New models for dynamic choice contexts: further steps towards bridging choice modelling with mathematical psychology

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@leeds.ac.uk

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

ABSTRACT

Modelling decisions in dynamic choice contexts, where the decision maker has to make a single or a series of choices in a changing scenario (e.g. driving manoeuvres, route choice, investment decisions, etc.) is much more challenging than in static scenarios. These decisions have typically been based on models that split time into discrete intervals, estimate utilities for each time point, then use lagged variables to account for the dynamic nature of the scenario. This is despite the simultaneous development of 'sequential sampling' (accumulator) models, that attempt to capture the accumulation of preference over time. The application of such models have however been limited to scenarios with static attributes. Decision field theory (DFT), one such model which assumes that the preferences for different alternatives update stochastically over time, has recently made the transition into mainstream choice modelling. Whilst this model has been effective at modelling, for example, route choice, it has not yet been tested on choice scenarios where the attributes of alternatives update over time. In this paper, we discuss the relative benefits of using an accumulator model for dynamic choice contexts, as well as developing and operationalising both DFT and models based on quantum probability frameworks for dynamically changing attributes in the context of driving behaviour problems. We demonstrate that analytical and simulated versions of DFT as well as quantum models provide a better account of behaviour than models based on more traditional methods. Whilst the models developed in this paper are preliminary in nature, we demonstrate that there is immense potential in the development of such models for dynamic choice scenarios.

1. INTRODUCTION

The vast majority of applications in choice modelling look at individual choices in isolation, where a decision maker chooses amongst a finite set of alternatives from a static choice set, i.e. one where the alternatives and their attributes do not change over time. A relatively small departure from that context comes in work looking at the relationship between individual choices over time, largely modelled through state dependence (Kitamura and Bunch, 1990; Bjørner and Leth-Petersen, 2005).

A rather different set of circumstances applies in the context where a choice is made in a scenario where the composition of the choice set and/or the attributes describing the alternatives changes during the time interval when the choice takes place. Work with dynamic discrete choice models (DDCM) looks at situations where the set of alternatives or their attribute levels can change (Cirillo and Xu, 2011). Dynamic problems are often formulated as Markov decision processes (Rust, 1994) under which the next system state is dependent on a Markov transition probability. This has typically been applied to consumer problems such as the purchase of digital camcorders (Rysman et al., 2009) or the purchase of new automobiles (Schiraldi, 2011). To a large extent, much of this work is based on the idea of splitting the choice process into a number of sequential choices. This is however quite different from the situation where only a single choice is made, but where the characteristics of the choice set evolve during the process of the decision maker making that choice. In addition, choice contexts in some cases can change much more rapidly than for the longer term scenarios discussed above. This is especially the case in transport. In particular, the decision to merge into a different lane is dependent on the size of the gap in the target lane, which changes rapidly as the vehicles in front and behind change speed. Previous efforts at modelling merging decisions and gap acceptance problems have also included logit models (Polus et al., 2005; Rossi et al., 2013; Marczak et al., 2013), but are largely static in nature. A different application of dynamic models comes in route choice modelling, with Fosgerau et al. (2013) demonstrating that a dynamic model for sequential link choices can be used effectively to capture route choice in a network and Mai et al. (2015) further demonstrating that correlation across error terms can be included for such problems. This work however also just replaces a single choice of an overall route by a recursive set of individual decisions, where in addition, the attributes of the network are stable over time.

In DDCMs, the continuous time horizon is typically split into separate individual choice situations, where the decision to take a given action (possibly out of a set of actions) or not, is modelled at each point in time (Aguirregabiria and Mira, 2010). In contrast to these 'static' models, many 'dynamic' choice models have been developed in mathematical psychology, where the preference of an alternative updates over time as the decision-maker deliberates on the set of alternatives (Busemeyer and Townsend, 1992; Ratcliff and Rouder, 2000; Usher and McClelland, 2001; Bhatia, 2013). In these 'accumulator' models, however, the dynamic process can thus be estimated as a single continuous process, where the preferences change over time and where an action is taken when a certain preference threshold is reached.

These 'process' or 'sequential sampling' models aim to better understand the choice by attempting to additionally consider the decision-maker's deliberation process that occurs whilst they consider their alternatives. This results in models that can accurately predict response times (Brown and Heathcote, 2008; Diederich and Busemeyer, 2003) as well as choice outcomes. Typically, a key focus of accumulator models has been on how well they predict various contextual effects (Tsetsos et al., 2010; Trueblood et al., 2014; Noguchi and Stewart, 2018), with the transition of such models into 'mainstream' choice modelling limited due to computational complexities of such models (Otter et al., 2008). However, recent work has demonstrated that decision field theory (DFT), a stochastic, dynamic model for understanding choices, and the multiattribute linear ballistic accumulator model (MLBA), a model where the preference for an alternative 'drifts' linearly, can both be used effectively to model travel and consumer choice problems (see (Hancock et al., 2019a)).

The focus in the application of these models has however been on the dynamic updating of preferences (i.e. internal accumulation of evidence) rather than looking at dynamically changing attributes of the alternatives/choice sets. Whilst these models should theoretically be able to easily incorporate updating of attributes, applications thus far have been limited to basic scenarios such as changes in dot direction in a dot perception task (Holmes et al., 2016). While the use of DDCMs may be appropriate for longer term choices, where a decision maker over time returns to a specific choice context and *thinks about it again*, it seems in our view a less natural solution for situations where a decision maker focusses on a specific choice for a relatively short period of time but where the attributes of that choice are not static. Numerous examples exist. A driver aiming to merge onto a motorway may have to assess a number of gaps, where the attributes of the gap (e.g. gap-size) are changing over time, before finding a suitable gap to merge into. Alternatively, a traveller may be tracking hotel prices over multiple days, where the prices are changing dynamically, before making a booking. Outside of transport, stock market investment decisions are evaluated over time, while share prices change continuously. This paper aims to bring together dynamic datasets and dynamic models, testing whether models and ideas developed in mathematical psychology can be used to better model decision probabilities in dynamic environments, where we particularly focus on decision field theory (DFT).

Under DFT, a decision-maker is assumed to consider different attributes for the different alternatives over the course of the deliberation process. In the standard implementation of DFT, this *sampling* takes information from a static set, i.e. the values for the different attributes and alternatives are kept constant during the deliberation. This is appropriate for example in the context where a decision maker looks at a shelf of products in a supermarket and keeps alternating between price and some quality attribute for comparing the alternatives, where these attributes remain static. It is however less appropriate for truly dynamic setting as it in essence assumes that a decision maker takes a snapshot of the information describing the alternatives at the outset, and then keeps comparing the alternatives internally using that information. It seems much more likely that there is in addition some *external* updating of the information, notwithstanding the possibility that the *internal* deliberation may operate at a higher frequency (i.e. thinking repeatedly about a specific gap size and speed differential before updating the information on this after it changes).

A simple example where a decision-maker considers the choice between two alternative routes helps illustrate the concept. The first alternative has a cost of £5 and travel time of 30 minutes and the second costs £10 and takes 20 minutes. This might result in the preferences for alternatives updating over time in line with the values given in Figure 1. In each timestep, one attribute is considered, and the alternative which performs better for that attribute gains in preference value. For example, in the 1st, 3rd and 4th deliberation (preference updating) timesteps, the decision-maker considers cost, and consequently the cheaper alternative quickly becomes favourable, even though the faster alternative was originally preferred. If the first alternative to reach a preference value of 5 is chosen, the decision-maker would then choose the cheaper alternative. If, however, the decision-maker did not come to a conclusion, they might reach a point where the attributes change. The example in Figure 1 demonstrates what might happen if tickets were reduced in price for the faster alternative (after 5 deliberation timesteps). This results in the faster alternative quickly becoming the preferred option.

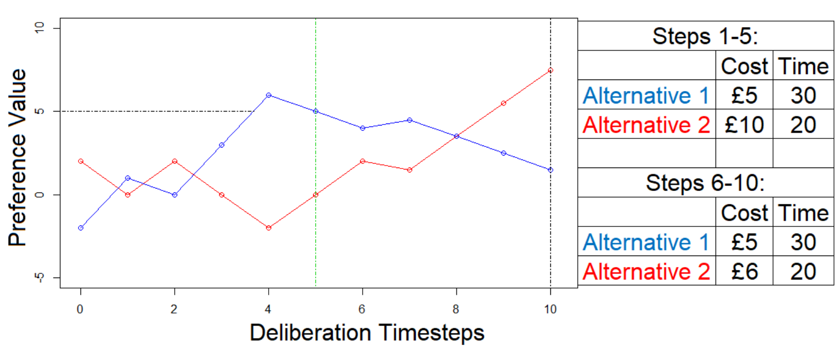

FIGURE 1 : An example deliberation process when attribute values change, under a DFT model

Thus far, the only previous attempt to fit such a model to dynamic data (as far as the authors are aware) was made by Holmes et al. (2016). They used an adapted version of the linear ballistic accumulator model (where the preferences for alternatives grows linearly until the preference for one alternative reaches a threshold) to model decisions about moving dot perception tasks. However, many real-world dynamic choices are more complex, thus ready-made accumulator models with analytical solutions for calculating the probability of alternatives being chosen may not be possible. The aim of this paper is to introduce an effective dynamic choice model to predict choices under dynamic environments. We consider the use of decision field theory for understanding a driver's decision as to whether or not to merge onto a motorway. As well as testing a version with attributes staying constant for a set number of deliberation timesteps before updating (as depicted in Figure 1 and in line with the earlier point about internal deliberation operating at a higher frequency than external updating), we also compare a simulated approach where the attributes change with every deliberation timestep. We also compare these to more traditional 'static' models that use exponential and hyperbolic discounting functions to account for 'remembered' preferences evaluated at previous attribute values.

An alternative approach is to add a dynamic element to quantum probability models. These models are based on the operationalisation of quantum logic to create a different probability system to generate the probabilities with which choices (or actions) are made. Models with such a structure have recently made a significant impact in cognitive psychology (c.f. Busemeyer and Bruza 2012 for a review) and consequently have been demonstrated to show significant promise in choice modelling (Hancock et al., 2019c). The key assumption under quantum logic is that some pairs of decisions are 'incompatible', meaning that the probabilities of choosing different alternatives or actions in one decision are impacted by the choice made in the other. This results in the adoption of quantum logic allowing for convenient and elegant solutions to, for example, ordering effects (Trueblood and Busemeyer, 2011; Wang et al., 2014). In this paper, we consider a model based on Lipovetsky (2018)'s quantum choice model. Hancock et al. (2019c) demonstrate that quantum rotations can be used to capture changes in choice contexts. In this paper, we explore whether these rotations can also be used to capture changing choice contexts. Thus a natural extension is to test whether these models can also capture dynamically changing choice contexts.

The remainder of this paper is organised as follows. First, we give an outline of what a dynamic accumulator model for dynamic choice contexts could theoretically include. Next, we detail mathematically how DFT and quantum models could account for changing choice contexts. We then describe the US101 merging dataset and detail how our models are adapted for modelling the choice to merge into a new lane or not. We then report the results of applying different variations of the models to the data. We finish by giving some conclusions as well as detailing future steps for improving dynamic models.

2. DYNAMIC MODELS: THEORY AND ELEMENTS TO INCLUDE

In this section, we first consider a theoretical example of preferences for alternatives changing over time whilst attributes of the alternatives also change. We then detail a number of features that a dynamic model could consider, before detailing how some of these can be implemented under a decision field theory model. Finally, we demonstrate how a model based on quantum probability can capture dynamic elements through quantum rotations.

2.1. A theoretical example

Whilst one method for visualising the preference for alternatives changing over time is demonstrated by Figure 1, another possibility is to consider differences between preferences, with many models already framing alternatives in this way (Busemeyer and Townsend, 1992; Noguchi and Stewart, 2018; Krajbich et al., 2012) such that thresholds are based on relative differences between the preferences for alternatives. For example, it is simple to reformulate our original example in terms of differences between the preferences, where this is shown in Figure 2, using the difference between the preferences for alternative 2 and alternative 1.

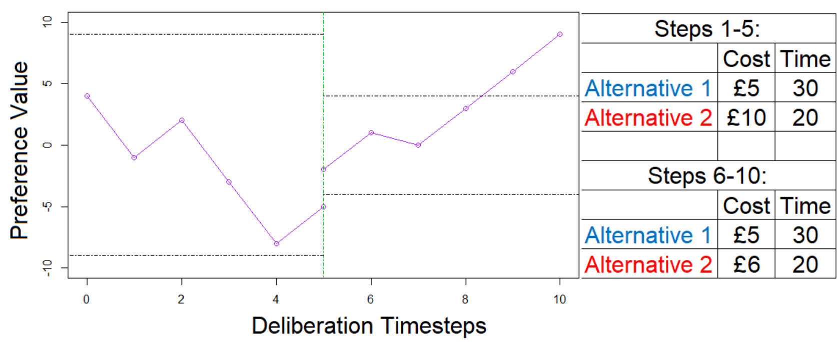

FIGURE 2 : DFT example with difference in preference values in a dynamic setting

This example additionally considers the possibility of updating threshold values and the impact of dampening effects when attributes change. These are just two of a number of factors that a dynamic model could theoretically take into account, and we now look at these in turn:

Impact of deliberation time: For the model to be dynamic, it needs to have some element that can be related to real decision time. If it includes some number of 'preference updating steps' (as depicted in Figure 2) then the decision-maker's response time taken to make a choice can be incorporated by setting the number of deliberation steps as some function of this time. This has been considered for decision field theory for static attributes (Hancock et al., 2019b).

Different rates for attribute updates compared to preference updates: In the above example, there are 5 deliberation steps for each of the two different sets of attribute values. There could clearly be a large number of deliberation steps for a slowly changing scenario such as the one above, and this is particularly the case when the acquisition of new information requires some action or comes at some cost, with many models already incorporating information search costs (Drugowitsch et al., 2012; Kim et al., 2016; González-Valdés and de Dios Ortúzar, 2018). The number of deliberation steps used for each new set of attribute values may be much lower in a continuously changing scenario, and the internal processing speed may in some cases go as low as the external updating. The rate at which an individual samples the real world information may also change over time, with for example more rapid updating the closer the person comes to making a decision. Finally attributes may not be updated from visual stimuli with every preference updating step but could be interpolated towards an individual's expectation of how the attribute will change.

Processing speeds: In a rapidly changing context, the processing of attribute changes will also be a function of how quickly an individual can process new information. For example, if there is suddenly a change in price for an ebay item very close to the ending time of the auction, a consumer may not react in time to increase their bid. For this reason, choices made over a short timeframe (such as driving decisions, for example) may be subject to both mental processing speeds and physical reaction times, with Holmes et al. (2016) finding a strong delay effect before a decision-maker integrated new information in a moving dot perception task.

Significant threshold changes: In Figure 2, the upper boundary represents the threshold for alternative 2 being chosen and the lower boundary represents the threshold for alternative 1 being chosen. Theoretically, when the attributes change, the thresholds may also change. Drift diffusion models have often considered threshold changes or 'collapsing boundaries' (Zhang et al., 2014), although recent evidence suggests that these models do not have an empirical advantage over models with a fixed threshold (Hawkins et al., 2015; Voskuilen et al., 2016).

Gradual threshold changes: Thresholds may also change more gradually than depicted in Figure 2. For example, a careful driver's preference may initially need to reach a high threshold for merging lanes if they only consider merging when the gap is particularly large. However, as the driver approaches the end of the slip road, their threshold may gradually decrease if they otherwise risk not merging into the correct lane, with drivers more likely to accept a gap if they are under stress (Paschalidis et al., 2018).

Dampening effects: If attributes change gradually, the impact of dampening may be small. However, a more significant change in attribute values such as shown in the example in Figure 2 may result in preference values resetting to zero or to their initial starting values. For example, an individual considering purchasing an item on ebay may reconsider an alternative that they had previously ruled out if the price is significantly reduced. Many process models already include 'decay' parameters (Roe et al., 2001; Usher and McClelland, 2001), but separate parameters may be required for gradual changes compared to sudden changes.

Changes in the set of available alternatives: While we have focussed on changes in attribute values, it is also possible that whilst the decision-maker is deliberating, one of the alternatives becomes unavailable. This may result in preferences resetting, or a large shift towards the most similar alternative. A typical transport example could be a cancelled train service.

Initial preferences towards alternatives: It is easily possible that there may be initial biases or preferences towards an alternative. For example, commuters will likely choose the same route to work that they chose the day before. With random utility models adept at capturing such effects and decision field theory also able to account for initial preferences (see Hancock et al. 2018), a dynamic model that accounts for changing attributes should also be able to account for initial preferences. This is particularly important given that more complex choice scenarios result in more frequent choice of the status quo (Boxall et al., 2009).

In the following section, we look at a steps towards including some of these factors in models based on decision field theory.

2.2. Decision field theory

One model that can already account for several of the above factors is decision field theory (DFT), which was first developed by Busemeyer and Townsend (1992, 1993). It is a stochastic, dynamic model for understanding choices between multiple alternatives. Under DFT, each alternative has a 'preference value' that stochastically updates over time as the decision-maker considers the different attributes for the alternatives. The decision-maker considers their alternatives for some number of deliberation timesteps until one of the alternatives reaches some internal threshold value or until the decision-maker runs out of time upon reaching some external threshold. These thresholds are represented graphically in Figure 1 by the horizontal dashed line (preference value = 5) and the vertical dashed line (deliberation timesteps = 10) respectively.

Many variations of DFT exist, with structures for models with internal thresholds (Busemeyer and Townsend, 1993) rather different to versions for multiple alternatives and multiple attributes, which have external thresholds (Roe et al., 2001; Berkowitsch et al., 2014). For the work in this paper, we consider models based on DFT models with external thresholds as this version offers a clearer method for the implementation of updating attribute values. Although analytical solutions for internal DFT are available (albeit only for choice scenarios with two alternatives, see Busemeyer and Townsend 1993), similar extensions for changing attribute values do not result in clear methods for calculating the probabilities with which the different alternatives are chosen. This is a result of the fact that there is no equivalent component to the number of preference updating steps in the calculation for the probability of alternatives under internal DFT (see Hancock et al. (2019b)). Instead, simulated approaches have to be considered for models utilising an internal threshold. Whilst an advantage of simulated approaches is that additional features as described in Section 2.1 can easily be added, the main advantage of external DFT is that we can analytically calculate the expected preference values after any number of deliberation timesteps. These calculated values can then be used as the starting values for the next DFT calculation at the point where the attribute values change.

Under DFT with external thresholds, the preference values update as follows (Roe et al., 2001):

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

where P_{t-1} and P_t are column vectors representing the previous and updated preference values, V_t represents a 'valence' vector (in effect, the utility gained during one deliberation step) and S is a feedback matrix. The feedback matrix is typically specified to include decay and sensitivity parameters that allow for the impact of contextual effects such as attraction, compromise and similarity effects (Berkowitsch et al., 2014). However, for simpler cases where there are only two alternatives, we do not need to control for these effects as previous applications of DFT have found that the sensitivity parameter can becomes meaningless for choice scenarios with only two alternatives (See Hancock et al. 2019a). For such cases the feedback matrix can simply act as a decay parameter:

$$S = \phi_2, \tag{2}$$

with $0 \leq \phi_2 \leq 1$. This additionally allows for a simpler calculation for the probabilities with which the alternatives are chosen (as defined by Equations 4-10).

Finally, the valence vector, V_t, is determined based on the attribute attended to by the decision-maker at deliberation timestep t. It is defined:

$$V_t = C \cdot M \cdot W_t + \varepsilon_t, \tag{3}$$

where C is a contrast matrix used to rescale the attribute values M around 0. It is defined to have diagonal elements of 1 and off-diagonal elements of $-1/(x-1)$, where x is the number of attributes. W_t is a column matrix of zeros with a single 1 in the row representing the attribute that is attended to at time t. A DFT model will hence estimate a set of weights, w_a, w_b,.., w_k, with w_k giving the probability that attribute k is attended to in a given preference updating step. To allow for adjustments to the attributes M, as well as avoiding the requirement of a priori knowledge of the attributes (See Hancock et al. (2019a) for details of this), we fix each of these weights to $1/x$, and estimate scaling parameters for the attributes. This allows for the application of, for example, income effects to the attributes for DFT in the same way that these effects are added to typical random utility models. Additionally, there is a noise parameter ε_t, which adds normally distributed values to the valence for each alternative (with a mean of zero, independent draws and a standard deviation that is estimated).

To calculate the probabilities for the different alternatives, a DFT model simply requires the expected value of the preference vector after some number of deliberation timesteps, Φ_t, and the covariance of the preference vector Ω_t (Roe et al., 2001). With the feedback parameter S now simply represented by a constant, ϕ_2, the expectation of the preference vector after t preference updating steps becomes:

$$E[P_t] = \xi_t = \frac{1-\phi_2^t}{1-\phi_2} \cdot \mu + \phi_2^t \cdot P_0, \tag{4}$$

where P_0 is the initial preference matrix, which can be used to calculate underlying baseline preferences towards an alternative (Hancock et al., 2018) and μ is the expectation of the valence vector, V_t.

The covariance matrix for a single set of attributes, M, also simplifies to become:

$$Cov[P_t] = \Omega_t = \frac{1-\phi_2^{(2\cdot t)}}{1-\phi_2^2} \cdot \Phi, \tag{5}$$

where Φ is the covariance matrix for V_t (Roe et al., 2001).

However, it is worth noting that if $P_0 = 0$, t and ϕ_2 cannot be separately identified as the full set of probabilities generated by all possible combinations of t and ϕ_2 can be generated by t alone with $\phi_2 = 0$.

First, note that when $\phi_2 = 0$, the calculated values for the expectation and covariance are $\xi_t = t \cdot \mu$ and $\Omega_t = t \cdot \Phi$. Then as P_t converges to a multivariate normal distribution (Roe et al., 2001), the same set of probabilities are generated with μ and Φ or $x \cdot \mu$ and $x^2 \cdot \Phi$, where x is some constant. Consequently, the set of probabilities that are calculated with ξ_t and Ω_t are also generated by Φ and:

$$\frac{1-\phi_2^t}{1-\phi_2} \cdot \mu \cdot \sqrt{\left(\frac{1-\phi_2^2}{1-\phi_2^{(2\cdot t)}}\right)}, \tag{6}$$

which can be rearranged to:

$$\sqrt{\left(\frac{(1+\phi_2)\cdot(1-\phi_2^t)}{(1+\phi_2^t)\cdot(1-\phi_2)}\right)} \cdot \mu. \tag{7}$$

Then, as $t > 1$ and $0 < \phi_2 \leq 1$, this is equivalent to $y \cdot \mu$, where $y \geq 1$. Consequently, we only need to estimate one parameter, y, which is equivalent to estimating $\sqrt{(t)}$ when $\phi_2 = 0$. Note that if $P_0 \neq 0$, then the value of ϕ_2 will have an impact. It is however simpler to estimate y as before, as well as some new factor, α_d, which is the factor which previous preference values are multiplied by. Consequently, if the attribute values change (for example from M_1 to M_2) then we can define the expected preference values at the end of the deliberation process to be:

$$\xi_{2t} = t \cdot \mu_2 + \alpha_d \cdot t \cdot \mu_1 + \alpha_d^2 \cdot P_0, \tag{8}$$

where μ_1 and μ_2 are the expectations of the valence vectors for the different attribute matrices M_1 and M_2, and the number of preference updating steps for M_1 and M_2 are assumed to be equal.

To generalise this for R different sets of attribute matrices M_1, M_2,..., M_r, the total expected value can be calculated recursively. This results in a total expectation of:

$$\xi_{rt} = t \cdot \sum_{i=1}^{R} \left(\mu_i \cdot \alpha_d^{R-i}\right) + \alpha_d^r \cdot P_0 \tag{9}$$

To allow for an appropriate amount of decay for the covariance matrix relative to the expected preference, we instead multiply the covariance by α_d^2. We can then calculate the covariance matrix after R different sets of attribute matrices, for which Equation 5 can be expanded to:

$$\Omega_{rt} = t \cdot \sum_{i=1}^{R} \left(\Phi_i \cdot \alpha_d^{R-i}\right) \tag{10}$$

Once the expectation and covariance matrices are calculated for the preference vector, the probability with which the different alternatives are chosen can be calculated using multivariate normal distributions (Roe et al., 2001).

2.3. Quantum rotation models

Another alternative structure for a dynamic choice model is to use quantum rotation models. Under these models, each alternative is represented by a vector in some multidimensional Hilbert space. A 'state' vector is then used to define the probabilities of the decision-maker choosing each of the alternatives. Hancock et al. (2019c) demonstrate that for each alternative, a 'projection length' must be estimated to generate this state vector. These projection lengths can then be used to estimate the probability with which the alternative is chosen, given that the sum of squared projection lengths must sum to one (Hancock et al., 2019c). For a binary scenario, one possibility is to consider quantum projection lengths based on trigonometric functions (Lipovetsky, 2018), where we define the probability of the first choice as:

$$Prob(Alt_1) = cos^2\left(min\left(\frac{\pi}{2}, max\left(0,z\right)\right)\right), \tag{11}$$

where z corresponds to an aggregate of predictors[1]. The probability of the second alternative is simply calculated with a sine rather than cosine function. The functional form of z can then be based on any relevant attributes of the alternatives to generate appropriate probabilities for each alternative being chosen. To add a dynamic element to the model, we then require the use of a quantum rotation. A general notation for a quantum rotation is given by Equation 15 of Hancock et al. (2019c) for changes in choice context. The extension to dynamic changes in choice context can be made by specifying functional forms for elements of the rotation matrix, M^*, based on dynamic components:

$$M^* = \begin{bmatrix} 1 & df * r_{1,2} \\ df * r_{2,1} & 1 + df \end{bmatrix}, \tag{12}$$

where $r_{1,2}$ and $r_{2,1}$ are estimated rotation parameters and df is based on some dynamic function of contextual effects. Figure 3 gives an illustration of how the state vector may change under dynamic choice settings. Initially, it is of unit length. Attributes of the alternative impact the angle of the state vector with the vectors for each of the alternatives (represented by the axes in the figure).

The quantum rotation is then denoted by the pink arrow, which demonstrates how the choice context may impact the state vector. Crucially, as sine and cosines can be used to calculate the

[1] Note that the use of minimum and maximum functions here restricts $cos^2(x)$ into being a strictly monotonic function.

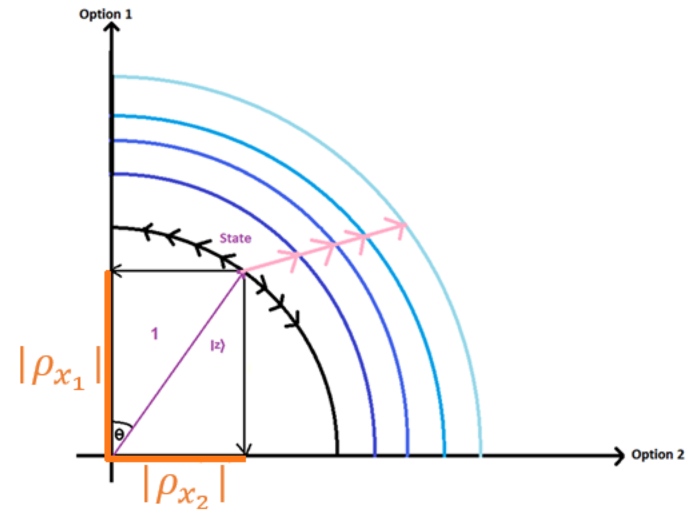

FIGURE 3 : An illustration of how the probabilities of alternatives may change under dynamic settings in quantum choice models.

probability of alternatives based on the angle between the state vector and the vectors for the alternatives, the state vector does not have to remain of unit length to generate probabilities. Thus, in a dynamic choice scenario, changes in the attributes of alternatives are represented by the black arrows on the arc corresponding to circle of radius one and changes in the choice context are represented by the pink arrows which move away from this arc.

3. EMPIRICAL FRAMEWORK

In this section, we detail the driving dataset that is used in the empirical applications in this paper. We discuss how it is implemented for our logit, probit, quantum and DFT models, as well as describing how we add a dynamic element to logit, probit and quantum such that a fair comparison can be made with DFT (which is dynamic by nature).

3.1. Data for case study

For testing our dynamic models, we use a dataset collected on site on the southbound section of the 101 Highway in Los Angeles in June 2005 (Cambridge Systematics Inc., 2005b). It comprises of a 45 minute video recording of vehicles travelling across a 640 metre cross-section of the road, which has 5 lanes as well as an auxiliary lane which connects an on-ramp and off-ramp (see Figure 4).

This provides ideal dynamic scenarios for which naturally dynamic models such as decision field theory can be applied to. For testing the models, we look at the lane merging behaviour by the drivers who join the US-101 from the Ventura on-ramp. These drivers can join lane 5 immediately at the end of the on-ramp, or wait to merge at some point as they travel down the auxiliary lane (lane 6 in Figure 4). The factors impacting this decision (such as the size of the gaps on lane 5,

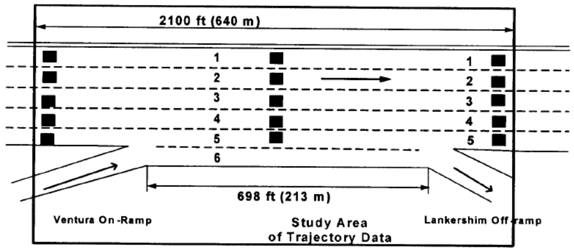

FIGURE 4 : The US-101 data collection site (Reprinted from Choudhury 2007).

speed and acceleration of the vehicles involved, etc.) are constantly changing, as a result of both the vehicles trying to join lane 5 and those already on it changing speed. Over the course of the 45-minute video, 399 vehicles start on the on-ramp and merge onto lane 5. We use trajectory data which details the exact location and speeds of all the vehicles that feature in the video (Cambridge Systematics Inc., 2005a) at intervals of 0.1 seconds, as well as giving information on the length, width and type of vehicle. This results in hundreds of observations for each vehicle, with an average time of 8.73 seconds taken by a vehicle to join the US-101 from the Ventura on-ramp from the first point at which they could make the merge.

Some of the vehicles merge immediately from the on-ramp to lane 5, whilst others do not merge until they reach the Lankershim off-ramp. Whilst there is likely to be significant variation in driver perception-reaction times (Fu et al., 2016; Wood and Zhang, 2017; Paschalidis et al., 2019), we simplify our models by assuming that drivers take a second to react to visual stimuli. Thus, we define that the driver chooses to merge lane due to visual stimuli at a point a second before their vehicle physically changes lane. The drivers also have the external constraint that they must merge by at a certain point which acts as an external threshold in the formulation. In order to estimate the model to predict the probability that a driver will choose to merge at a certain time step, we also require observations in which they choose not to merge lanes. For the majority of the applications in this paper (with the exception of results in Section 4.5), we use observations at 1 second intervals before the moment where the driver chooses to merge (but do not consider any moments after the driver has made the decision to merge lanes). This results in a total of 3,293 observations across 395 vehicles.

3.2. Explanatory variables

For all the variations of models that we use, we define two utilities/preference values, one for merging and one for staying in the current lane. Then, for each observation in the dataset, we use four key attributes to define these utilities/preference values:

1. Time headway in the target lane in front of the merging vehicle (in seconds).
2. Time headway in the target lane behind the merging vehicle (in seconds).
3. The velocity of the merging vehicle (feet per second).

4. Distance (in feet) to the point at which merging lane is no longer possible (in this case, shortly after the driver reaches the Lankershim off ramp).

3.3. 'Static' models

3.3.1. Logit models

In our logit models, the utility for a driver n choosing to merge at some time point t is assumed to be a function of the time headways and velocity of the merging vehicle:

$$V_{MERGE_{nt}} = U_{GF_{nt}} + U_{GB_{nt}} + U_{VEL_{nt}} + \varepsilon_{nt}, \tag{13}$$

where $U_{GF_{nt}}$ is the utility gained from the gap in front, $U_{GB_{nt}}$ is the utility gained from the gap behind, $U_{VEL_{nt}}$ is a utility that is dependent on the velocity of the merging vehicle and ε_{nt} is a random error term. Again, other specifications would be possible, such as interacting the gaps with the velocity, but this is an exploratory first step.

To calculate this utility, we then need to define specifications for the three deterministic utility components. Firstly, given that the impact of time headway is very non-linear (with increases in headway being much more important at small base values), we use a logistic transform to convert the time headway to a subjective utility[2] for the gap in front ($U_{GF_{nt}}$):

$$U_{GF_{nt}} = \beta_{GF} * \left(2 * \frac{exp(\alpha_{GF} * TH_{F_{nt}})}{1 + exp(\alpha_{GF} * TH_{F_{nt}})} - 1\right), \tag{14}$$

where β_{GF} is an estimated weight for the relative importance of the gap in front, α_{GF} is an estimated scaling parameter and $TH_{F_{nt}}$ is the time headway in front of the merging vehicle. We multiply this logistic transformation by 2 and take away 1 such that this utility can be both positive and negative. Note that some observed time headways are negative, which means that the vehicle in front in the target lane is partially adjacent to the merging vehicle (see an example of this in Figure 5).

As the time headways only account for the distance between the cars in one-dimension, along the length of the direction of travel (for example, in Figure 5, the orange arrow represents the (negative) gap between the back of the car in front, and the front of the merging car. This gap does not account for distances between the vehicles perpendicular to the direction of travel (the dashed line in Figure 5), thus we do not automatically translate negative time headways to extremely negative subjective values, as a vehicle may begin to merge whilst being adjacent to the vehicle in front if the relative speeds of the vehicles are very different.

The time headway in front ($TH_{F_{nt}}$) of the vehicle is defined as the gap size between the front of the merging vehicle and the back of the vehicle in front in the target lane, with respect to the

[2]Note that different linear and non-linear functional forms for mapping the four attributes to utilities were tested based on literature and the final functional form has been selected based on empirical testing.

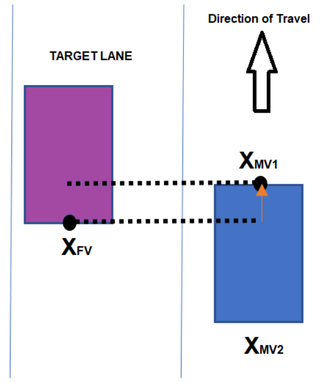

FIGURE 5 : A merging car adjacent to the vehicle in front on the target lane.

speed of the vehicle in front. Thus for decision-maker n at time point t, we have:

$$TH_{F_{nt}} = \frac{x_{VF_{nt}} - x_{MV1_{nt}}}{max(1, vel_{VF_{nt}})}, \tag{15}$$

where $x_{VF_{nt}}$ is the location of the back of the vehicle in front, $x_{MV1_{nt}}$ is the location of the front of the merging vehicle and $vel_{VF_{nt}}$ is the velocity of the vehicle in front. The gap size is divided by a minimum value of 1 such that the gap is not transformed to a large value should the vehicle in front be stationary or moving extremely slowly (with a velocity of less than 1 foot per second).

We then estimate parameters β_{GB} and α_{GB} for an equivalent adjustment for the gap behind to that of the adjustment for the gap in front as defined by Equation 14, to get a utility U_{GB} for the gap behind. For this, the time headway behind (TH_B) the vehicle is defined as the gap size between the back of the merging vehicle and the front of the vehicle behind in the target lane, with respect to the speed of the vehicle behind:

$$TH_{B_{nt}} = \frac{x_{MV2_{nt}} - x_{VB_{nt}}}{max(1, vel_{VB_{nt}})}, \tag{16}$$

where $x_{VB_{nt}}$ is the location of the front of the vehicle behind, $x_{MV2_{nt}}$ is the location of the back of the merging vehicle and $vel_{VB_{nt}}$ is the velocity of the vehicle behind.

Finally, the utility for choosing to merge also accounts for the speed of the merging vehicle:

$$U_{VEL_{nt}} = \beta_{VEL} * vel_{MV_{nt}}, \tag{17}$$

with β_{VEL} an estimated parameter and $vel_{MV_{nt}}$ the velocity of the merging vehicle. The driver may alternatively decide to stay in the same lane. We define the utility of staying in the same lane as a function of the distance traversed thus far:

$$U_{STAY_{nt}} = \alpha_{STAY} + \beta_{dist} * MV_{DE_{nt}} + \alpha_{AL} * MV_{AL_{nt}} + \alpha_{OR} * MV_{OR_{nt}}, \tag{18}$$

where $MV_{DE_{nt}}$ is the distance between the merging vehicle and the last possible point at which they can merge onto the US101 and β_{dist} an estimated parameter for the relative impact of this distance. We then additionally estimate a constant α_{STAY} for choosing not to merge lanes and two further parameters α_{AL} and α_{OR}, which control for the impact of the driver being in the auxiliary lane (for which $MV_{AL_{nt}} = 1$) or on the off ramp ($MV_{OR_{nt}} = 1$).

The final specification for our logit model uses 9 estimated parameters. An assumption of type I extreme value distributions for ε_{nt} results in typical MNL choice probabilities.

3.3.2. Probit models

A first alternative model considered is a restricted form of probit model. With the assumption of no correlation between the error terms, a probit model also resembles a restricted form of decision field theory. A DFT model comprises of two key sources of error: the variation generated by which attribute the decision-maker attends to at each preference updating step, as well as the normally distributed error terms added to the valence vectors (see Equation 3). The first of these, however, becomes insignificant if the estimate for the number of deliberation timesteps is very high (which results in attribute attendance variation being averaged out). This leaves the source of error in the model coming solely from the normal error terms, the sum of which is also a normal (and hence the model becoming a restricted probit).

3.4. Dynamic specification of 'static' models

In this section we consider features for going beyond the basic models defined in the previous section. Whilst models that attempt to capture the accumulation of preference over time may provide a more 'natural' account of dynamic choice contexts, models for which the preference (or more typically, utility) for an alternative is calculated for an instantaneous moment can incorporate previous preferences through lagged variables. For example, attributes of previously chosen alternatives can be entered as a function of how similar they are to the alternatives for the current choice (Erdem, 1996), though special care is of course needed in relation to endogeneity risks. Lagged variables have also been used in the context of consumer preferences (Seetharaman, 2004) and environmental economics (Swait et al., 2004). More specifically, the latter of these examples defines 'meta-utilities' for an alternative as a function of the product of utilities in current and previous periods. Thus dynamic discrete choice models can include parameter(s) similar to the decay parameter used in DFT, which can be adopted into our logit, probit and quantum models. The most basic approach for the incorporation of a decay parameter is to set the total utility/preference (Q_t) at a time point t as the sum of the utility/preference at time t, (P_t), combined with the utility/preference at previous time points which are subject to a decay parameter $0 \leq \alpha_d \leq 1$:

$$Q_t = \sum_{r=0}^{t-1} \left(\alpha_d^r P_{(t-r)} \right), \tag{19}$$

where the instantaneous-only model is obtained with $\alpha_d = 0$.

Another alternative is to use hyperbolic discounting functions (Mazur, 1987). Here we consider the generalised hyperbolic discounting function (Green et al., 1994), which results in the total utility at time t being defined:

$$Q_t = \sum_{r=0}^{t-1} \left(\frac{P_{(t-r)}}{(1+\alpha_d \cdot r)^s} \right), \tag{20}$$

where s is some estimated factor for the discount rate. Examples of these discounting curves are given in Figure 6, with $\alpha_d = 0.5$ and $s = 1$ for the hyperbolic discounting function and $s = 2$ for the generalised hyperbolic discounting function.

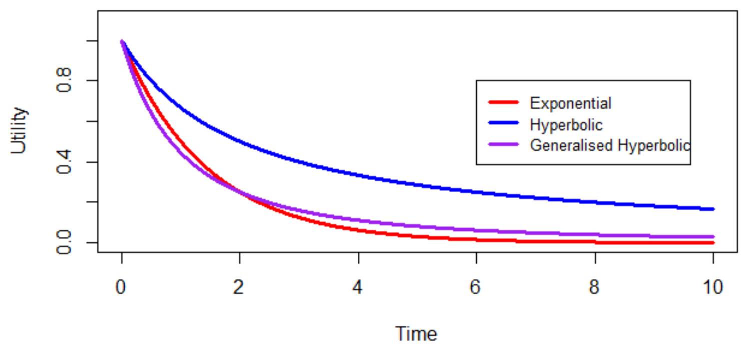

FIGURE 6 : An example of each of the discount curves tested in this paper.

Both discounting functions have been tested in dynamic choice models previously, with an exponential discounting factor used by Swait et al. (2004) and Aguirregabiria and Mira (2010) and a hyperbolic discounting function used by Fang and Wang (2015). Whilst hyperbolic functions have empirical support in a number of different fields (Kirby and Maraković, 1995; Green and Myerson, 1996; Madden et al., 1999), examples from intertemporal choice illustrate that they do not always outperform exponential discounting functions (Read, 2001; Rubinstein, 2003).

The inclusion of these dynamic elements in the otherwise static models bridge the gap with truly dynamic models, yet do not move us away from the fact that these approaches still rely upon the stitching together of individual choices.

3.5. DFT specification

As far as the authors are aware, decision field theory in its current form has never been applied to driving behaviour, although Zhao et al. (2011) used a version 'rule-based decision field theory' in their model considering driver speed control. Whilst the choice to merge lanes is determined by the driver (and thus is more akin to a decision-maker reaching an internal threshold), there is also an external threshold in that the driver will pass a point after which merging lanes is no longer possible. Additionally, as we have trajectory data for intervals of the same length, a DFT approach as described by Equations 1-10 fits well.

A key distinction arises between DFT in estimation and in application. In application, with a given set of parameters, we can simulate the accumulation of preference until a threshold is exceeded and a decision is made, or until the decision maker runs out of deliberation time. In estimation however, we are concerned with finding values for the model parameters that explain the behaviour observed in the data. Just as with 'standard' models, this requires the maximisation of an objective function such as log-likelihood, and we need to contrast observed events with estimated probabilities. We thus need to not only estimate the probability of merging at the point where a decision is made, but similarly the probability of not merging at the time points before this event. Under DFT with an external threshold, the chosen alternative is the one with the higher preference value at some time point t. Thus, a threshold for choosing to merge can be equated to the point at which the preference for merging overtakes the preference for staying in the same lane. Hence, to make a fair comparison of DFT with alternative models, we estimate the probability that the preference value for merging is higher than the preference value for staying in the same lane after each 1 second interval. The resulting required transition from a logit to a DFT is then relatively simple.

We use attribute scaling parameters rather than attribute weights, which means that time headway in front of the vehicle, time headway behind the vehicle, the velocity of the merging vehicle and the distance (Equations 15 - 18) can be entered as four different attributes that update over time. Thus, the attribute matrix, M to be used in the DFT model can be based on the utility functions defined previously:

$$M = \begin{bmatrix} U_{GF} & U_{GB} & U_{VEL} & 0 \\ 0 & 0 & 0 & U_{STAY} \end{bmatrix}. \tag{21}$$

A dynamic version of DFT with updating attribute values can then be estimated using Equations 1 - 10. As well as the same nine attribute parameters used for MNL, this model then requires three process parameters. These are the standard deviation of the valence error, the number of deliberation timesteps and a decay parameter α_d to account for the dampening of preferences when attributes update (see Equations 7-10). This allows for the current attributes for the choice as to whether to merge or not to have more impact than previous attributes, with the preferences from old time points decaying systematically over time. Note that the decay parameter ϕ_2 cannot be separately identified for our DFT model (see Section 2.2). Additionally, as we use scaling parameters, the standard deviation of the error must be fixed to avoid confounding. This leaves just two additional parameters.

The above use of DFT relies on the closed form calculation of probabilities, which is computationally desirable, but where the use of expectations limits the additional dynamic complexities that can be introduced. We can alternatively adopt a simulated approach, for which only Equations 1-3 apply. We have two updating preference values, one for the preference to merge, and one for the preference to stay in the same lane. The decision-maker repeatedly evaluates the attributes which results in the preference values updating over time. For some preference updating steps, the decision-maker 'samples the world' and thus the attribute matrix M in Equation 1 is updated. At other preference updating steps, the decision-maker does not update M, thus instead effectively 'internally resamples' the information that they already have. In this case, in the line with the other models tested, we allow updates for M to happen once per second (though of course insights from cognitive psychology could be used to inform and test different rates in future applications). We then trial different numbers of internal preference updating steps. At each preference updating step, one of the four attributes is selected randomly (with equal probability), with the preference values for merging and staying in the same lane updated in line with Equations 1 - 3. After some set number of preference updating steps, the probability of merging at each second interval can be estimated based on the mean and standard deviations for the difference between the preference values (which will also converge due to the central limit theorem), meaning that the probabilities can be calculated directly. For example, 200 simulations will give 200 estimates for the preference difference at each second interval. These values can then be used to generate a distribution which can in turn be used to calculate the probability of choosing to merge.

3.6. Quantum model specifications

For our quantum model specifications, we do not consider the two specifications we have previously developed in Hancock et al. (2019c), as these both rely on the comparison of attributes across alternatives. Given that the decision to merge or not is a binary choice, we instead use Equation 11. We can then define two different quantum models.

The first is 'static'. For this quantum model (QM), we set $z = \frac{\pi}{4} + U_{STAY_{nt}} - U_{MERGE_{nt}}$, such that Equations 13 and 18 can be implemented into the quantum model (where the addition of $\frac{\pi}{4}$ results in base probabilities of 0.5). Adjustments to $U_{STAY_{nt}}$ and $U_{MERGE_{nt}}$ through the use of discounting curves can then add dynamic elements, exactly as implemented for the logit and probit models.

The second variation is based on the quantum rotation model (QRM) framework described in Section 2.3. In this case, z is adjusted to $z = \frac{\pi}{4} - U_{MERGE_{nt}}$. The factors that enter $U_{MERGE_{nt}}$, which are all based on how far the decision-maker is from the end of the slip road, are instead used to represent the dynamic 'changing choice context.' Thus a rotation matrix as defined by Equation 12 is used, where $df = U_{MERGE_{nt}}$.

4. RESULTS

In this section, we apply decision field theory models accounting for changing attribute values to the US-101 dataset, as well as providing comparisons based on logit, probit, quantum and quantum rotation models. We start by considering models that do not account for the dynamic nature of the dataset, which treat each choice at one second intervals as completely independent of the choices made by the same driver at later time points. We then apply both analytical and simulated versions of DFT models, comparing their performance as well as testing the impact of changing the number of attribute updates per second. Next, we apply alternative model structures based on 'static' models with variables that account for 'remembered' utilities (see Section 3.4), which adds a dynamic component to these models. Finally, we look at the inclusion of initial starting preferences, testing whether these additional parameters impact all models or just DFT models.

4.1. Basic static models

For our first set of models, we analyse the choices independently, without including any memory parameters (as detailed in Section 3.4) in any of the models. We thus test a basic version of MNL as described by Equations 13 and 18 as well as probit and quantum models, as described in Section 3.3.2 and Section 3.6. We can also test a 'static' DFT, for which the parameter α_d is fixed to a value of 0, meaning that the choices at different timepoints are not linked. For all models, we use R packages maxLik (Henningsen and Toomet, 2011) and apollo (Hess and Palma, 2019) for estimation of the log-likelihood functions.

TABLE 1 : Basic static model results

Model variation	Free Parameters	Log-likelihood	BIC
Logit	9	-862.59	1798
Probit	9	-858.59	1790
Quantum	9	-856.21	1785
DFT-1	10	-858.61	1798
DFT-2	9	-858.81	1791

The performance is relatively similar across models, where the best performing model at this point is the quantum model. The estimate for the number of preference updating steps for the first DFT model (DFT-1 in Table 1) is very high and consequently there is no significant difference for a DFT model (DFT-2) for which the number of steps is fixed to 1000. As it has a fit that is equivalent to the probit model, it appears that there is no gain that can be attributed to stochastic variability that is generated through attention to different attributes.

For all models, we see that the size of the gap behind tends to be more important than the gap in front, with Figure 7 giving the contribution of U_{GF} and U_{GB} as a function of the time headways in front and behind the merging vehicle. Whilst all models appear to show very similar results, the damping in the quantum model is less extreme, with some differences between 1 and 2 seconds. This is likely a key reason for the better fit for this model.

4.2. Comparison of simulated and analytical DFT models

For our first comparison of simulated and analytical DFT models, we test the impact of fixing the number of preference updating steps. For all of the simulated models, we run 200 simulations with the use of the R package RCPP (Eddelbuettel et al., 2011). For each model, we assume that the decision-maker samples the world at one second intervals (thus attribute matrix M is updated once per second) and test rates of 1, 5, 10, 20, 50 and 100 (internal) preference updating steps per second (where for analytical DFT, this rate determines the value of t in Equations 9 and 10, and for simulated DFT, there are simply t iterative updates of Equation 1 per second), obtaining the results

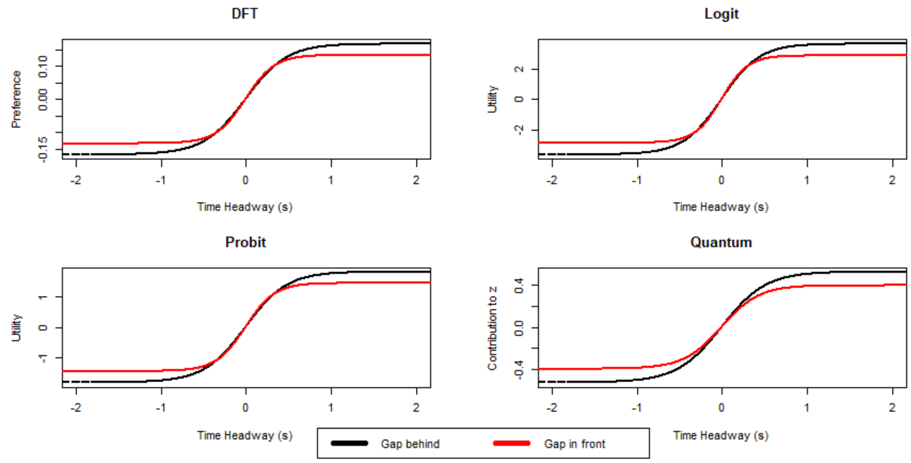

FIGURE 7 : The conversion of time headway into preferences/utilities

given in Table 2.

TABLE 2 : A comparison of log-likelihoods from analytical and simulated DFT models with different numbers of preference updating steps

Preference updates per second	Sim DFT	Analytical DFT	Difference
1	-1,038.51	-1,122.82	-84.30
5	-931.00	-934.99	-3.99
10	-877.33	-883.55	-6.22
20	-860.50	-856.08	4.42
50	-851.56	-843.75	7.81
100	-843.83	-840.40	3.43

The results demonstrate that analytical DFT has a worse performance for a low number of timesteps but does better when the number of steps is high. Notably, analytical DFT is reliant on the central limit theorem (for the preference values to converge to a multivariate normal distribution) and it appears that this becomes a poor approximation when the number of steps is low. Whilst the best performance in this case is for an analytical model, these results imply that models for which the estimated number of preference updating steps is low (which was not the case for DFT-1 in Table 1) may perform better with a simulated version of DFT. For example, if we instead assumed that individuals sampled the world at more frequent intervals than once per second (and thus there were more frequent updates for attribute matrix M and consequently a smaller number of steps t for each set of attribute matrices if the total number of preference updating steps per second is to remain the same) then simulated models may be preferable. The simulated approach also of course makes it easier to incorporate additional dynamic elements, as discussed earlier in the paper.

An illustration of some of these models is given in Figure 8, which shows 200 sets of simulated preference values for one of the drivers under a simulated DFT model (with $t = 100$).

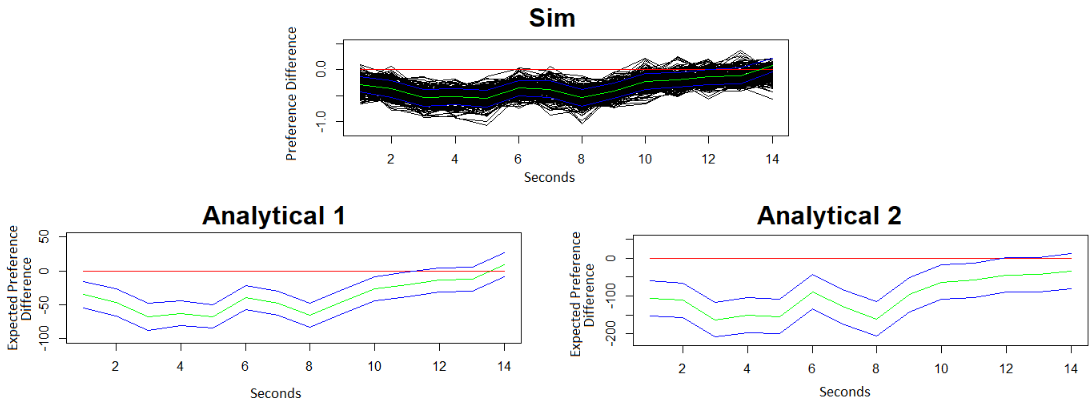

FIGURE 8 : Evolution of preference values under DFT models

The figure also shows the expected preference values for two versions of analytical DFT, a model with a memory parameter (Analytical 1) and one without (Analytical 2). For all models, preference differences above the threshold of zero (which is indicated by the red horizontal line) imply that the driver will choose to merge. The green lines give the average and expected preference values for the simulated and analytical versions of DFT respectively. Finally, the blue lines give the standard deviations of these values (which are calculated directly for simulated DFT, and estimated based on the covariance of the preference values for analytical DFT). For the simulated model, all 200 simulations are given. For example, we can see that after 4 seconds not a single simulation finds a positive difference (which corresponds to choosing to merge) This additionally illustrates why we do not simply evaluate the probabilities of merging through summing the number of simulations for which the difference is positive, as this would result in zero probabilities. Note that these probabilities are not cumulative, but are the probability of a driver merging at that moment. Notably, the key difference between models without a memory parameter (Analytical 2) and the other models is that the probability of choosing to merge at the final time point (at which point the driver chooses to merge) is lower. This is also demonstrated in Table 3, which gives the average probability of observing the chosen alternatives under three different DFT models (corresponding to the three models in Figure 8), as well as the full set of probabilities for the same driver. As a contrast, all models perform similarly for the timepoints where the driver does not merge.

4.3. Models accounting for decay parameters

Whilst DFT provides a natural method for the evolution of preferences over time, models that calculate probabilities based on a single utility calculation can also capture the impact of previous preferences through decay parameters (see Equations 19 and 20). This results in a significant improvement in model fit for all of the models, with the log-likelihoods of models implementing an exponential discounting function given in Table 4, and analytical DFT models (with free timestep parameters) also displayed for comparison. We simplify all models by assuming that the unobserved errors are independent over time.

For the models with and without a decay parameter, there is very little difference between the different types of models, although the quantum model is no longer the best performing model

TABLE 3 : Probabilities for 14 'choices' by one driver under four different DFT models

Model		Simulated	Analytical 1	Analytical 2
deliberations per second		100	100	1000
α_d		0.35	0.25	0.00
Log-likelihood		-843.83	-840.40	-858.81
Average stay probability		0.9075	0.9077	0.9061
Average merge probability		0.3245	0.3230	0.3057
Seconds	1	0.9693	0.9627	0.9890
	2	0.9886	0.9887	0.9920
	3	0.9997	0.9996	0.9998
	4	0.9996	0.9997	0.9995
	5	0.9997	0.9999	0.9997
	6	0.9921	0.9877	0.9741
	7	0.9892	0.9961	0.9978
	8	0.9995	0.9999	0.9998
	9	0.9979	0.9951	0.9827
	10	0.9323	0.9323	0.9173
	11	0.9140	0.8672	0.8974
	12	0.8310	0.7747	0.8312
	13	0.7847	0.7449	0.8279
	14	0.2611	0.3107	0.2272

TABLE 4 : Performance and model outputs for models with basic decay parameters

	Without memory	With memory		
	Log-likelihood	Log-likelihood	α_d estimate	β_{GF}/β_{GB}
Logit	-862.59	-840.36	0.227	0.732
Probit	-858.59	-838.29	0.237	0.755
Quantum	-856.21	-837.67	0.235	0.775
DFT	-858.81	-837.48	0.250	0.759

when a decay is included. Additionally, there is now also a small difference between DFT and probit, although this difference is not significant. The estimate for the decay parameter is slightly higher for DFT than the other models. Table 4 also gives the ratio β_{GF}/β_{GB}, which gives the estimated relative importance of the gap in front of the vehicle compared to the gap behind the vehicle. The results here demonstrate that all models find that the gap behind is more important, with the quantum model assigning the highest importance to the gap in front relative to the other models.

Next, we consider improving the models by increasing the flexibility of the decay parameters. For all models, we consider hyperbolic discounting functions as well as utilising different discounting parameters for the utility to stay and the utility to merge. The results of these models are given in Table 5. For the hyperbolic functions, we consider two variations. The first of these is based on the original hyperbolic discounting function (for which s is fixed to 1 in Equation 20). The second uses both hyperbolic discounting parameters.

Three clear patterns emerge from these models. Firstly, quantum models perform slightly but not significantly better than probit models, but significantly better than logit models. Secondly, additional parameters allowing for different decay rates for the utility for merging compared to the utility for staying in the same lane significantly improve model fit. Finally, exponential decay models perform better than hyperbolic decay models, but worse than generalised hyperbolic decay models (for which s in Equation 20 is a free parameter). The shape of the discount curves estimated are given for the best exponential and best hyperbolic versions of each model in Figure 9.

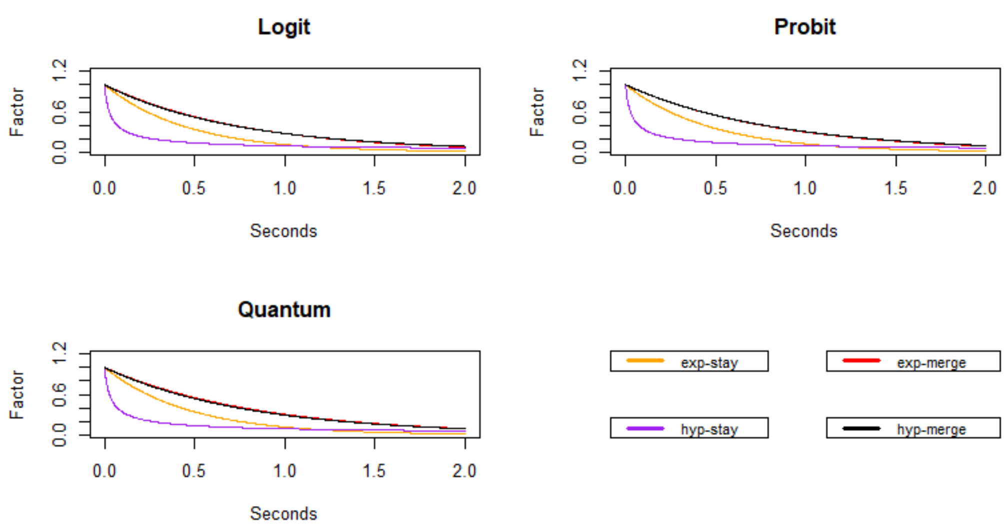

FIGURE 9 : The discount curves estimated by the different models

Notably, the same patterns emerge for all of the different models. There is no difference in the decay for the utility to merge, regardless of whether hyperbolic or exponential discounting decays are used, but there is a substantial difference for the decay for the utility to stay. This implies that the gain in model fit found by the hyperbolic decay models is possibly due to the increased flexibility that allows for an initially sharp decay before a more gradual decrease after 1 second

TABLE 5 : Log-likelihoods from models utilising additional decay parameters

Discount function	Exponential	Exponential	Hyperbolic	Hyperbolic	Hyperbolic
Memory parameters	1	2	1	2	4
Equivalent decay	yes	no	yes	no	no
Logit	-840.36	-834.91	-847.05	-840.57	-830.45
Probit	-838.29	-831.97	-843.81	-838.46	-827.78
Quantum	-837.67	-830.53	-842.13	-837.93	-828.73

(the purple lines in Figure 9. Furthermore, Figure 10, which gives the evolution of utilities under logit, probit and quantum models (for the same driver whose preferences are illustrated in Figure 8), suggests that the addition of decay parameters appears to decrease the probability of choosing to merge. As before with DFT, the models predict a greater probability of choosing to merge at the point where the driver does actually choose to merge.

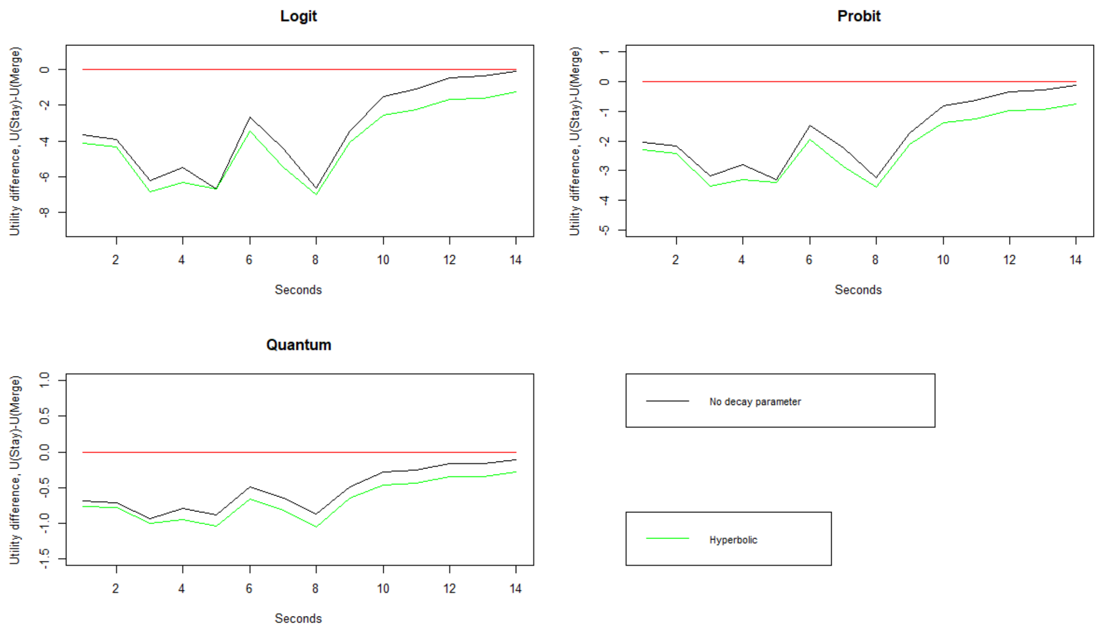

FIGURE 10 : Evolution of utilities under logit, probit and quantum models

4.4. Inclusion of initial preferences

Thus far, the alternative specific constant that we have utilised in our models has been applied for every new set of attribute values. Whilst this parameter helps capture the baseline preferences for alternatives (note that the drivers choose to merge at only 12% of the timepoints), it does not capture a baseline preference at the start of the decision process. For example, a driver may initially be keen to merge lanes, but then quickly realise that it is not possible upon approaching the motorway. For DFT models, this initial preference can easily be captured by estimating P_0 in Equation 1. For the other models, this initial preference can be captured through the addition of a constant in one of the utilities/projection lengths (that is applied at the first time point only). The results of adding this parameter are given in Table 6. The logit, probit and quantum models without the initial preference parameters are based on the best fitting models displayed in Table 5. We also give the results for quantum rotation models, where there are no memory parameters, but dynamic effects are captured through the incorporation of quantum rotation matrices. Finally, for analytical DFT, the model without an initial preference parameter has the number of deliberation timesteps fixed to 1000.

Crucially, DFT is the only model that benefits from the addition of an initial preference parameter. It is also worth noting that a simulated version of DFT improves significantly more than an

TABLE 6 : The impact of adding an initial preference parameter on the log-likelihood for the different models

Model	Decay Type	without	+ P0	Difference
Simulated DFT	DFT	-840.76	-824.57	16.19
Analytical DFT	DFT	-837.48	-831.44	6.04
Probit	Exp	-831.97	-831.39	0.58
Probit	Hyp	-827.78	-827.31	0.47
Logit	Exp	-834.91	-834.34	0.57
Logit	Hyp	-830.45	-830.42	0.03
Quantum	Exp	-830.53	-829.59	0.94
Quantum	Hyp	-828.73	-828.67	0.06
Quantum Rotation	Rotation	-836.53	-835.46	1.07

analytical version of DFT. This implies that some of the benefit of an initial parameter is averaged out by an analytical DFT model.

4.5. Impact of information update rate

Whilst all of the results thus far have only used observations at 1 second intervals before the moment where the driver chooses to merge, we also consider models applied to the data where more observations are included. This results in different frequencies of 'visual information updating speeds' for the decision-makers (i.e. different rates for how often the attributes of alternatives are updated). The key benefit of testing across different rates is that it allows us to test how the estimate for the number of preference updating steps in a DFT model is impacted by the information updating rate. This is a result of the fact that under our implementation, it is assumed that an individual considers each set of attribute values for the same number of preference updating steps. We estimate full (analytical) DFT models on datasets that include 1, 2, 5 and 10 data updates per second. The results of these models are given in Table 7.

TABLE 7 : Results from DFT models applied with different visual update rates

Data Updates per second	1	2	5	10
Observations	3,293	6,391	15,718	31,245
Log-likelihood	-831.75	-1,110.15	-1,473.87	-1,751.41
t estimate:	347	202	114	43

Crucially, the estimate for the number of preference updating steps per visual update decreases as the of data updating steps increases, resulting in approximately the same total of preference updating steps per second. A visual representation of this result is given in Figure 11.

Additionally, we also compare the results of our final models across the dataset with two different assumptions for the visual update rate. Table 8 gives the log-likelihood of models where

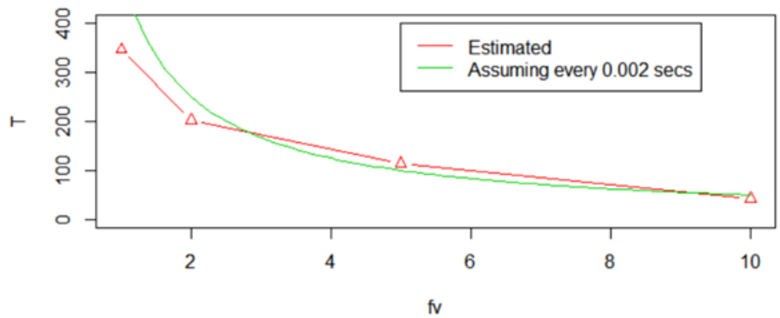

FIGURE 11 : Relationship between the frequency of visual (fv) and mental updating steps (T)

0.1 and 1 second intervals are used for data update rates (with exponential discount rates used for the ‘static’ models).

TABLE 8 : Performance of the different models for different data update frequencies

		0.1s interval		1s interval	
	pars	LL	BIC	LL	BIC
Logit	12	-1,768.58	3,661.36	-830.42	1,785.04
Probit	12	-1,753.21	3,630.61	-831.39	1,786.98
DFT	12	-1,750.80	3,625.79	-824.57	1,773.34
Quantum	12	-1,759.13	3,642.46	-829.59	1,783.38
Quantum Rotation	14	-1,756.04	3,656.98	-835.46	1,815.81

The relative performance of the models changes depending on which update rate is used. Whilst the quantum rotation model has the worse fit than the quantum model for a 1 second interval data, it has better fit for 0.1 second interval data. Similarly, probit performs relatively better for 0.1 second data, obtaining the second best model fit. Across both datasets, however, DFT is the best performing model.

5. CONCLUSIONS AND FUTURE STEPS

This work provides a first step towards bridging the gap between mathematical psychology and choice modelling for dynamic data, for which the attributes for alternatives change over time. In particular, we focus on the implementation of a dynamic decision field theory model for such data. As DFT models already attempt to capture the deliberation process, the inclusion of attribute changes whilst the decision-maker deliberates on their alternatives does not imply fundamental changes to the mathematical structure of the model, but has never before been implemented, as far as the authors are aware. In the context of driving decisions, we find that a dynamic DFT model can capture the choice of whether to merge lanes or not. In addition, a simulated DFT model can account for initial preferences, which do not appear to be picked up by alternative model structures.

The work conducted in this paper, however, is just a first step, as there are many more features that a model for dynamic data could include. Given that the data tested here is based on video recorded data, we have not included any sociodemographics or individual-specific attributes. It is easily possible, for example, that a driver's reaction times and risk propensity could have a significant impact on the choice of whether to accept a gap or not. Additionally, an individual's speed at which they process new information will likely impact how they perceive the new information, with previous results suggesting that there is a strong delay effect on the processing of changing information (Holmes et al., 2016). Furthermore, the separation of processing and action times will likely impact models. There is however, a clear additional advantage of a simulated DFT model, in that the addition of such features to control for heterogeneous driving behaviour can easily be accommodated, as adding distributions for new parameters will not result in longer model runtimes. These further steps would help create a *truly* dynamic model, as at present DFT operates similarly to the 'static' models in that it is reliant on the decay parameter α_d.

Additionally, a decision field theory model for changing attributes could also be compared to alternative model structures based on alternative accumulator models. For example, the linear ballistic accumulator model could easily be adapted to include drift rates that are applied for some amount of time before they are updated when the attributes change.

Furthermore, whilst the emphasis of this paper is not on the driving behaviour itself, future work could compare the models developed in this paper with models traditionally used for gap acceptance and lane merging tasks, for which a large array of models have previously been specifically developed (Brilon et al., 1999). Additionally, further tests of the models developed in this paper should consider transferability to other driving behaviour datasets, for which logit models already perform well (Rossi et al., 2013). A number of other factors that influence merging behaviour could also be added, with, for example, neither vehicle accelerations nor traffic conditions taken into account in the models specified in this paper. The work in this paper also assumes that the choice to merge or not is a binary decision. An alternative approach would be to attempt to capture the distance travelled in the direction of the target lane at each time interval, as the action of merging can take a few seconds and a driver may start merging before changing their mind. Accumulator models accounting for changing attributes may also look very different for differing choice contexts. In the dataset in this paper, the pattern of choices is always the same (the driver 'chooses' not to merge for some number of time intervals before 'choosing' to merge in the final time interval). Thus the key impact of the decay parameters in these models may simply be accounting for this fact. More interesting behavioural insights may be generated by a more 'random' pattern of choices, for which the impact of decay parameters may be very different. The dynamic logit, probit and quantum models in this paper also rely on the assumption that there is no correlation in the unobserved factors over time. Whilst this is an approach that has previously been applied to a probit model with lagged variables (Papatla and Krishnamurthi, 1992), and this approach may be suitable for larger time intervals, it becomes less likely that this assumption is necessarily realistic nor valid for shorter time intervals of, for example, 0.1 seconds. Further work should consider methods for treating possible correlations of unobserved factors over time.

Overall, the work in this paper suggests that accumulator models that account for changing attributes could provide a useful tool for the study of rapidly changing choice contexts, demonstrating that there is clearly extensive scope for future developments of such models.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support by the European Research Council through the consolidator grant 615596-DECISIONS.

REFERENCES

Aguirregabiria, V. and Mira, P. (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics*, 156(1):38–67.

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.

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

Bjørner, T. B. and Leth-Petersen, S. (2005). Dynamic models of car ownership at the household level. *International Journal of Transport Economics/Rivista internazionale di economia dei trasporti*, pages 57–75.

Boxall, P., Adamowicz, W. L., and Moon, A. (2009). Complexity in choice experiments: choice of the status quo alternative and implications for welfare measurement. *Australian Journal of Agricultural and Resource Economics*, 53(4):503–519.

Brilon, W., Koenig, R., and Troutbeck, R. J. (1999). Useful estimation procedures for critical gaps. *Transportation Research Part A: Policy and Practice*, 33(3-4):161–186.

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 Bruza, P. D. (2012). *Quantum models of cognition and decision*. Cambridge University Press.

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.

Cambridge Systematics Inc. (2005a). NGSIM US 101 Data Analysis. http://www.ngsim.fhwa.dot.gov.

Cambridge Systematics Inc. (2005b). NGSIM US 101 Data Analysis: Summary Report. *Prepared for Federal Highway Administration*.

Choudhury, C. F. (2007). *Modeling driving decisions with latent plans*. PhD thesis, Massachusetts Institute of Technology.

Cirillo, C. and Xu, R. (2011). Dynamic discrete choice models for transportation. *Transport Reviews*, 31(4):473–494.

Diederich, A. and Busemeyer, J. R. (2003). Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *Journal of Mathematical Psychology*, 47(3):304–322.

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11):3612–3628.

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.

Erdem, T. (1996). A dynamic analysis of market structure based on panel data. *Marketing science*, 15(4):359–378.

Fang, H. and Wang, Y. (2015). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review*, 56(2):565–596.

Fosgerau, M., Frejinger, E., and Karlstrom, A. (2013). A link based network route choice model with unrestricted choice set. *Transportation Research Part B: Methodological*, 56:70–80.

Fu, C., Zhang, Y., Bie, Y., and Hu, L. (2016). Comparative analysis of driverâĂŹs brake perception-reaction time at signalized intersections with and without countdown timer using parametric duration models. *Accident Analysis & Prevention*, 95:448–460.

González-Valdés, F. and de Dios Ortúzar, J. (2018). The stochastic satisficing model: A bounded rationality discrete choice model. *Journal of choice modelling*, 27:74–87.

Green, L., Fry, A. F., and Myerson, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological science*, 5(1):33–36.

Green, L. and Myerson, J. (1996). Exponential versus hyperbolic discounting of delayed outcomes: Risk and waiting time. *American Zoologist*, 36(4):496–505.

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.

Hancock, T. O., Hess, S., and Choudhury, C. F. (2019a). An accumulation of preference: two alternative dynamic models for understanding transport choices. *Submitted*.

Hancock, T. O., Hess, S., and Choudhury, C. F. (2019b). A careful respondent or an uncertain response: disentangling confounding sources of increased deliberation time using decision field theory. *submitted*.

Hancock, T. O., Hess, S., and Choudhury, C. F. (2019c). Quantum probability: A new method for modelling travel choices. *submitted*.

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.

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

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

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.

Kim, J. B., Albuquerque, P., and Bronnenberg, B. J. (2016). The probit choice model under sequential search with an application to online retailing. *Management Science*, 63(11):3911–3929.

Kirby, K. N. and Maraković, N. N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human decision processes*, 64(1):22–30.

Kitamura, R. and Bunch, D. S. (1990). Heterogeneity and state dependence in household car ownership: A panel analysis using ordered-response probit models with error components.

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.

Lipovetsky, S. (2018). Quantum paradigm of probability amplitude and complex utility in entangled discrete choice modeling. *Journal of choice modelling*, 27:62–73.

Madden, G. J., Bickel, W. K., and Jacobs, E. A. (1999). Discounting of delayed rewards in opioid-dependent outpatients: exponential or hyperbolic discounting functions? *Experimental and clinical psychopharmacology*, 7(3):284.

Mai, T., Fosgerau, M., and Frejinger, E. (2015). A nested recursive logit model for route choice analysis. *Transportation Research Part B: Methodological*, 75:100–112.

Marczak, F., Daamen, W., and Buisson, C. (2013). Key variables of merging behaviour: empirical comparison between two sites and assessment of gap acceptance theory. *Procedia-Social and Behavioral Sciences*, 80:678–697.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Commons, ML.; Mazur, JE.; Nevin, JA*, pages 55–73.

Noguchi, T. and Stewart, N. (2018). Multialternative decision by sampling: A model of decision making constrained by process data. *Psychological review*, 125(4):512.

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.

Papatla, P. and Krishnamurthi, L. (1992). A probit model of choice dynamics. *Marketing Science*, 11(2):189–206.

Paschalidis, E., Choudhury, C. F., and Hess, S. (2018). Modelling the effects of stress on gap-acceptance decisions combining data from driving simulator and physiological sensors. *Transportation research part F: traffic psychology and behaviour*, 59:418–435.

Paschalidis, E., Choudhury, C. F., and Hess, S. (2019). Combining driving simulator and physiological sensor data in a latent variable model to incorporate the effect of stress in car-following behaviour. *Analytic Methods in Accident Research*, 22:100089.

Polus, A., Shiftan, Y., and Shmueli-Lazar, S. (2005). Evaluation of the waiting-time effect on critical gaps at roundabouts by a logit model. *European Journal of Transport and Infrastructure Research*, 5(1):1–12.

Ratcliff, R. and Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human perception and performance*, 26(1):127.

Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of risk and uncertainty*, 23(1):5–32.

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.

Rossi, R., Meneguzzer, C., and Gastaldi, M. (2013). Transfer and updating of logit models of gap-acceptance and their operational implications. *Transportation Research Part C: Emerging Technologies*, 28:142–154.

Rubinstein, A. (2003). âĂIJeconomics and psychologyâĂİ? the case of hyperbolic discounting. *International Economic Review*, 44(4):1207–1216.

Rust, J. (1994). Structural estimation of markov decision processes. *Handbook of econometrics*, 4:3081–3143.

Rysman, M., Gowrisankaran, G., et al. (2009). *Dynamics of Consumer Demand for New Durable Goods*. National Bureau of Economic Research.

Schiraldi, P. (2011). Automobile replacement: a dynamic structural approach. *The RAND journal of economics*, 42(2):266–291.

Seetharaman, P. (2004). Modeling multiple sources of state dependence in random utility models: A distributed lag approach. *Marketing Science*, 23(2):263–271.

Swait, J., Adamowicz, W., and van Bueren, M. (2004). Choice and temporal welfare impacts: incorporating history into discrete choice models. *Journal of Environmental Economics and Management*, 47(1):94–116.

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. and Busemeyer, J. R. (2011). A quantum probability account of order effects in inference. *Cognitive science*, 35(8):1518–1552.

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

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

Voskuilen, C., Ratcliff, R., and Smith, P. L. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. *Journal of Mathematical Psychology*, 73:59–79.

Wang, Z., Solloway, T., Shiffrin, R. M., and Busemeyer, J. R. (2014). Context effects produced by question orders reveal quantum nature of human judgments. *Proceedings of the National Academy of Sciences*, 111(26):9431–9436.

Wood, J. and Zhang, S. (2017). Evaluating relationships between perception-reaction times, emergency deceleration rates, and crash outcomes using naturalistic driving data. Technical report, Mountain-Plains Consortium, North Dakota State University, Fargo, ND.

Zhang, S., Lee, M. D., Vandekerckhove, J., Maris, G., and Wagenmakers, E.-J. (2014). Time-varying boundaries for diffusion models of decision making and response time. *Frontiers in Psychology*, 5:1364.

Zhao, G., Wu, C., and Ou, B. (2011). Mathematical modeling of average driver speed control with the integration of queuing network-model human processor and rule-based decision field theory. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 55, pages 856–860. Sage Publications Sage CA: Los Angeles, CA.