2.

The results from our MNL models are shown in Table 5.

Whilst MNL has vastly worse fit for the route choice dataset, the MNL models have very similar log-likelihoods to their corresponding DFT models[14] for the conservation and accommodation

[14] Note that DFT-v2 is the same model as DFT-T2 in Table 2.

TABLE 5 : Results from the MNL models, with log-likelihoods also shown for equivalent DFT models

model	version	response time	weights difference	classes	Route		Conservation		Accommodation	
					f.p	LL	f.p	LL	f.p	LL
MNL	1	no	no	1	3	-8,769.75	4	-1,958.68	4	-1,329.09
	2	yes	no	1	5	-8,767.11	6	-1,934.59	6	-1,313.95
	3	no	no	2	5	-7,395.42	6	-1,841.38	6	-1,316.49
	4	no	yes	2	7	-7,105.15	9	-1,749.63	9	-1,265.66
	5	yes	yes	2	12	-7,089.40	14	-1,739.94	14	-1,257.27
DFT Time Threshold	1	no	no	1	6	-6,883.18	7	-1,959.20	7	-1,320.59
	2	yes	no	1	8	-6,874.37	9	-1,917.96	9	-1,307.36
	3	no	no	2	8	-6,742.23	9	-1,830.76	9	-1,310.15
	4	no	yes	2	13	-5,983.38	15	-1,745.76	15	-1,255.22
	5	yes	yes	2	17	-5,976.97	19	-1,734.77	19	-1,238.64
DFT Evidence Threshold	1	no	no	1	3	-6,930.75	-	-	-	-
	2	yes	no	1	5	-6,924.73	-	-	-	-
	3	no	no	2	5	-6,745.06	-	-	-	-
	4	no	yes	2	7	-5,986.41	-	-	-	-
	5	yes	yes	2	11	-5,976.97	-	-	-	-

choice datasets. In particular, the impact of including response time is similar for MNL compared to DFT.

Consequently, it appears that adjustments in the specification for the scale parameter in MNL have very similar impacts to the equivalent adjustments for the time parameters in DFT. Additionally, the time parameter estimates for the single class version of MNL and MDFT for all three choice sets appear to be similar in that they always have the same sign (see Table 6). The t_1 and t_2 estimates for MNL are approximately half that of those for the corresponding MDFT models. The impacts of this are particularly clear when considering Figure 6, in which the distributions for the scale parameter estimates for the route and accommodation datasets very closely resemble that of the estimated number of preference accumulation steps for MDFT in Figure 3.

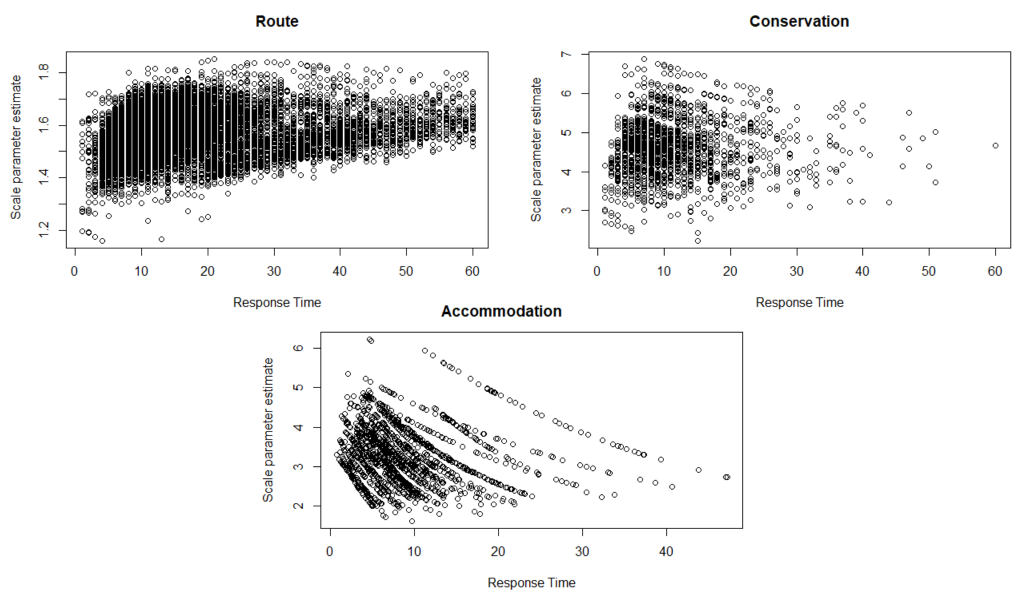

FIGURE 6 : The scale parameter estimates in the MNL models depending on the response times.

This implies that the parameter for the number of preference accumulation steps in a MDFT model can perform the same function as the scale parameter in a MNL model. The consequences of this are that if we only have the choices and response times of a decision-maker, then MDFT may not make any further gain than a far simpler MNL model (with the exception being the single class version for the conservation dataset, although this gain is lost once two classes are considered). This suggests that attempting to capture the deliberation process, in this case, does not add anything to the performance of the model.

Whilst the same cannot be said of MNL and DFT-I, which have somewhat different estimates (see Table 6), the results for DFT-I suggest that response time cannot be directly linked the time parameter. This is because the latent class DFT-I model has a significant negative estimate for t_2 for one of the classes, suggesting that individuals in this class are more deterministic if they have a faster mean response time.

TABLE 6 : Time parameter estimates from DFT and MNL models

Dataset			MNL LC1	MDFT LC1	DFT-I LC1	MNL LC2		MDFT LC2		DFT-I LC2	
Route		Log-likelihood	-8,767.11	-6,874.37	-6,924.73	-7,089.40		-5,976.97		-6,007.09	
	t_0	estimate	0.00	0.16	0.27	0.00	-3.95	2.45	-2.43	1.07	0.66
		rob. t-ratio	**fixed**	0.36	1.56	**fixed**	-4.53	4.39	-1.18	10.77	4.45
	t_1	estimate	-0.01	-0.02	0.02	-0.13	-0.06	0.03	-0.02	-0.02	0.00
		rob. t-ratio	-0.41	-0.58	1.99	-2.56	-1.12	0.45	-0.19	-0.74	0.01
	t_2	estimate	0.17	0.54	0.15	0.12	0.56	-0.19	2.11	-0.02	0.03
		rob. t-ratio	2.17	3.22	2.31	0.64	2.13	-0.92	2.43	-2.67	3.18
Accommodation		Log-likelihood	-1,313.95	-1,307.35	-	-1,257.27		-1,238.64		-	
	t_0	estimate	0.00	4.40	-	0.00	0.16	5.66	6.90	-	-
		rob. t-ratio	**fixed**	4.85	-	**fixed**	0.07	7.64	1.22	-	-
	t_1	estimate	-0.28	-0.44	-	-0.17	-0.23	-0.34	-0.82	-	-
		rob. t-ratio	-4.58	-2.75	-	-1.68	-1.58	-2.10	-3.30	-	-
	t_2	estimate	0.36	0.72	-	0.04	0.34	0.93	-0.20	-	-
		rob. t-ratio	3.11	3.53	-	0.04	1.02	3.63	-0.25	-	-
Conservation		Log-likelihood	-1,934.59	-1,917.96	-	-1,739.94		-1,734.77		-	
	t_0	estimate	0.00	-11.18	-	0.00	0.38	1.87	-1.21	-	-
		rob. t-ratio	**fixed**	-1.82	-	**fixed**	0.33	2.04	-0.48	-	-
	t_1	estimate	-0.08	1.05	-	-0.22	-0.15	-0.38	-0.01	-	-
		rob. t-ratio	-2.09	1.62	-	-0.74	-1.31	-3.09	-0.03	-	-
	t_2	estimate	0.55	8.51	-	-0.05	0.23	0.36	0.93	-	-
		rob. t-ratio	3.06	1.98	-	-0.03	1.16	0.88	0.65	-	-

4. CONCLUSIONS

The work in this paper was motivated by the recent improvements (in Hancock et al. (2018)) in the computational mechanisms underlying decision field theory models with an external threshold (MDFT). With it now being easily possible to incorporate decision response time into a MDFT model, this paper considers the impact of this on three datasets, as well as testing the inclusion of response time in models based on the original specification of DFT (labelled DFT-I in this paper) by Busemeyer and Townsend (1993), in which individuals conclude a decision when they reach an internal (evidence) threshold.

In favour of the notion of preference accumulation, we find that MDFT models estimate a larger number of preference accumulation steps for decision-makers who have a longer mean response time relative to other decision-makers. However, contradicting the notion of evidence accumulation, we find some negative estimates for the impact of increased response time meaning that more deterministic decisions are made in choices where a decision-maker responds more quickly than their mean response time across all of their choice tasks. Additionally, the impact of the average response time for an individual diminishes with latent class MDFT models, with two separate estimates for the number of preference accumulation steps controlling for the effect of a longer mean response time. With MDFT parameters for the number of preference accumulation steps behaving very similarly to MNL scale parameters, it appears that the step parameter cannot be directly linked to response time. This means that a MDFT model with response time may not in fact capture a choice deliberation process. This is perhaps unsurprising given that MDFT models could be considered inappropriate for the choice tasks in this paper, for which no time limit is imposed (contrary to the key assumption of MDFT, that a decision is made upon reaching a specific time limit). However, we do not computationally impose a strict time limit on our MDFT models (with the number of preference accumulation steps being adjusted according to the response times) and results from our DFT models with internal thresholds are also unfavourable for the notion of preference accumulation. Our latent class DFT-I model results in both positive and negative estimates for the effect of a participant's mean response time on how deterministic the choice is. This implies that the core assumption of DFT, that choice probabilities become more extreme with a higher threshold (evidence or time), cannot be predicted when considering response times. However, the impact of response time on the parameter for the number of preference accumulation steps and the time parameter might look very different if a DFT model could control for choice certainty in a more direct way. Thus, further work could consider additionally linking choice certainty as well as response times to the time parameters.

Whilst our results suggest that the deliberation process may not be truly captured by an DFT model, this work does provide evidence that incorporating response times into choice models can improve model fit. The fact that response time can be directly included in calculating the probability of alternatives may have some impact in stated preference studies but is far more likely to have an impact in work involving revealed preference data, where decisions such as what to order in a restaurant, who to vote for and which route to choose at a roadway junction are more likely to be impacted by time pressure due to the nature of the choices. Additionally, previous work has demonstrated that the linear ballistic accumulator model (Brown and Heathcote, 2008), can be adjusted such that changing information can be incorporated into the model (Holmes et al., 2016).

Similarly, it is possible that as well as incorporating response time, a decision field theory model could incorporate changing attributes, such as in a dynamic price setting (e.g. auctions, flight booking websites) or travel times for different routes. This could also prove useful for studying how commuters change their route choice when forced to do so due to a change to their original schedule, such as a delayed train. In particular, dynamic models such as DFT may prove useful for econometric forecasting, particularly if we do not have much information on the decision-maker but there is some indication on how long they might take to make the decision. However, DFT and other accumulation models may only have advantages over basic models such as multinomial logit by considering such data complexities, without which, results from this paper suggest that MNL can perform just as well for multi-alternative, multi-attribute choice. Thus, a DFT model for dynamic data may also require information such as choice certainty if it is to capture the deliberation process accurately.

Further work could also consider latent constructs, where latent variables are used to predict both response times and the number of preference accumulation steps. Certain individuals may process information at different rates, meaning that there is likely to be variation in the estimated number of iterations of preference updating per second across individuals. It may not be that individuals who spend longer considering a decision are necessarily considering the alternative in more detail, as implied by the structure we impose on the models in this application. Using random parameters to capture this difference across individuals therefore may have much more explanatory power than when point estimates for parameters are used.

Whilst these results give various implications and provide many directions for future work, it is clear that, where possible, analysts should record response times in choice decisions, as it is likely that both dynamic models such as DFT and simpler models such as multinomial logit are able to use this information to better predict the choices made. This may lead to more accurate estimations and forecasts in many situations.

ACKNOWLEDGEMENTS

The authors acknowledge the financial support by the European Research Council through the consolidator grant 615596-DECISIONS. We thank Andrew Cohen, Romain Crastes dit Sourd and Maria Börjesson for allowing us to use their datasets for this research.

REFERENCES

Abramowitz, M. and Stegun, I. A. (1965). *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, volume 55. Courier Corporation.

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.

Bhattacharya, R. N. and Waymire, E. C. (1990). *Stochastic Processes with Applications*, volume 61. SIAM.

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

Börjesson, M. and Fosgerau, M. (2015). Response time patterns in a stated choice experiment. *Journal of choice modelling*, 14:48–58.

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

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.

CMC (2017). CMC choice modelling code for R, Choice Modelling Centre, University of Leeds, UK.

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. and Zachary, S. (1978). Improved multiple choice models. *Determinants of travel choice*, 335:357.

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.

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.

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

Hess, S., Stathopoulos, A., and Daly, A. (2012). Allowing for heterogeneous decision rules in discrete choice models: an approach and four case studies. *Transportation*, 39(3):565–591.

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.

Kamakura, W. A. and Russell, G. J. (1989). A probabilistic choice model for market segmentation and elasticity structure. *Journal of marketing research*, pages 379–390.

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

Mahieu, P.-A., Crastes, R., Louviere, J., Zawojska, E., et al. (2016). Rewarding truthful-telling in stated preference studies. Technical report.

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

McFadden, D. (1978). Modeling the choice of residential location. *Transportation Research Record*, (673).

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.

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.

Raab, M. and Johnson, J. G. (2004). Individual differences of action orientation for risk taking in sports. *Research quarterly for exercise and sport*, 75(3):326–336.

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. and Black, I. R. (2006). Means matter, but variance matter too: Decomposing response latency influences on variance heterogeneity in stated preference experiments. *Marketing Letters*, 17(4):295–310.

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

Schall, J. D. (2003). Neural correlates of decision processes: neural and mental chronometry. *Current opinion in neurobiology*, 13(2):182–186.

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.

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.

Tsetsos, K., Usher, M., and Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological Review*, 117(4):1275.

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

Williams, H. C. (1977). On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and planning A*, 9(3):285–344.