A Dual-Process Diffusion Model

Carlos Alós-Ferrer

Forthcoming in *Journal of Behavioral Decision Making* (2016)

A Dual-Process Diffusion Model

Carlos Alós-Ferrer
Department of Economics, University of Cologne, Albertus-Magnus Platz, D 50923 Cologne, Germany

This paper presents a simple formal-analytical model delivering qualitative predictions for response times in binary-choice experiments. It combines a dual-process/multi-strategy approach with the standard diffusion model, modeling a utility decision process and a heuristic decision process as diffusion processes of evidence accumulation. For experiments with objective alternatives (including many tasks in judgment and decision making), the model predicts that errors will be quicker than correct responses in case of process conflict and slower in case of alignment, capturing a well-documented asymmetry regarding slow or fast errors. Further, the model also predicts that correct responses are slower in case of conflict than in case of alignment, capturing the well-known Stroop effect. The model is also extended to cover experiments with subjective alternative evaluations, i.e., preferential choice. In this case, results depend on whether trials are hard or easy, i.e., on whether the heuristic can be interpreted as relatively automatic or not.

Keywords: dual processes, diffusion model, response times.

Response times are a central tool for the study of psychological processes underlying human decisions. They are especially important in two-alternative choice experiments, which in turn are ubiquitous in several branches of psychology. Many formal-analytical models have been developed to account for regularities in the data, one of the most prominent thereof being the diffusion model of evidence accumulation (Ratcliff, 1978, 2013; Ratcliff & Rouder, 1998; Smith, 2000). This model has harvested many successes, chief among them being able to account for the speed-accuracy trade-off (Wickelgren, 1977). However, the mathematical complexity of diffusion processes is relatively high, and analytical results are scarce, leading many researches to rely on simulations for model evaluation.

This paper proposes a parsimonious formal-analytical model for response times in binary-choice decisions, the Dual-Process Diffusion Model or DPDM. My objective is not to develop a parametric fit for individual decision makers, but rather to obtain analytically founded, easily testable predictions on conditional, average response times. The target are experimental paradigms in which decisions can be usefully characterized as those arising from either a "conflict" or an "alignment" of decision criteria, strategies, or processes. A simple example is the Stroop Task (Stroop, 1935), where incongruent trials correspond to a choice where naming the word and naming the color produce different responses, and congruent trials are those where both responses are aligned. Another example from the consumer choice literature is the decision to buy consumer items which differ along two main attributes, e.g., brand and price. Analogously, in judgment and decision-making experiments where a particular heuristic or decision bias is targeted, the heuristic can either prescribe a response contrary to the normative one (a conflict), or one coinciding with it (alignment). Examples of the latter kind include the interaction between intuitive and normative responses (e.g., the conflict in the LINDA problem, Tversky & Kahneman, 1983) and more complex interactions between behavioral rules. For instance, a reinforcement heuristic might prescribe to repeat whatever has been successful in the past, which, depending on the problem at hand, might or might not be normatively optimal (Achtziger & Alós-Ferrer, 2014).

The simplest way to capture the idea of conflict and alignment is to postulate two different decision processes relying

Forthcoming in **Journal of Behavioral Decision Making**. This Revision: January 2016.

The author thanks Johannes Buckenmaier, Michael Dambacher, Holger Gerhardt, Andreas Gloeckner, Georg Granić, Hauke Heekeren, Johannes Kern, David Laibson, Julian Marewski, three anonymous referees, and participants in the Symposium on Interdisciplinary Perspectives on Decision Making at the Wissenschaftszentrum Berlin (WZB) and the First Symposium in Motivation and Self-Control at the University of Cologne, for helpful and stimulating comments, both on the paper and the overall research agenda. Financial support from the Research Unit "Psychoeconomics" (FOR 1882), financed by the German Research Foundation (DFG), is gratefully acknowledged.

Correspondence concerning this article should be addressed to Carlos Alós-Ferrer, Department of Economics, University of Cologne, Albertus-Magnus Platz, D 50923 Cologne, Germany. Phone: +49 2214708303, Fax: +49 2214704119, E-mail: carlos.alos-ferrer@uni-koeln.de

on different cues, analogously to the way dual-process models (e.g., Alós-Ferrer & Strack, 2014; Evans, 2008; Sloman, 1996; Strack & Deutsch, 2004) are built. An alternative interpretation is to define conflict and alignment with respect to the prescriptions of multiple strategies in a multi-attribute decision making context, where each decision strategy relies on a different set of attributes (e.g., Söllner, Bröder, Glöckner, & Betsch, 2014; Weber & Johnson, 2009).

The approach is also conceptually related to Rottenstreich and Hsee (2001) and Hsee and Rottenstreich (2004) (see also Hsee, Rottenstreich, & Xiao, 2005). These authors argued that decisions are often the result of two different kinds of valuations. The first, *valuation by calculation*, essentially corresponds to the economic concepts of preference, utility, or willingness to pay. The second, *valuation by feeling*, captures more affective predispositions for the alternatives. In the spirit of this distinction, and for the sake of concreteness, the model will be formulated in terms of two decision processes, the *utility process* and the *heuristic process*. The former is meant to capture the more computational-normative aspects of decision making, while the latter broadly corresponds to the evaluation of intuitive-affective attributes. In both cases, the word "process" is used for concreteness, but can be taken to mean either "strategy" or "criterion."

The model is built on the premise that the actual decision arises from the interplay of the utility process and the heuristic process. Each of the processes, however, is treated as a mathematical diffusion process as in Ratcliff's (1978) model. The model hence combines foundations from different strands of the literature and partakes from their respective strengths. The structure, however, is simple enough to allow deriving analytical predictions which translate directly into empirically testable hypotheses.

The architecture of the DPDM places it within the category of multi-strategy approaches to decision making (Marewski & Link, 2014), in the sense that decisions are ultimately made by one of two strategies, captured by the heuristic and utility processes. However, the modeling of each of the two processes relies on evidence accumulation models (Ratcliff, 1978) which are commonly viewed as an example of single-strategy approaches, in the sense that they postulate a domain-general mechanism. In this sense, the DPDM is a parsimonious multi-strategy approach. Alternatively, the DPDM can be seen as a minimalistic extension of a prominent single-strategy framework to accommodate the interplay of processes which defines dual-process theories.

The objective of this work is not to build a detailed, flexible model capable of fitting experimental data from every conceivable paradigm. Rather, the aim is to derive general, testable predictions independent of parameter fitting and of the specific processes selected. This point is important, because a generalized criticism of multi-strategy approaches is that, by tailoring the strategies to the experimental task, the general approach might become excessively flexible and hard to refute (Glöckner & Betsch, 2011; Marewski & Link, 2014). On the other hand, a typical criticism of the single-strategy approach is that such models come with a large number of free parameters which can be adjusted by the researcher to mimic different, conceivable strategies (Marewski & Link, 2014). In contrast, the objective of this paper is to derive *qualitative* predictions which hold independently of the exact values of the parameters which are part of the model. This way, the DPDM remains falsifiable from the onset.

The basic version of the model targets *decisions from inference*. Those correspond to experiments with fixed alternatives that can be objectively evaluated and where there is a clear notion of what an error is. Examples include tasks in the domain of probability judgment, where individuals should rely on objectifiable attributes, and many tasks in cognitive psychology. The model is used to obtain two different kind of predictions. The first concerns the relative speeds of errors and correct responses, where an asymmetric prediction is obtained. In case of conflict, correct responses will be slower than errors. However, in case of alignment, errors will be slower than correct responses. Both predictions are already supported by empirical evidence from the literature. The second kind of predictions concerns the relative speeds of correct responses in case of alignment and in case of conflict. The prediction of the model in this case captures the well-known Stroop effect: correct responses are slower in case of conflict (incongruent tendencies) than in case of alignment (congruent tendencies).

The second, more elaborated version of the model targets *preferential choice*. This corresponds to experiments with variable alternatives, where individuals' choices rely on subjective attributes. Examples include tasks from the literature on attitudes and attitude change, and many examples in economic decision making, e.g., consumer choice in the presence of multiple salient attributes or the extensive literature on lottery choice. In these experiments, each participant typically faces a large collection of binary choices, which might uncover an underlying heuristic conflicting with a more normative (but still subjective) evaluation. Trials can then be classified in those where the heuristic poses a serious challenge by virtue of delivering an intuitive response (hard trials) and those where it does not (easy trials). For hard trials, the predictions from the basic model can be extended to this setting in a mathematically straightforward way. For easy trials, however, errors will take longer both in case of alignment and in case of conflict.

Each version of the model is presented in a different section below, with separate sections addressing the predictions for the speed of errors and the Stroop effect. A further section before the discussion presents an example applying the (second) model to a consumer-choice paradigm. The discussion briefly addresses the relation to other models in the literature, possible generalizations, and avenues for future research.

Fixed Alternatives: A Model of Inference

Many experiments in decision making involve two abstract choices, with participants' payment (if at all) depending on their estimation of which of them is correct. In many probability inference tasks, participants are asked to choose the correct answer to a given question. Examples include the classical questions used to demonstrate the representativeness heuristic or, more generally, base-rate neglect (De Neys & Glumicic, 2008; De Neys, Vartanian, & Goel, 2008; Kahneman & Tversky, 1972, 1973), including the celebrated lawyers-engineers problem. Further examples include the LINDA problem used to study the conjunction fallacy (Tversky & Kahneman, 1983), or even the questions part of the Cognitive Reflection Test (Frederick, 2005). Experiments in the field of probability biases and belief updating often employ designs where risk or uncertainty is represented by urns from which balls are extracted, but the prior information is varied from decision to decision. Examples include designs capturing both conservativeness and the representativeness heuristic (Achtziger, Alós-Ferrer, Hügelschäfer, & Steinhauser, 2014; Grether, 1980, 1992), and experiments studying biases derived from reinforcement (Achtziger & Alós-Ferrer, 2014; Charness & Levin, 2005), among many others. Beyond judgment and decision making, the description also fits many paradigms in the fields of attention and cognition, starting with the flanker task (Eriksen & Eriksen, 1974) and the Stroop task (MacCleod, 1991; Stroop, 1935), if we consider the named color and the actual color as the two options. In all these experiments, the choices themselves are *fixed* and *neutral* (objective), in the sense that the framing remains constant, and the participants' task does not involve a subjective evaluation of the choices themselves. Often, the target of the experiment is a specific bias or heuristic, to be studied against the benchmark of normatively rational behavior.

This section presents the simplest version of the dual-process diffusion model (DPDM), which is tailored to situations as described above. There are two fixed alternatives, A and B, and two decision processes, called the *utility process* (U) and the *heuristic process* (H) for concreteness. These processes are modeled through mathematical diffusion processes, following sequential sampling models form cognitive psychology (Ratcliff & Smith, 2004). That is, each of them is captured by a Brownian motion with drift $X(t)$,

$$dX(t) = \mu dt + \sigma dW(t),$$

where μ and $\sigma > 0$ are the drift and diffusion parameters, respectively, and $dW(t)$ is the increment of a standard Brownian motion. The process is endowed with two barriers ℓ, h with $\ell < 0 < h$. The interval $]\ell, h[$ is called the *corridor* of the diffusion process. The decision process starts at $X(0) = 0$ and selects A or B if the upper or lower barrier is hit first, respectively. The idea is that evidence gradually accumulates in favor of one or the other option, with the drift describing both the rough direction and the swiftness of the process, and the diffusion parameter describing an inherent randomness.

The key of the model is the concept of *favored response.* In the study of heuristics and biases, and also in the study of perception and cognition, it is usually easy to associate a typical response with each conceptualized decision process. For example, for the Stroop task, let the utility process capture the full-attention ideal. The favored response of this process is the actual color. The heuristic process, however, has the named color as the favored response, creating a conflict in incongruent trials. In tasks capturing the representativeness heuristic, the heuristic process favors the stereotypical response, while the utility process favors the normatively correct option derived from Bayes' rule.

Within the context of a diffusion process, say that A (resp. B) is the favored response of the process if $\mu > 0$ (resp. $\mu < 0$). That is, the favored response is that reflecting the default trend of the process, although due to the stochastic nature of evidence accumulation, the actual response can differ from the favored one. By renaming alternatives (or simply exchanging their position), however, we can always focus on the case $\mu > 0$. From now on, formally all diffusion processes in the model are endowed with a positive trend. This does not mean, however, that they all favor the same response. For instance, if process H favors option A but process U favors option B, we will have parameters $\mu_H > 0$ and $\mu_U > 0$, but process H is understood to have an upper barrier leading to A and process U an upper barrier leading to B.

To provide increasingly better fits to the data, models building on diffusion processes as above have studied asymmetric barriers (Ratcliff, 1978) and enriched the basic model by, e.g., allowing for random starting points (Ratcliff, 1978, 1981; Ratcliff & Rouder, 1998). The model presented here, however, explains empirical regularities without the recourse to such extensions. Hence, I will concentrate on the *unbiased* case $\ell = -h$ and normalize $h = 1$.

The basic assumption in the literature of dual processes is that there are two types of processes, automatic and controlled. Automatic processes are quicker and reflect quick associations, while controlled processes are slower and partially reflect deliberation or cognition. In the terms of the model, we can simply think that automatic processes are those which, by virtue of relying on clear associations, have a stronger tendency to select the favored response, or, in the terms above, a larger drift. The assumption of the model is that the heuristic process H is endowed with a drift μ_H and the utility process with a drift μ_U, with $\mu_H > \mu_U > 0$ (and both have the same diffusion parameter σ). As will be seen below (Corollary 1), this assumption immediately implies that the heuristic (automatic) process is faster than the utility (controlled) process, and also behaves more in a stimulus-response way, selecting its own favored response more often than the utility process. In other words, the heuristic process is "swifter."

The simplification that positive drift is always towards the favored response allows using μ_H, μ_U directly as the measure of swiftness of the process. However, the processes can still share the same favored response or deliver opposed prescriptions. An *alignment situation* occurs if the response favored by the heuristic process is the same as the one favored by the utility process, and a *conflict situation* occurs if the favored responses are different. For example, in the Stroop task, incongruent trials where the named color is printed in a different color are conflict situations, and congruent trials where the actual color matches the named one are alignment situations. The situation posed by many inference questions used to study decision biases is one of conflict, as, e.g., in the case of the LINDA problem.

The response favored by the utility process is, by definition, normatively correct. Accordingly, decisions where this option is selected will be referred to as *correct* decisions, and decisions where the other response is selected will be called *errors*. In general, however, the names "correct decision" and "error" should be taken merely as a convenient terminology, since the analysis also applies for situations where the response of the heuristic process is, for whatever the reason, considered to be desirable.

To complete the model, we need to detail the interaction of both processes. I will adopt a simple, probabilistic approach. Assume that, with probability $0 < \Delta < 1$, the decision follows the heuristic process and with probability $1 - \Delta$ it follows the utility process. It is natural to assume that Δ might be influenced by factors as mood, cognitive load, monetary incentives, etc. However, the properties proven below are independent of the exact value of Δ, and hence this parameter will not be fleshed out any further here.

A standard prediction of dual-process models is that total response time in case of conflict is increased due to the existence of the conflict itself. This is in agreement with widespread evidence from cognitive psychology and neuroscience indicating that the detection of a decision conflict (by the central executive) is time-consuming.[1] This additional factor plays no role if one compares only response times conditional on conflict or alignment, as in Theorem 1 below. However, this possibility needs to be explicitly taken into account if one compares response times in case of conflict and those in case of alignment, as in Theorem 2 below. It is natural to capture this element by specifying an additional *conflict detection time* D_i which is weakly larger for conflict situations ($i = 1$) than for alignment situations ($i = 0$), i.e., $D_1 \geq D_0$. The particular case $D_1 = D_0$ is equivalent to neglecting implications from conflict detection time.

Interpretation of the Basic Model

Figure 1 illustrates the basic workings of the DPDM. A mechanistic, direct interpretation is as follows. For any given decision problem, the heuristic process and the utility process run in parallel, each of them following a diffusion process with its own drift parameter. Simultaneously, a central executive process selects one of the two processes: the heuristic process with probability Δ, and the utility process with the remaining probability. The decision of the selected process is implemented as soon as that process delivers a response. In case of alignment, both processes favor the same, correct response, that is, both upper boundaries correspond to a correct decision. In case of conflict, the upper boundary of the heuristic process corresponds to an error, while the upper boundary of the utility process corresponds to a correct response.

This interpretation is close to the recent model of Pleskac and Wershbale (2014) for the Balloon Analogue Risk Task (see the Discussion section for details). Of course, this mechanistic reading of the model has to be seen as a simplification. For instance, the response time is the response time of the selected process, without direct or indirect interference of the other process. This means that process selection time is assumed either to occur before the actual process time of to be short enough not to interfere with it. A more nuanced version of the model could have process selection depending on initial, early signals extracted from the two processes. However, it is easy to see that a "race model" where the first process to finish is selected would have low empirical validity, since this would imply that a majority of our decisions follow heuristic processes (which, fortunately, seems not to be the case for most human decision makers).

The model can also be reframed in terms of multi-attribute decision making as follows. Suppose there are two options which differ along two attributes, e.g., price and branding. The attributes have an attached evaluation scale, e.g., some prices are lower than others and some brands are more well-known than others. Hence, the attributes can either point towards the same option (alignment) or towards different options (conflict). Assume one of the attributes captures the relevant evaluation (e.g., select the cheapest option), while the other is a heuristic/automatic evaluation. The relevant evaluation always points towards the correct response (although, being a stochastic process, it might still generate an error). In case of conflict, the heuristic evaluation points towards an error, but, in case of alignment, it points towards the correct response. This interpretation brings the model closer to the

[1] For example, a brain imaging study by De Neys et al. (2008) found ACC activation in case of conflict in the lawyer-engineer problem of Kahneman and Tversky (1972), independently of whether the heuristic response was adopted or not, indicating that (time-consuming) conflict detection is independent of the process and response ultimately selected.

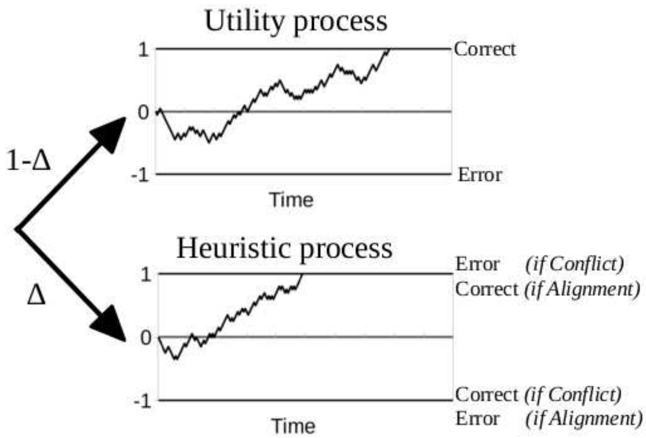

Figure 1. Schematic representation of the model.

multi-attribute model of Diederich (1997) (see the Discussion section).

Dual-Process Properties

Given a single decision process as those described above (i.e., either the heuristic or the utility process), two quantities are of interest. The first is the probability that the favored response is selected, P. The complementary probability $1 - P$ will be called the *fault probability*, i.e., the probability that the process selects the non-favored response. The second quantity of interest is the expected time until the process leaves the corridor, T, i.,e. the (process) expected response time.

Explicit formulae for P and T are known (and become relatively simple in the symmetric case) and have been derived in the literature (see, e.g., Grasman, Wagenmakers, & van der Maas, 2009; Palmer, Huk, & Shadlen, 2005)[2] using standard techniques from stochastic calculus (see Karatzas & Shreve, 1998). The probability that the process selects its favored option is given by

$$P(\mu) = \frac{1}{1 + e^{-\frac{2\mu}{\sigma^2}}} > \frac{1}{2} \qquad (1)$$

and the expected response time is given by

$$T(\mu) = \frac{1}{\mu}(2P - 1). \qquad (2)$$

with $T(0) = \lim_{\mu \to 0} T(\mu) = 1/(\sigma^2)$. A surprising but very convenient property of the diffusion model in the symmetric case is that the process response time conditional on either response is identical (and hence equal to T; see Palmer et al., 2005, Appendix). This fact will greatly simplify the computations in the sequel since one does not need to consider conditional response times *for a given process*.

Fix the diffusion parameter σ and consider a family of processes indexed by μ. Given a process as described, the following comparative statics result is straightforward (and well-known).

Lemma 1. *For a diffusion process with corridor* $]-1, 1[$, *both the fault probability and the response time are strictly decreasing in* μ.

Proof. The expression of $P(\mu)$ above shows that this quantity is strictly increasing in μ, hence $1 - P(\mu)$ is strictly decreasing in μ. It will be shown that the derivative of $T(\mu)$ is strictly negative, proving that $T(\mu)$ is strictly decreasing. A direct computation shows that

$$P'(\mu) = \frac{2}{\sigma^2} e^{-\frac{2\mu}{\sigma^2}} \left(1 + e^{-\frac{2\mu}{\sigma^2}}\right)^{-2}$$

and

$$T'(\mu) = \frac{1}{\mu^2}\left[2P'(\mu)\mu - (2P(\mu) - 1)\right]$$

Hence, $T'(\mu) < 0 \iff 1 + 2P'(\mu)\mu < 2P(\mu)$. Writing $z = \frac{2\mu}{\sigma^2}$, a computation shows that this condition holds if and only if

$$\left(1 + e^{-z}\right)^2 + 2ze^{-z} < 2\left(1 + e^{-z}\right)$$

which simplifies to $2z < e^z - e^{-z}$. This last inequality holds for any $z > 0$, as can be easily established taking a Taylor expansion of the function $f(z) = e^z - e^{-z}$ around 0. Hence, $T'(\mu) < 0$ and it follows that T is strictly decreasing in μ. ■

The intuition of this result is simple. A process with a larger drift parameter tends towards the favored response in a swifter way, resulting in increased speed and also on a lower probability of going astray. This fact is the key for the interpretation of the model in dual-process terms.

Applying Lemma 1, we immediately obtain the following result.

Corollary 1. *Let a heuristic and a utility process with* $\mu_H > \mu_U > 0$, *common* σ, *and corridor* $]-1, 1[$. *Then,*

(a) $T_H < T_U$, *that is, the expected response time of the heuristic process is shorter than that of the utility process, and*

(b) $(1 - P_H) < (1 - P_U)$, *that is, the fault probability of the heuristic process is lower than that of the utility process.*

[2]Interestingly, these processes have also been considered in finance for the analysis of *double-barrier options*, and the same formulae have been independently derived there (see, e.g., Douady, 1998).

The properties identified in the last corollary are both intuitive and natural within a dual-process approach. They justify the interpretation of the heuristic process as more automatic and the utility process as more controlled. The heuristic process is relatively fast and has a relatively low fault probability, and can hence be interpreted as being closer to a stimulus-response mapping. In contrast, the controlled process is slower and has a larger fault probability, corresponding to the idea of a deliberative process.

It should be remarked at this point that part (b) in the last corollary does not mean that the probability of an error is smaller for the heuristic process. This is not true, because whether the favored response of the heuristic process is correct or an error depends on whether the situation is one of conflict or of alignment. In case of conflict, the heuristic process has a larger probability of resulting in an error, because the process' favored response is an error. In case of alignment, however, the heuristic process is a quick and efficient process which delivers the correct response more often than the utility process.

Remark 1. The DPDM aims to build a framework which is as simple as possible. Hence, a number of structural assumptions are made. For instance, the model abstracts from many elements which add flexibility to evidence-accumulation models. Specifically, it assumes no starting-point variability and no starting bias. Also, it assumes symmetric barriers, $\ell = -h$ (the normalization $h = 1$ is of course inconsequential and just for convenience). More importantly, it assumes that both processes have the same barrier. The model can be enriched to accommodate such possibilities, but it is my position that the predictions of the basic, simpler model should be understood first. Hence, this task is left for future research.

However, it should be pointed out that the qualitative predictions of the model are not very sensitive to small variations in the approach. While the formulation of the DPDM in terms of two diffusion processes with different drift rates is natural, the predictions derived below for response times depend only on the two properties identified in Corollary 1. That is, the predictions discussed here still obtain if one replaces the modeling of the heuristic and utility process by a different formulation such that those two properties still hold. For instance, consider the more general case with symmetric but different barriers for the two processes, that is, let the heuristic and the utility processes have corridors $]-h_H, h_H[$ and $]-h_U, h_U[$, respectively. An examination of the analytical formulae for fault probabilities and expected response times in this case shows that the predictions hold if the heuristic process has a weakly larger drift rate and a weakly lower barrier than then utility process ($\mu_H \geq \mu_U$ and $h_H \leq h_U$), as long as $\mu_H \cdot h_H > \mu_U \cdot h_U$. The predictions would not hold if, e.g., $\mu_H = \mu_U$ and $h_H < h_U$, because in this later case the heuristic process would have a larger fault probability than the utility process.

Fast and Slow Errors

There are many experimental paradigms where decisions can be classified into correct responses and errors, and where response times have been analyzed as a tool to better understand the underlying decision processes. The empirical relation between the response times of errors and correct responses varies, often within the same class of paradigms. For instance a classical pattern found in choice detection paradigms is that errors are typically slower than correct responses if discrimination is hard or accuracy is emphasized, but it is not uncommon to find that errors are faster than correct responses if discrimination is easy or speed is emphasized (for a discussion, see Townsend & Ashby, 1983, Chapter 9).

Correctly predicting experimentally observed patterns has proven a difficult task (see, e.g., Ratcliff & Rouder, 1998, for a discussion). In particular, the empirical observation of slow errors for certain tasks was originally considered puzzling. Fitting the diffusion model to this observation was one of the motivations for introducing, e.g., drift variability (Ratcliff & Rouder, 1998). However, the model discussed here delivers this prediction in a natural way while sticking to the simplest (unbiased) case.

The DPDM does of course not aim to organize the whole literature on response times. The intention is more modest. The predictions derived below aim at a particular subset of of paradigms, namely those where two different decision processes or decision criteria can be identified, *and* error rates are high enough to allow for experimental testing. Under these conditions, as we will see below, the DPDM delivers a clear-cut prediction: fast errors arise in case of conflict, and slow errors in case of alignment.

Clear experimental evidence confirming the predictions above has been obtained in the domain of judgment and decision making. For instance, De Neys (2006) examined response times in a conjunction-fallacy problem and found that errors were faster than correct responses. Since the conjunction fallacy corresponds to a conflict situation, this agrees with the prediction above. Achtziger and Alós-Ferrer (2014) conducted a decision-making experiment which endogenously created situations of conflict and of alignment. Participants had to choose one of two covered urns, in which balls of two colors were present; one and only one of the colors resulted in monetary payment. The decision of interest was made after a previous urn choice and ball extraction (with replacement). Hence, a heuristic reinforcement process was activated, of the form "choose the same urn again if I won, switch if not". Crucially, however, the number of balls of each color in each urn depended on an unknown state of the world, in such a way that the previous ball extraction allowed the participant to determine (through Bayesian updating of beliefs) which state was more likely and which urn choice was optimal. In some of the trials, the optimal decision was the same as prescribed by reinforcement (alignment); in other trials, it was the opposite (conflict). The results of the experiment indicated that errors were faster than correct responses in case of conflict and slower than correct responses in case of alignment.

Tasks from the field of attention and cognition, as, e.g., the Stroop task, often fulfill the first criterion mentioned above (with congruent trials generating alignment and incongruent trials corresponding to conflict), but empirically observed error rates are typically too low to detect significant differences, especially in the case of congruent trials. For the flanker task White, Ratcliff, and Starns (2011, Figure 1, Tables 1, 4, and 7) reported mean response times of errors and correct responses in several flanker tasks and observed an asymmetry which agrees with the DPDM predictions: for congruent trials, errors are always faster than correct responses, but for incongruent trials, the opposite pattern is often observed.

We now prove the first main result, which, as announced, establishes the predictions of the model regarding expected response times conditional on conflict or alignment. That is, the predictions refer to the actually observable response times of the dual-process model, as opposed to those of an individual process, which are unobservable.[3]

Theorem 1. *Consider the DPDM for fixed alternatives, with parameters Δ, μ_U, and μ_H with $\mu_H > \mu_U > 0$.*

(a) *In case of alignment, the expected response time of correct responses is smaller than the expected response time of errors.*

(b) *In case of conflict, the expected response time of correct responses is larger than the expected response time of errors.*

Proof. Let Q_H be the probability that the heuristic process selects the *correct* response. That is, in case of alignment we have $Q_H = P_H$, but in case of conflict $Q_H = 1 - P_H$, i.e., the probability that this process selects the correct response is actually the fault probability, because the erroneous response is favored. To simplify notation, let $P_\Delta = \Delta Q_H + (1-\Delta)P_U$. The expected response time conditional on the response being right or wrong is given by

$$T(\text{Correct}) = \frac{1}{P_\Delta}\left(\Delta Q_H T_H + (1-\Delta)P_U T_U\right)$$

$$T(\text{Error}) = \frac{1}{1-P_\Delta}\left(\Delta(1-Q_H)T_H + (1-\Delta)(1-P_U)T_U\right)$$

Then we compute

$$\begin{aligned}
&T(\text{Error}) > T(\text{Correct}) \iff \\
&T_H\Delta\left((1-Q_H)P_\Delta - Q_H(1-P_\Delta)\right) > \\
&\qquad T_U(1-\Delta)\left(P_U(1-P_\Delta) - (1-P_U)P_\Delta\right) \iff \\
&T_H\Delta\left(P_\Delta - Q_H\right) > T_U(1-\Delta)\left(P_U - P_\Delta\right) \iff \\
&T_H\Delta(1-\Delta)\left(P_U - Q_H\right) > T_U(1-\Delta)\Delta\left(P_U - Q_H\right) \iff \\
&T_H\left(P_U - Q_H\right) > T_U\left(P_U - Q_H\right)
\end{aligned}$$

That is, we have that $T(\text{Error}) > T(\text{Correct})$ holds if and only if $T_H(P_U - Q_H) > T_U(P_U - Q_H)$.

In case of alignment, $Q_H = P_H > P_U$ by the last corollary and hence $P_U - Q_H < 0$. We conclude that $T(\text{Error}) > T(\text{Correct})$ if and only if $T_H < T_U$. The latter holds by the last corollary. This proves part (a).

In case of conflict, $Q_H = 1 - P_H < 1 - P_U$ by the last corollary. Since $P_U > \frac{1}{2}$, hence $1 - P_U < P_U$, it follows that $Q_H < P_U$ and hence $P_U - Q_H > 0$. We conclude that $T(\text{Error}) < T(\text{Correct})$ if and only if $T_H < T_U$ (and vice versa). The latter holds again by the last corollary, thus $T(\text{Error}) < T(\text{Correct})$, proving part (b). ■

The intuition of this result is as follows. For the conflict case, the utility process favors the correct response, while the heuristic process favors the incorrect one. Hence most correct answers originate in the utility process, and those are slower since $T_U > T_H$. The conclusion in case of alignment might seem counterintuitive at first glance. However, recall that in this case most responses generated by the heuristic process are actually correct. Since $P_H > P_U$, in this case the heuristic process is a quick, efficient shortcut which selects the incorrect answer with a *smaller* probability than the utility process. It follows that most errors are generated by the (slower) utility process. Hence, response times conditional on errors tend to be longer in case of alignment.

The model of Ratcliff (1978) can also account for differences in the relative speeds of correct responses and errors by allowing the parameter μ to vary across trials (see Ratcliff, 2013 for a recent discussion). In the original model, the drift parameter μ itself is normally distributed. However, as Ratcliff and Rouder (1998, Fig. A) observed in a numerical example, if μ takes only two positive values, the expected time of the favored answer can be shorter than that of the alternative one, fitting the prediction. A single process with two possible positive drifts selected randomly is mathematically equivalent to two randomly selected processes with fixed but different drifts, and hence this conclusion also follows from Theorem 3. This enables a reinterpretation of the model as a (single-process) variant of the standard diffusion model. The

[3]A version of Theorem 1 (which, however, was not formulated in terms of diffusion processes) was also briefly discussed in Achtziger and Alós-Ferrer (2014).

problem with this interpretation is that, in order to simultaneously fit both predictions above, i.e., the interaction of response times for correct and incorrect responses with conflict or alignment, the model would need to specify intertrial parameter variation in such a way that alignment and conflict correspond to randomization among drift terms of the same or different sign, respectively. Hence, it is more natural to interpret the two possible values of the drift parameter as two processes of different nature.

The Stroop Effect

One of the most well-known empirical observations concerning experimental observation of response times is the *Stroop effect* (Stroop, 1935), where (correctly) naming the color of a word occurs more slowly when the word names a different color (incongruent trials) than when it names the same color (congruent trials). This effect is empirically robust and easily reproduced (see, e.g., MacCleod, 1991, for a review). In terms of the DPDM, the Stroop effect translates into the assertion that the total expected time of correct responses in case of conflict should be strictly longer than the total expected time of correct responses in case of alignment. This prediction arises naturally and is proven now.

Theorem 2. *Consider the DPDM for fixed alternatives, with parameters Δ, μ_U, μ_H with $\mu_H > \mu_U > 0$, D_0, and D_1 with $D_1 \geq D_0$.*

The expected response time of correct responses in case of alignment is smaller than the expected response time of correct responses in case of conflict.

Proof. The expected response times (excluding conflict detection times D_i) are

$$T(\text{Correct}|\text{Alignment}) = \frac{\Delta P_H T_H + (1-\Delta)P_U T_U}{\Delta P_H + (1-\Delta)P_U}$$

$$T(\text{Correct}|\text{Conflict}) = \frac{\Delta(1-P_H)T_H + (1-\Delta)P_U T_U}{\Delta(1-P_H) + (1-\Delta)P_U}$$

Note that both expressions are convex combinations of T_H and T_U with different weights. Since $T_H < T_U$ by Corollary 1, the statement follows if the weight of T_H in case of alignment is larger than the weight of T_H in case of conflict, that is,

$$\frac{\Delta P_H}{\Delta P_H + (1-\Delta)P_U} > \frac{\Delta(1-P_H)}{\Delta(1-P_H) + (1-\Delta)P_U}.$$

This later statement holds immediately because the real-valued function $f(x) = \frac{x}{(x+(1-\Delta)P_U)}$ is strictly increasing in x and $P_H > (1-P_H)$ since $P_H > 1/2$.

In case of conflict, the total expected response time is incremented by D_1. In case of alignment, it is incremented by $D_0 \leq D_1$. Hence, the claim follows. ■

The intuition of the result is as follows. Both in case of conflict and in case of alignment, in expected terms the same number of decisions will be made by the utility process, and the same fraction of those will result in correct responses. If one thinks in terms of an experiment, the expected number of decisions coming from the utility process contributing to the sample of response times (correct responses) is identical in both cases. The same is not true for the heuristic process. The same expected number of decisions will be made by the heuristic process in both cases, but in case of alignment most of them (the favored responses) will be correct and hence contribute to the sample of response times, while in case of conflict only the less numerous non-favored responses will contribute to the sample. That is, in the sample (correct responses), the *frequency* of response times that originate from the quicker heuristic process is larger in case of alignment.

It is noteworthy that this result holds even if we do not assume any differences in conflict detection and resolution between conflict and alignment, i.e., if $D_1 = D_0$. That is, the effect result does not hinge on the particular assumption $D_1 \leq D_0$. It might also be reasonable to assume that in case of alignment more decisions are made by the heuristic process, i.e., to disentangle Δ in Δ_1 (for the case of conflict) and Δ_0 (for alignment) with $\Delta_0 \geq \Delta_1$. It is easy to see that the last result also holds under this more general assumption, since essentially this again increases the contribution of the heuristic process to the sample of response times in case of alignment.

Variable Alternatives: A Model of Preferential Choice

This section extends the model to settings involving subjective choices. A large part of the literature in judgment and decision making, overlapping heavily with microeconomics, is concerned with utilitarian choices, that is, choices where there is no objectively correct or wrong response, but rather the choice itself is the object of study. A first example is the consumer choice literature, which frequently examines the influence of different strategies and the relevance of difference item attributes on choices. A particular case is the influence of branding on preference-based choices (Philiastides & Ratcliff, 2013; Thoma & Williams, 2013). A second example is the extensive literature on decisions under risk. The experimental workhorse in this literature, ranging back to early examples as Lichtenstein and Slovic (1971), is the choice among different bets or *lotteries*, where optimal choices depend, e.g., on the participant's attitude towards risk. In a typical experiment, each participant makes a series of binary choices, where every choice involves two different lotteries. Hence, the options are not fixed, but change from trial to trial, and the actual task is to evaluate the alternatives and choose whichever is subjectively best. A third example is the literature on attitudes and post-decisional attitude change (e.g., Ariely & Norton, 2008; Brehm, 1956), which has motivated influential theories as cognitive dissonance (Festinger, 1957) and self-perception theory (Bem, 1967). In experiments as those following Brehm (1956) or recent variants as Alós-Ferrer, Granić, Shi, and Wagner (2012), participants are presented with a series of alternatives, say, artistic paintings or hypothetical holiday destinations, and are asked to provide both evaluations (ratings or rankings) and actual binary choices.

The model in the previous sections is now extended to settings of this kind. For every given trial, the decision maker faces a binary decision with two possible alternatives, A and B. During the decision two processes are active, a *utility process* U and a *heuristic process* H. As in the previous sections, each process follows a diffusion process with symmetric barriers normalized to -1 and 1. However, unlike in the previous section, the drift rates are not given, but rather depend on the subjective evaluations of the alternatives.

The utility process operates on the basis of true underlying utilities u_A and u_B for A and B, respectively. This process always favors the "preferred" alternative, i.e it favors A if $u_A > u_B$ and favors B if $u_B > u_A$. The drift parameter μ_U is a function of $u = |u_A - u_B|$. Specifically, $\mu_U : \mathbb{R}_+ \to \mathbb{R}_+$ is a continuous and strictly increasing function with $\mu_U(0) = 0$. Given u, the probability $P_U(u)$ that the favored response is chosen and the expected response time $T_U(u)$ are obtained by replacing μ by $\mu_U(u)$ in equations (1) and (2). The identity of the favored option depends on the sign of $u_A - u_B$, but $P_U(u)$ and $T_U(u)$ depend only on its magnitude.

Since in a model of preferential choice there will typically not be a normative criterion to determine one response as "correct," we will adopt here the terminology "preferred" and "non-preferred alternative." That is, the preferred alternative is simply that which is preferred by the utility process, in the sense of having a larger utility. For a comparison of results with the fixed-alternatives DPDM, the reader might equate "preferred" in this part of the article with "correct" in the previous one. The heuristic process favors the option which is superior according to some other criterion (e.g., valuation by feeling), captured by real numbers v_A and v_B, the *heuristic utilities*. This might be either the preferred alternative or the non-preferred one. An example might be, e.g., the maximum amount to win in a lottery, ignoring probabilities. Another example might be the familiarity (recognition) of the brand of a consumer item. Analogously to the utility process, the drift parameter μ_H of the heuristic process is a continuous and strictly increasing function of $v = |v_A - v_B|$, specifically $\mu_H : \mathbb{R}_+ \to \mathbb{R}_+$, with $\mu_H(0) = 0$. Given v, the probability $P_H(v)$ that the preferred option is chosen and expected response time $T_H(v)$ are obtained replacing μ by $\mu_H(v)$ in equations (1) and (2).

The assumption that drift rates are increasing in utility differences is a natural, empirically motivated one. It is a long-standing, well-established fact (both in psychology and in economics) that decisions which are closer to indifference take longer (e.g., Dashiell, 1937; Mosteller & Nogee, 1951). This fact has also been observed in the risky-choice studies of Wilcox (1993) and Moffatt (2005). As captured by Lemma 1, drift rates closer to zero are associated with longer response times. That is, a larger utility difference leads to a clear, easier decision resulting in a swifter path for the decision process.

Since the alternatives and their evaluations change from trial to trial within an experiment, we need to specify how the true and the heuristic utilities arise. Since experiments generally aim to randomize the presentation of options, we assume that utilities u_A and u_B are drawn from continuous random variables. Similarly, v_A and v_B are assumed to be drawn from continuous random variables. Since the whole analysis depends only on utility differences, all that is needed is that the differences $u = u_a - u_B$ and $v = v_A - v_B$ are independently distributed according to some (possibly different) density functions, g and h. As in the previous section, the model is completed with a parameter Δ capturing the probability that the heuristic process is selected and determines the actual response. Note that the case of experiments with fixed options from the previous section arises as a special (limit) case in which utilities are constant across trials.

Consider a given trial where the choice between two alternatives A and B is offered. Define an *easy trial* to occur if $\mu_U(|u_A - u_B|) > \mu_H(|v_A - v_B|)$. In this case, the utility process, which favors the preferred response, is actually faster than the heuristic process. Conversely, a *hard trial* is determined by $\mu_U(|u_A - u_B|) < \mu_H(|v_A - v_B|)$. In this case, the heuristic process is faster. This distinction points to a limitation of simple versions of dual-process thinking. For, if one wishes to interpret the processes as more automatic or more controlled, the same process will be more automatic depending on whether the trial is easy or hard.

An alternative interpretation of the DPDM for the case of preferential choice is as follows. Suppose the utility process concentrates on the actually relevant evidence in a multi-attribute decision problem. Hold the drift rate at a medium, positive value. The heuristic process then reflects a distractor attribute, which might be more or less salient. If the distractor attribute is very salient, the drift rate of the heuristic process is high, corresponding to the hard trials. If the distractor corresponds to some background, less-relevant noise, the drift rate is low. This corresponds to the easy trials.

Preferential Choice: Hard Trials

Consider first the case of hard trials. The following theorem derives testable predictions for expected response times. The result is analogous to Theorem 1 for the setting of variable, evaluated alternatives, and the core of its proof is an application of that result.

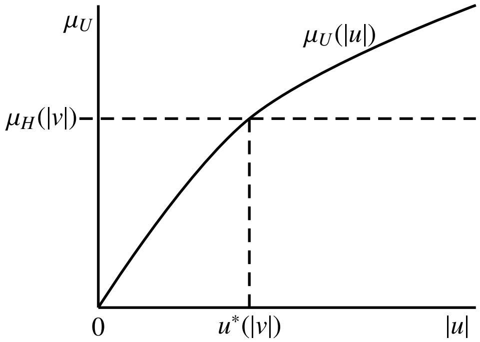

Figure 2. For a fixed value of $v = v_A - v_B$, the fact that μ_H, μ_V are increasing and continuous allows to find a cutting point $u^*(|v|)$ such the pairs (u, v) corresponds to a hard trial for $|u| < u^*(|v|)$ and to an easy trial if $|u| > u^*(|v|)$.

Theorem 3. *Consider the dual-process diffusion model for variable alternatives. For hard trials, the following holds.*

(a) In case of alignment, the expected response time of a non-preferred response is strictly longer than that of a preferred one.

(b) In case of conflict, the expected response time of a preferred response is strictly longer than that of a non-preferred one.

Proof. Since μ_U and μ_H are continuous and increasing and take the value 0 at 0, it follows that for each fixed difference $v = v_A - v_B$ there exists a threshold $u^*(|v|) > 0$ such that $\mu_U(|u|) < \mu_H(|v|)$ for all $|u| < u^*(|v|)$ (leading to a hard trial) and $\mu_U(|u|) > \mu_H(|v|)$ for all $|u| > u^*(|v|)$ (leading to an easy trial). See Figure 2 for an illustration. Formally, $u^*(|v|)$ might be infinite, corresponding to the case where $\mu_U(|u|) < \mu_H(|v|)$ for all u. However, as long as it takes positive values, the threshold function $u^*(\cdot)$ is continuous and strictly increasing, and it fulfills $u^*(0) = 0$.

Consider the event that a realization of utility differences (u, v) gives rise to a hard trial and a case of alignment. Alignment means that u and v have the same sign. For $v > 0$, a hard trial occurs for $0 \leq u < u^*(v)$ (since $|v| = v$). For $v < 0$, it occurs for $-u^*(-v) < u \leq 0$ (since $|v| = -v$). This describes a measurable region in the (u, v)-plane. We can analogously describe all relevant events. See Figure 3 for an illustration.

For each realization of u and v, we obtain fixed values $\mu_U = \mu_U(|v|)$, $\mu_H = \mu_H(|v|)$. The corresponding expected response times $T(\text{Preferred}|u, v)$ and $T(\text{Nonpref}|u, v)$ are then as computed in the proof of Theorem 1, substituting "preferred" and "non-preferred" for "correct" and "error," respectively. Consider the case of hard trials in case of conflict. Denote this event by HC. The expected response times conditional on this event are given by the respective double integrals over the area of HC.

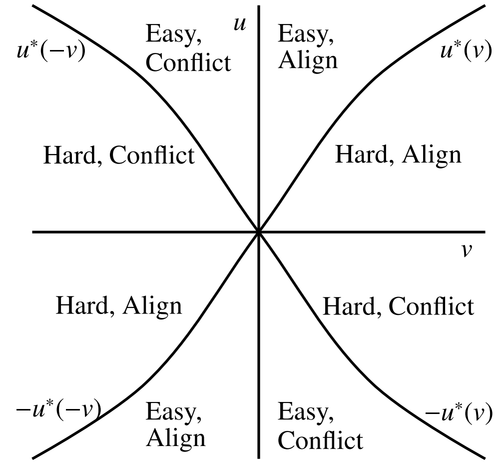

Figure 3. Depiction of the events combining hard or easy trials with alignment or conflict in the space of realizations of utility differences u and heuristic utility differences v.

$$T(\text{Preferred}|HC) = \frac{1}{\Pr(HC)} \iint_{(u,v)\in HC} T(\text{Preferred}|u, v) g(u) h(v) du dv$$

$$T(\text{Nonpref}|HC) = \frac{1}{\Pr(HC)} \iint_{(u,v)\in HC} T(\text{Nonpref}|u, v) g(u) h(v) du dv$$

where g and h are the density functions of u and v, respectively.

For $(u, v) \in HC$, $\mu_U(|u|) < \mu_H(|v|)$ and Theorem 1(b) applies, hence $T(\text{Preferred}|u, v) > T(\text{Nonpref}|u, v)$ for all $(u, v) \in HC$. It follows that $T(\text{Preferred}|HC) > T(\text{Nonpref}|HC)$, proving part (b). The proof of part (a) is completely analogous, integrating over the event corresponding to hard trials in case of alignment and applying Theorem 1(a). ■

The intuition for this result is identical to that of Theorem 1. The reason is that, for the case of hard trials, the heuristic process is always swifter. Hence the intuition conforms to the interpretation that this process is more automatic.

Preferential Choice: Easy Trials

The next result considers the case of easy trials. In case of alignment, the prediction is identical to the one in Theorem 3. In case of conflict, the prediction is the mirror image of the one in that theorem.

Theorem 4. *Consider the dual-process diffusion model for variable alternatives. For easy trials, the following holds.*

(a) In case of alignment, the expected response time of a non-preferred response is strictly longer than that of a preferred one.

(b) In case of conflict, the expected response time of a non-preferred response is strictly longer than that of a preferred one.

This last result needs no formal proof, because it can be derived from Theorem 3 by a symmetry argument. Consider an easy trial. The utility process is swifter than the heuristic process. Exchanging the names of the processes, we can consider if as a hard trial for the processes with exchanged names, with the difference that, because of renaming, in case of conflict what is now "preferred" was non-preferred for the original processes. In case of alignment, the renaming of the processes does not affect what is "preferred". Hence the result follows.

The statement of Theorem 4 is that non-preferred responses will be slower than preferred responses both in case of alignment and in case of conflict. This prediction is particular to the case of easy trials, that is, it obtains when the heuristic process reflects noise due to attributes or factors which are less-relevant than the crucial ones captured by the utility process. For instance, consider perceptual experiments where accuracy is emphasized. In this case non-preferred responses are actually errors, and one could assume the distractor heuristic to be particularly weak, so that most trials are "easy" in the sense defined above. The prediction of Theorem 4 then agrees with the general observation that in this case errors are slower than correct responses independently of whether there is alignment or conflict (Luce, 1986; Ratcliff & Rouder, 1998).

Without further analysis, however, the conclusion of Theorem 4 does not necessarily mean that non-preferred responses will be slower without conditioning on either conflict or alignment. The unconditional expected time of non-preferred responses is a convex combination of the expected time of non-preferred responses in the two cases, weighted by the probabilities of conflict and alignment given a non-preferred response, and analogously for the expected time of preferred responses. Since those probabilities depend, e.g., on the actual probabilities of conflict and alignment in the paradigm at hand, in principle the result could be overturned for the unconditional response times. However, with two minimal additions the model does imply the inequality for unconditional response times. The first addition is simply the assumption that $D_1 \geq D_0$, i.e., conflict detection and resolution time is larger in case of conflict (already discussed above; recall Theorem 2). The second addition is technical. Suppose that the density functions g, h of the utility differences u, v are symmetric, i.e. $g(-u) = g(u)$ and $h(-v) = h(v)$ (which is reasonable if options are randomized). In this case, we obtain the following preliminary result, which is conceptually analogous to Theorem 2.

Lemma 2. *Consider the DPDM for variable alternatives alternatives, with parameters Δ, D_0, and D_1 with $D_1 \geq D_0$, and symmetric density functions for utility differences u, v.*

For easy trials, the expected response time of non-preferred responses in case of alignment is shorter than the expected response time of non-preferred responses in case of conflict.

Proof. For a fixed, easy trial, one obtains $T_H > T_U$ and $P_H < P_U$. A completely analogous to the argument in Theorem 2 establishes that, for such values, $T(Nonpref|Alignment) < T(Nonpref|Conflict)$. Consider a fixed (u, v) leading to an easy trial (hence $T_H(\mu_H(|v|) > T_U(\mu_U(|u|))$ and a case of alignment. Then, $(-u, v)$ also leads to an easy trial but a case of conflict (Figure 3). Since both (u, v) and $(-u, v)$ lead to the same drift rates of both processes, we obtain $T(Nonpref|u, v) < T(Nonpref| - u, v)$. Integrating this inequality over all (u, v) leading to easy trials in case of alignment is equivalent to integrating over all $(-u, v)$ leading to easy trials in case of conflict (since $g(-u) = g(u)$). The integrals then deliver the desired result (note that, by symmetry, of g and h, the total probability mass of the easy-alignment and the easy-conflict areas is identical). ■

This lemma can then be used to show that, in the case of easy trials, non-preferred responses will also be slower than preferred ones without conditioning on alignment or conflict.

Corollary 2. *Consider the DPDM for variable alternatives and easy trials, with parameters Δ, D_0, and D_1 with $D_1 \geq D_0$, and symmetric density functions for utility differences u, v. The unconditional expected response time of non-preferred responses is strictly longer than the unconditional expected response time of preferred ones.*

Proof. All quantities below are conditional on easy trials; I skip this conditioning to ease notation. Let $p = \Pr(\text{Conflict}|\text{Nonpref})$ and $q = \Pr(\text{Conflict}|\text{Preferred})$. It will be shown that $p > q$, that is, it is more likely to be in a situation of conflict if a non-preferred response is observed than if a preferred one is observed. To see this, note that for fixed values of u, v, setting $P_H = P_H(\mu_H(|v|))$ and $P_U = P_U(\mu_U(|u|))$,

$$\begin{aligned}&\Pr(\text{Nonpref}|\text{Conflict}, u, v) = \\ &\quad \Delta P_H + (1-\Delta)(1-P_U) > \Delta(1-P_H) + (1-\Delta)(1-P_U) \\ &\quad = \Pr(\text{Nonpref}|\text{Alignment}, -u, v)\end{aligned}$$

simply because $P_H > \frac{1}{2}$. As in the proof of the last lemma, note that (u, v) results in an easy trial and a case of conflict if and only if $(-u, v)$ results in an easy trial and a case of alignment. Integrating over all values of (u, v) yielding a conflict in easy trials (and using symmetry of g

and h; recall Figure 3) shows that Pr(Nonpref|Conflict) > Pr(Nonpref|Alignment). Analogously, we obtain that Pr(Preferred|Conflict) < Pr(Preferred|Alignment).

From the first inequality and the equation

$$\begin{aligned}\Pr(\text{Nonpref}) &= \Pr(\text{Nonpref}|\text{Conflict})\Pr(\text{Conflict})\\ &\quad + \Pr(\text{Nonpref}|\text{Alignment})\Pr(\text{Alignment})\end{aligned}$$

it follows that

$$\frac{\Pr(\text{Nonpref})}{\Pr(\text{Nonpref}|\text{Conflict})} < 1$$

and analogously

$$\frac{\Pr(\text{Preferred})}{\Pr(\text{Preferred}|\text{Conflict})} > 1.$$

Then,

$$\begin{aligned}p &= \frac{\Pr(\text{Nonpref}|\text{Conflict})\Pr(\text{Conflict})}{\Pr(\text{Nonpref})}\\ &> \frac{\Pr(\text{Preferred}|\text{Conflict})\Pr(\text{Conflict})}{\Pr(\text{Preferred})} = q\end{aligned}$$

establishing that $p > q$.

Let T_{NA}, T_{NC}, T_{PA}, and T_{PC} be the expected response time of nonpreferred responses conditional on alignment, of nonpreferred responses conditional on conflict, of preferred responses conditional on alignment, and of preferred responses conditional on conflict, respectively. By Theorem 4, we have that $T_{NA} > T_{PA}$ and $T_{NC} > T_{PC}$. By Lemma 2, we have that $T_{NC} > T_{NA}$. Let T(Nonpref) and T(Preferred) denote the expected response times of nonpreferred and preferred responses in easy trials, respectively.

$$\begin{aligned}T(\text{Nonpref}) &= p \cdot T_{NC} + (1-p)T_{NA} = T_{NA} + p(T_{NC} - T_{NA}) >\\ &\quad T_{NA} + q(T_{NC} - T_{NA}) = q \cdot T_{NC} + (1-q)T_{NA} >\\ &\quad q \cdot T_{PC} + (1-q)T_{PA} = T(\text{Preferred})\end{aligned}$$

where the first inequality follows from $T_{NC} - T_{NA} > 0$ and $p > q$, and the second inequality follows from $T_{NA} > T_{PA}$ and $T_{NC} > T_{PC}$. This completes the proof. ■

A Consumer-Choice Example

As an example of application, I analyzed data from a recent study on brand recognition and product ratings by Thoma and Williams (2013). Participants were presented with a variety of choices among consumer items (see Thoma & Williams, 2013, for details). Crucially for our purposes, the items varied in two attributes. The first was the brand of the product. Stimuli were labeled with famous and unknown brands, and hence brand heuristic would lead to the selection of the famous brand. However, the second stimuli was a consumer rating in the form of a certain number of stars. Stimulus pairs created three within conditions, according to whether the product from a famous brand was labeled with more, less, or the same stars as the product from a nonfamous brand. In the terms of the previous sections, the first and the second classes correspond to alignment and conflict, respectively.

Thoma and Williams (2013) recorded response times for all choices, but used them for a different purpose. I reanalyzed their data in order to test the model. The steps of the analysis were as follows. First, Thoma and Williams (2013) checked whether the famous brands were actually recognized by the participants and reported this in their data. I hence concentrated on the trials in which there was actual recognition. Second, response times above 10 seconds were excluded (the grand average of response times in the study was 2.86 s, with a standard deviation of 3.00 s; some response times were above 30 s, one around 85 s). Then, individual average response times were computed conditional on conflict and alignment (those were taken to be the unit of analysis since the predictions derived from the DPDM in this work refer to expected conditional response times), and conditional on the actual response, i.e., whether the product from the famous brand was chosen or not. For each test below, only participants who made at least two choices of each type were considered.

In terms of the DPDM, this example fits the case of preferential choice, because the actual choices varied among very different products from trial to trial. Brand recognition is the natural candidate for the heuristic process in the model, and the utility process can be taken to be based on the number of stars observed. This kind of recognition is often assumed to be rather quick and the experiment always used famous vs. nonfamous brands, thus the difference in heuristic utilities should be large. Hence it is reasonable to assume that, in the terminology of the model for preferential choice, the experiment generated hard trials. Hence the predictions are derived from Theorem 3, taking "preferred" to mean the product with a larger number of stars.

Consider the case of conflict, i.e., trials where the product from the famous brand was given less stars than the alternative product. The prediction from Theorem 3(b) is that preferred responses should be slower than non-preferred ones. In this case, choosing the famous product is "nonpreferred." The mean of the individual average response times when choosing non-famous (preferred) products was 2.81 s, against 2.53 s when choosing famous (non-preferred) products. The difference in distributions was significant according to a Wilcoxon Signed-Rank test ($N = 26$, $z = 1.968$, $p = .049$).

The case of alignment corresponds to trials where the product from the famous brand was given more stars than the alternative product, and the "preferred" response was to choose the famous product. The prediction from Theorem

3(a) is that non-preferred responses should be slower than preferred ones. The mean of the individual average response times when choosing famous (preferred) products was 2.51 s, against 3.00 s when choosing non-famous (non-preferred) products. The difference in distributions was weakly significant according to a Wilcoxon Signed-Rank test ($N = 26$, $z = -1.892$, $p = .059$).

Remark 2. Brand recognition can be considered as a particular case of the recognition heuristic (see, e.g., Gigerenzer & Goldstein, 2011; Glöckner & Bröder, 2011; Goldstein & Gigerenzer, 2002; Pachur & Hertwig, 2006). The recognition heuristic has been used as a case study for sophisticated cognitive architectures as ACT-R (Marewski & Mehlhorn, 2011). I use the study of Thoma and Williams (2013) here merely as an example of multi-attribute choice. In this sense, the processes assumed by the DPDM in this example remain "black boxes" whose favored response can be readily identified. In this and many other cases, the DPDM has to be understood as an "as if" model.

Discussion

The first subsection below compares the DPDM to other available models from the literature. The second subsection discusses the limitations of the model and points out possible extensions.

Comparison with Other Models

Decision Field Theory. The model presented here is conceptually related to Decision Field Theory, or DFT for short (see Busemeyer & Diederich, 2002, for a review). As presented in Busemeyer and Townsend (1993), DFT is a stochastic-dynamic model of choice behavior which was originally developed for binary choices. Diederich (1997) extended the theory to multi-attribute decision making, and this latter extension bears several conceptual similarities with the model presented here. DFT assumes that decision makers update "strength of preference" values in time by adding valence inputs. When those valences come from several dimensions (multi-attribute setting), the update is based on weighted evaluations. As in weighted additive utility models, the weights represent the relative importance or attention of the attributes. However, the weights themselves change over time according to a stationary stochastic process $W(t)$ (a Markov process in Diederich, 1997). In the optional stopping time version of DFT, a response is generated when one of the preference strengths reaches a given threshold.

Consider a two-attribute, binary choice model. Each of the attributes, which provide input for the dynamic updating of preference strengths, could be considered to be analogous to a decision process. The ideas in the model presented here can be recast by assuming the two attributes to represent a heuristic evaluation and a relevant-normative one, respectively. A situation of conflict is one where the two attributes point toward different options.

In DFT, however, preference strengths are updated for each option, and the attributes combine and influence both at every point in time. If the DPDM is reinterpreted in terms of DFT, what is updated are preference-strength differences, one for each attribute (process). That is, while in DFT multiple attributes are constantly combined in order to generate weighted preferences for certain, given weights $W(t)$, in the DPDM the attributes are never combined into a signal measure. Rather, preference, although constantly changing in time, remains multidimensional, and the parameter Δ determines which dimension will eventually dominate.

The differences are best seen if one considers the continuous-time limit of the DFT, which is actually a diffusion process (see, e.g., Busemeyer & Diederich, 2002, for details). By construction, the total change in preference strength in DFT is zero, that is, whatever preference strength an option gains, another option loses. This means that the model reduces to a *single* diffusion process, with within-trial changes in the drift rate. In contrast, the DPDM operates with two independent diffusion processes with constant but different drifts. This raises an interesting possibility: for the case of binary choice, one could formulate a DPDM where one of the diffusion processes captures a DFT model, while the other remains a conceptually simple automatic process. The result would be a model where deliberative aspects of thinking follow DFT, but are perturbed by a more heuristic process.

Linear Ballistic Accumulation. The LBA model arose as a simplification of the leaky competing accumulator model of Usher and McClelland (2001). In that model, each possible option is associated to an accumulator which changes in time, gathering evidence for the option while past evidence decays. As in DFT, an option is chosen when one of the accumulators reaches a given threshold. In Usher and McClelland (2001), the rates of evidence accumulation (drift rates) were allowed to fluctuate across trials. Brown and Heathcote (2005) simplified the model by dropping the within-trial randomness, and Brown and Heathcote (2008) further assumed linear evidence accumulation. The resulting model is the LBA, which has the advantage of being analytically solvable and has been shown to capture many of the empirical regularities reflected in the leaky competing accumulator model.

The key characteristic that the LBA model shares with the DPDM is that both assume separate evidence accumulation processes with constant drift rates, rather than relying on weighted preferences with within-trial variability in drift. The key difference is that, as DFT, the LBA model considers evidence accumulation for each option, with the threshold parameter representing the finish line of a race, while the DPDM considers evidence accumulation for each attribute, with the process selection parameter determining the probability with which each process will ultimately be decisive.

However, the LBA model has two additional sources of randomness in comparison to the DPDM. First, the drift rates, while fixed, are initially drawn from normal distributions with given, different means for the different options. That is, the parameters for fixed alternatives are the means of the distributions of drift rates, while in the DPDM the parameters are directly the drift rates. Second, the starting points for evidence accumulators in the LBA are also random values, drawn from a uniform distribution on an interval. As the DPDM, the LBA can be used to predict relative speeds of correct and incorrect response times, but it does so relying on the interplay between start point variability and drift rate variability (Brown & Heathcote, 2008, p. 160), while the DPDM does not rely on either of them.

A nice illustration of the LBA was provided by Eidels (2012), who showed that the model can predict the Stroop effect by relying on four independent channels (presented word, presented color, non-presented word, and non-presented color) and ten parameters. The point of Eidels (2012) was to show (numerically) that the Stroop effect can be encompassed in a stochastic model not involving a cross-talk between the channels for word and color. The DPDM achieves the same effect (as an analytical result; recall Theorem 2) with two processes, one reflecting word recognition and another reflecting color recognition.

Stroop Counter Model. Trainham, Lindsay, and Jacoby (1997) postulate two processes corresponding to word and color recognition, respectively. The model evolves in finite discrete intervals. In each one, a piece of evidence is received and each process has a given probability, $W(t)$ and $C(t)$, of allocating the count to the target response counter (parameterized as a gamma density function and a cumulative gamma function, respectively). In case both processes receive the chance to allocate the count, the word recognition process dominates. This model shares with the DPDM the idea of modeling word and color recognition as different processes which are viewed as alternative (no evidence is aggregated). However, the model is very different. First, the processes' "dominant responses" are deterministic, that is, the word recognition process either allocates the count to the named color or delivers no evidence. In the DPDM, each process can potentially deliver each of the two possible responses. Second, although processes are viewed as alternatives, this occurs at the count level, while in the DPDM this occurs at the aggregate level (Δ determines which process determines the final answer).

The DSTP The Dual-Stage Two-Phase model of Hübner, Steinhauser, and Lehle (2010) targets conflict situations, especially in the flanker task. It incorporates two discrete, sequential stages (see also Gratton, Coles, & Donchin, 1992). The first one is a race model where two diffusion processes run in parallel. One process is for response selection and is called *early selection process*, the other is for target selection and is called *late selection process*. If the early selection process finishes first, it determines the decision and the trial ends. If the late selection process wins the race, the model enters the second stage, where the decision-selection process continues in isolation with a new, different drift rate (positive or negative depending on target selection). This represents the idea that the target of attention (central stimulus or flanker) has been selected and there is no further interference.

The DSTP of Hübner et al. (2010) is quite different from the DPDM, because its main emphasis is on the two discrete stages (hence "dual-stage" and not "dual-process"). Still, it has to be credited as the first to explicitly incorporate the basic dual-process idea of different processes in a diffusion-process setting. Hübner et al. (2010) provided evidence that the DSTP delivers a better data fit than a single-process model in the flanker task, but White et al. (2011) found that a single-process with inter-trial variation of the drift rate (which resembles the results of the DPDM) fares better.

Dual-response Bayesian Sequential Risk-Taking. Pleskac and Wershbale (2014) have recently developed a model to explain data from the Balloon-Analogue Risk Task (BART), where participants sequentially pump a balloon trying to inflate it as much as possible but stopping before it explodes. The model is an explicit dual-process model which shares a crucial feature with the DPDM. At each pump opportunity, there is an explicit probability that the decision to pump or not will be made by a slow, deliberative process. As in the DPDM, the deliberative process is modeled as a diffusion process which gradually gathers evidence in order to decide whether to pump or to stop. With the remaining probability, an automatic process takes over. In contrast to the DPDM, this automatic process is assumed to simply pump, i.e., it has zero response variability. Since the automatic process is not explicitly modeled, Pleskac and Wershbale (2014) directly assume an ex-Gaussian distribution for the response times arising from this process. This difference is crucial. Mathematically, for each isolated decision the model of Pleskac and Wershbale (2014) could be considered as a limit case of the DPDM where the drift rate of the heuristic process is taken to infinity, hence obtaining a zero fault probability. However, this would also imply a response time of zero for that process.

Pleskac and Wershbale (2014) defend the dual-process approach showing that the addition of a probability that an automatic process makes the decision considerably improves the fit to experimental data from the BART. Their approach hence shares a basic motivation with the DPDM. Further, the basic building block that a decision is made by either one or another process is common to both models. However, the model of Pleskac and Wershbale (2014) is tailored to the

BART as an example of sequential decisions. This difference is crucial and makes a direct comparison impossible. The DPDM is designed to model a static decision problem, corresponding to experiments with either isolated or repeated decisions with no intertemporal structure, e.g., a series of analogous decisions. In the BART, decisions are part of a sequence with explicit interdependences. The first decision to stop, or the first negative feedback, end the sequence, while each decision to pump opens the door to a further decision. As a consequence, in the model of Pleskac and Wershbale (2014) both the deliberative process' drift parameter and the probability that the automatic process makes the decision are assumed to depend on the trial number, reflecting learning as the number of pumps increases (making a decision after n pumps implies there have been n successful pumps). In contrast to the DPDM, the model is designed to fit response time evidence as a function of pump number.

A Default-Interventionist Model. Klauer and Voss (2008) discussed several models in the context of a weapon identification task where race is primed before a stimuli (gun or tool) is shown and has to be identified. The task is difficult to compare with the ones discussed here due to the sequential presentation of the stimuli and its reliance on priming. However, several of the models considered are of interest beyond the particular task. One such model is a "default-interventionist" dual-process model which bears a certain similarity to a simplified version of the DPDM (see Klauer & Voss, 2008, Figure 4). An automatic bias cues a default response, which is driven by the race prime but can be any of the two options with certain probabilities. With a fixed probability, the automatic process simply determines the response on the basis of the default. With the remaining probability, an intervention occurs and the controlled process delivers the correct response with certainty (tool or gun). That is, the controlled process has zero response variability. Response times are simply assumed to be fast for the automatic process and slow for the controlled process. Hence, errors arise only from the automatic process and are always fast, while correct responses are a mixture of fast and slow decisions. This simple model gives a stylized explanation for faster errors in dual-process terms, but cannot explain asymmetries as that predicted by the DPDM (recall Theorem 1). Interestingly, however, if the model was enriched to allow for response variability for the controlled process, the result would be a version of the DPDM where processes are simple biased coin tosses instead of diffusion processes. Still, to reproduce the predictions of the DPDM one would need to specify fault probabilities fulfilling the properties identified in Corollary 1; this was exactly the approach taken in Achtziger and Alós-Ferrer (2014).

Outlook and Future Extensions

This work has proposed a parsimonious two-process model building on the simplest version of the well-known diffusion model. The model fits available empirical evidence on response times and is analytically tractable, in the sense that testable predictions can be derived without the recourse to numerical examples. It can be taken as a structuring framework for the study of response times in decision making, which can be put to the test in the laboratory and potentially expanded in a number of interesting directions.

The predictions translate into simple tests for response times conditional on correct and erroneous responses. As long as an experimental paradigm generates enough data, the ideal tests are within participants (paired samples), for they abstract from individual parameter variability. The paradigm, however, needs to identify which responses are correct and which are errors, or alternatively to clearly identify the favored responses of the two candidate processes (recall the consumer-choice example above). For inference models as those often used in judgment and decision making, this is rarely a problem, as there exist normatively correct answers. For preferential choice, special care should be taken to ensure that this is possible. For instance, studies in risky choice using classical paradigms might not be well suited for this purpose, because preference elicitation itself is known to be subject to a number of biases and be generally noisy (Butler & Loomes, 2007; Schmidt & Hey, 2004) (see Alós-Ferrer, Granić, Kern, & Wagner, 2015 for a model of response times in such settings). Further, the fact that choice might influence preferences within an experiment (Ariely & Norton, 2008; Brehm, 1956) presents a particular challenge for experimental design. Developing improved paradigms delivering appropriate data represents an interesting avenue of research. For instance, in the domain of risky choice this could be accomplished by estimating the preferences (e.g., a risk aversion parameter) in a separate set of lotteries before considering the actual choices affected by a particular heuristic process.

The model has of course a number of limitations, some of which could be addressed with extensions in a number of natural directions. First, the model targets only binary decisions at this point, and an extension to multiple alternatives would require an approach not unlike the assumption of binariness in decision theory. In contrast, an extension to more than two processes is straightforward, since the modeling unit is the process and not the alternative. Second, it does not account for speed-accuracy tradeoffs (Heitz, 2014), where time pressure results in an increase in error rates. For a single diffusion process, speed-accuracy tradeoffs can be modeled through the assumption that time pressure reduces the barrier (Ratcliff & Smith, 2004). The building block of the DPDM are diffusion processes with symmetric and identical barriers. Hence, it would be possible to incorporate a similar assumption in the DPDM, e.g., by assuming that only the barrier of the utility process is affected (recall, however, Remark 1). Third, dual-process theories assume that experimental manipulations as cognitive load or even time pressure can shift the balance among processes. This point has not been addressed in the current version of the DPDM, because the exogenous parameter Δ has been left unspecified. This is reasonable for the objectives of the work at hand, because the exact value of Δ does not affect the predictions of the formal results presented here. However, it would be natural to assume that, e.g., cognitive load increases Δ, resulting in a larger share of decisions accruing to the heuristic process. This could also provide a different channel for modeling speed-accuracy tradeoffs.

In the context of the preferential choice model, it is also natural to speculate that the probability with which a process is selected might depend, e.g., on whether a trial is hard or easy. More generally, extensions of the DPDM might fully endogenize the value of Δ in such a way that its value would be naturally linked to the kind of trial. A first possibility would be to consider a race among the two diffusion processes, analogous to race models among accumulators (LaBerge, 1962; Usher & McClelland, 2001; Vickers, 1970). However, such an extended model would have the unappealing property that most of the time the swifter process would win the race. An alternative to race models would be to endogenize process selection through a two-phase approach as in Hübner et al. (2010).

A further, practical limitation of the preferential choice model is that it implicitly assumes that in actual applications it will be easy to distinguish whether hard or easy trials are involved. This is of course a matter of proper experimental design, but still some paradigms might actually generate both, and, unless those can be distinguished, sharper assumptions will be needed to obtain unconditional predictions.

It should be remarked that, at this point, the DPDM is conceived as a simple, parsimonious formal-analytical framework delivering easily testable predictions of an ordinal nature (which response times are larger) at the population level. In particular, this work has abstracted from possible fitting to data at the individual level in order to estimate the actual parameters. The setting, however, is based on the diffusion model, for which such approaches have been enormously fruitful, and hence, given appropriate data, it would in principle be possible to perform such estimations.

To conclude, the DPDM is a formal, stylized model using single-strategy building blocks within a multi-strategy (dual-process) setting, which delivers predictions on response times without imposing sharp conditions on parameter values. Future research should build upon the basic structure of the model to exploit the advantages of both the multi-strategy approach and the single-strategy components. In this sense, the present article contributes to the literature on strategy selection by demonstrating how simple formal models can help build bridges between single-strategy and multi-strategy models.

References

Achtziger, A., & Alós-Ferrer, C. (2014). Fast or Rational? A Response-Times Study of Bayesian Updating. *Management Science*, *60*(4), 923-938.

Achtziger, A., Alós-Ferrer, C., Hügelschäfer, S., & Steinhauser, M. (2014). The Neural Basis of Belief Updating and Rational Decision Making. *Social Cognitive and Affective Neuroscience*, *9*(1), 55-62.

Alós-Ferrer, C., Granić, D.-G., Kern, J., & Wagner, A. K. (2015). Preference Reversals: Time and Again. *Journal of Risk and Uncertainty*, *forthcoming*.

Alós-Ferrer, C., Granić, D.-G., Shi, F., & Wagner, A. K. (2012). Choices and Preferences: Evidence from Implicit Choices and Response Times. *Journal of Experimental Social Psychology*, *48*(6), 1336-1342.

Alós-Ferrer, C., & Strack, F. (2014). From Dual Processes to Multiple Selves: Implications for Economic Behavior. *Journal of Economic Psychology*, *41*, 1-11.

Ariely, D., & Norton, M. I. (2008). How Actions Create – Not Just Reveal – Preferences. *Trends in Cognitive Sciences*, *12*, 13-16.

Bem, D. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, *74*, 183-200.

Brehm, J. W. (1956). Postdecision Changes in the Desirability of Alternatives. *Journal of Abnormal and Social Psychology*, *52*, 384-389.

Brown, S. D., & Heathcote, A. (2005). A Ballistic Model of Choice Response Time. *Psychological Review*, *112*(1), 117–128.

Brown, S. D., & Heathcote, A. (2008). The Simplest Complete Model of Choice Response Time: Linear Ballistic Accumulation. *Cognitive Psychology*, *57*, 153–178.

Busemeyer, J. R., & Diederich, A. (2002). Survey of Decision Field Theory. *Mathematical Social Sciences*, *43*, 345–370.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment. *Psychological Review*, *100*(3), 432–459.

Butler, D. J., & Loomes, G. (2007). Imprecision as an Account of the Preference Reversal Phenomenon. *The American Economic Review*, *97*(1), 277–297.

Charness, G., & Levin, D. (2005). When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. *American Economic Review*, *95*(4), 1300-1309.

Dashiell, J. F. (1937). Affective Value-Distances as a Determinant of Aesthetic Judgment-Times. *American Journal of Psychology*, *50*, 57-67.

De Neys, W. (2006). Automatic-Heuristic and Executive-Analytic Processing During Reasoning: Chronometric and Dual-Task Considerations. *The Quarterly Journal of Experimental Psychology*, *59*(6), 1070-1100.

De Neys, W., & Glumicic, T. (2008). Conflict Monitoring in Dual Process Theories of Thinking. *Cognition*, *106*(3), 1248-1299.

De Neys, W., Vartanian, O., & Goel, V. (2008). Smarter than We Think: When Our Brains Detect That We Are Biased. *Psychological Science*, *19*(5), 483-489.

Diederich, A. (1997). Dynamic Stochastic Models for Decision Making under Time Constraints. *Journal of Mathematical Psychology*, *41*(3), 345–370.

Douady, R. (1998). Closed Form Formulas for Exotic Options and Their Lifetime Distribution. *International Journal of Theoretical and Applied Finance*, *2*(1), 17-42.

Eidels, A. (2012). Independent Race of Colour and Word can Predict the Stroop Effect. *Australian Journal of Psychology*, *64*, 189-198.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of Noise Letters Upon the Identification of a Target Letter in a Nonsearch Task. *Perception & Psychophysics*, *16*, 143-149.

Evans, J. S. (2008). Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition. *Annual Review of Psychology*, *59*, 255-278.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Evanston, IL Row: Peterson.

Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, *19*(4), 25-42.

Gigerenzer, G., & Goldstein, D. G. (2011). The Recognition Heuristic: A Decade of Research. *Judgment and Decision Making*, *6*(1), 100-121.

Glöckner, A., & Betsch, T. (2011). The Empirical Content of Theories in Judgment and Decision Making: Shortcomings and Remedies. *Judgment and Decision Making*, *6*(8), 711-721.

Glöckner, A., & Bröder, A. (2011). Processing of Recognition Information and Additional Cues: A Model-Based Analysis of Choice, Confidence, and Response Time. *Judgment and Decision Making*, *6*(1), 23-42.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of Ecological Rationality: The Recognition Heuristic. *Psychological Review*, *109*, 75-90.

Grasman, R. P., Wagenmakers, E.-J., & van der Maas, H. L. (2009). On the Mean and Variance of Response Times Under the Diffusion Model With an Application to Parameter Estimation. *Journal of Mathematical Psychology*, *53*, 55-68.

Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Optimizing the Use of Information: Strategic Control of Activation of Responses. *Journal of Experimental Psychology: General*, *121*, 480-506.

Grether, D. M. (1980). Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *Quarterly Journal of Economics*, *95*, 537-557.

Grether, D. M. (1992). Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence. *Journal of Economic Behavior and Organization*, *17*, 31-57.

Heitz, R. P. (2014). The Speed-Accuracy Tradeoff: History, Physiology, Methodology, and Behavior. *Frontiers in Neuroscience*, *8*, Article 150.

Hsee, C. K., & Rottenstreich, Y. (2004). Music, Pandas, and Muggers: On the Affective Psychology of Value. *Journal of Experimental Psychology: General*, *133*(1), 23-30.

Hsee, C. K., Rottenstreich, Y., & Xiao, Z. (2005). When is More Better? On the Relationship Between Magnitude and Subjective Value. *Current Directions in Psychological Science*, *14*(5), 234-237.

Hübner, R., Steinhauser, M., & Lehle, C. (2010). A Dual-Stage Two-Phase Model of Selective Attention. *Psychological Review*, *117*(3), 759-784.

Kahneman, D., & Tversky, A. (1972). Subjective Probability: A Judgment of Representativeness. *Cognitive Psychology*, *3*, 430-454.

Kahneman, D., & Tversky, A. (1973). On the Psychology of Prediction. *Psychological Review*, *80*, 237-271.

Karatzas, I., & Shreve, S. E. (1998). *Brownian Motion and Stochastic Calculus, 2nd Ed.* Berlin, Heidelberg: Springer Verlag.

Klauer, K. C., & Voss, A. (2008). Effects of Race on Responses and Response Latencies in the Weapon Identification Task: A Test of Six Models. *Personality and Social Psychology Bulletin*, *34*, 1124-1140.

LaBerge, D. (1962). A Recruitment Theory of Simple Behavior. *Psychometrika*, *27*, 375-396.

Lichtenstein, S., & Slovic, P. (1971). Reversals of Preference Between Bids and Choices in Gambling Decisions. *Journal of Experimental Psychology*, *89*(1), 46-55.

Luce, R. (1986). *Response Times*. New York: Oxford University Press).

MacCleod, C. M. (1991). Half a Century of Research on the Stroop Effect: An Integrative Review. *Psychological Bulletin*, *109*(2), 163-203.

Marewski, J. N., & Link, D. (2014). Strategy Selection: An Introduction to the Modeling Challenge. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*, 39-59.

Marewski, J. N., & Mehlhorn, K. (2011). Using the ACT-R Architecture to Specify 39 Quantitative Process Models of Decision Making. *Judgment and Decision Making*, *6*(6), 439-519.

Moffatt, P. G. (2005). Stochastic Choice and the Allocation of Cognitive Effort. *Experimental Economics*, *8*(4), 369-388.

Mosteller, F., & Nogee, P. (1951). An Experimental Measurement of Utility. *Journal of Political Economy*, *59*, 371-404.

Pachur, T., & Hertwig, R. (2006). On the Psychology of the Recognition Heuristic: Retrieval Primacy as a Key Determinant of its Use. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 983-1002.

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The Effect of Stimulus Strength on the Speed and Accuracy of a Perceptual Decision. *Journal of Vision*, *5*, 376-404.

Philiastides, M. G., & Ratcliff, R. (2013). Influence of Branding on Preference-Based Decision Making. *Psychological Science*, *24*(7), 1208-1215.

Pleskac, T. J., & Wershbale, A. (2014). Making Assessments While Taking Repeated Risks: A Pattern of Multiple Response Pathways. *Journal of Experimental Psychology: General*, *143*(1), 142-162.

Ratcliff, R. (1978). A Theory of Memory Retrieval. *Psychological Review*, *85*, 59-108.

Ratcliff, R. (1981). A Theory of Order Relations in Perceptual Matching. *Psychological Review*, *88*, 552-572.

Ratcliff, R. (2013). Parameter Variability and Distributional Assumptions in the Diffusion Model. *Psychological Review*, *120*(1), 281-292.

Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, *9*(5), 347-356.

Ratcliff, R., & Smith, P. L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychological Review*, *111*(2), 333-367.

Rottenstreich, Y., & Hsee, C. K. (2001). Money, Kisses, and Electric Shocks: On the Affective Psychology of Risk. *Psychological Science*, *12*(3), 185-190.

Schmidt, U., & Hey, J. D. (2004). Are Preference Reversals Errors? An Experimental Investigation. *The Journal of Risk and Uncertainty*, *29*(3), 207–218.

Sloman, S. A. (1996). The Empirical Case for Two Systems of Reasoning. *Psychological Bulletin*, *119*(1), 3-22.

Smith, P. L. (2000). Stochastic Dynamic Models of Response Time and Accuracy: A Foundational Primer. *Journal of Mathematical Psychology*, *44*, 408-463.

Söllner, A., Bröder, A., Glöckner, A., & Betsch, T. (2014). Single-Process versus Multiple-Strategy Models of Decision Making: Evidence from an Information Intrusion Paradigm. *Acta Psychologica*, *146*, 84-96.

Strack, F., & Deutsch, R. (2004). Reflective and Impulsive Determinants of Social Behavior. *Personality and Social Psychology Review*, *8(3)*, 220-247.

Stroop, J. R. (1935). Studies of Interference in Serial Verbal Reactions. *Journal of Experimental Psychology*, *35*, 643-662.

Thoma, V., & Williams, A. (2013). The Devil You Know: The Effect of Brand Recognition and Product Ratings on Consumer Choice. *Judgment and Decision Making*, *8*(1), 34-44.

Townsend, J. T., & Ashby, F. G. (1983). *Stochastic Modeling of Elementary Psychological Processes*. Cambridge, MA: Cambridge University Press).

Trainham, T. N., Lindsay, S., & Jacoby, L. L. (1997). Stroop Process Dissociations: Reply to Hillstrom and Logan (1997). *Journal of Experimental Psychology: Human Perception and Performance*, *23*(5), 1579-1587.

Tversky, A., & Kahneman, D. (1983). Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment. *Psychological Review*, *90*(4), 293-315.

Usher, M., & McClelland, J. L. (2001). The Time Course of Perceptual Choice: The Leaky, Competing Accumulator Model. *Psychological Review*, *108*(3), 550-592.

Vickers, D. (1970). Evidence for an Accumulator Model of Psychophysical Discrimination. *Ergonomics*, *13*, 37-58.

Weber, E. U., & Johnson, E. J. (2009). Mindful Judgment and Decision Making. *Annual Review of Psychology*, *60*, 53-85.

White, C. N., Ratcliff, R., & Starns, J. J. (2011). Diffusion Models of the Flanker Task: Discrete Versus Gradual Attentional Selection. *Cognitive Psychology*, *63*, 210-238.

Wickelgren, W. A. (1977). Speed-Accuracy Tradeoff and Information Processing Dynamics. *Acta Psychologica*, *41*, 67-85.

Wilcox, N. T. (1993). Lottery Choice: Incentives, Complexity, and Decision Time. *Economic Journal*, *103*(421), 1397-1417.