Automatic biases in intertemporal choice

Wenjia Joyce Zhao[1] · Adele Diederich[2] · Jennifer S. Trueblood[3] · Sudeep Bhatia[1]

Abstract

Dual process theories of intertemporal decision making propose that decision makers automatically favor immediate rewards. In this paper, we use a drift diffusion model to implement these theories, and empirically investigate the role of their proposed automatic biases. Our model permits automatic biases in the response process, in the form of a shifted starting point, as well as automatic biases in the evaluation process, in the form of an additive drift rate intercept. We fit our model to individual-level choice and response time data, and find that automatic biases (as measured though the starting point and drift rate intercept in our model) are prevalent in intertemporal choice, but that the type, magnitude, and direction of these biases vary greatly across individuals. Our results pose new challenges for theories of intertemporal choice behavior.

Keywords Drift diffusion model · Intertemporal choice · Computational modelling · Automatic bias · Dual process theories

Introduction

Intertemporal choices, i.e., choices between rewards and punishments at different points in time, are often described as a product of automatic and controlled processes. Decision makers are assumed to be automatically biased to select immediate rewards. These biases may or may not be circumvented by a control process that monitors the decision and coordinates thoughts and actions with internal goals, which sometimes support the selection of delayed rewards (see, e.g., Burks et al., 2009; Figner et al., 2010; Hare et al.,

✉ Wenjia Joyce Zhao
    zhaowenj@sas.upenn.edu

    Adele Diederich
    a.diederich@jacobs-university.de

    Jennifer S. Trueblood
    jennifer.s.trueblood@vanderbilt.edu

    Sudeep Bhatia
    bhatiasu@sas.upenn.edu

[1]   University of Pennsylvania, Philadelphia, PA, USA

[2]   Jacobs University, Bremen, Germany

[3]   Vanderbilt University, Nashville, TN, USA

2009; Loewenstein et al., 2015; McClure et al., 2004; Peters & Büchel, 2011; Shamosh et al., 2008; but also see Franco-Watkins, Rickard, & Pashler, 2010 and Lindner & Rose, 2017 for contradictory evidence).

One approach to testing the predictions of such dual process theories has been to examine response time (RT) patterns for different types of rewards. If decision makers are automatically biased to choose a reward, trials in which that reward is chosen should have, on average, shorter RTs than corresponding trials in which the alternate reward is chosen (e.g., Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Rand, Greene, & Nowak, 2012; Rubinstein, 2007). In the domain of intertemporal choice, this would imply that immediate rewards are chosen faster than delayed rewards. However, using just RTs to infer automatic biases in intertemporal choices is problematic, as RTs also reflect choice difficulty. For example, an observation of shorter RTs for immediate rewards could be attributed not to the fact that the automatic response is to choose immediate rewards, but rather to the fact that such rewards are, on average, more attractive than delayed rewards, causing trials in which immediate rewards are chosen to be easier (and thus quicker) than those in which delayed rewards are chosen (see Evans & Stanovich, 2013; Hare, Hakimi, & Rangel, 2014 for a discussion).

In the present paper, we study automatic biases in intertemporal choice, with three novel insights: (1) We use the drift diffusion model (DDM) (Ratcliff, 1978; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998), a popular mathematical model of binary choice, to implement a dual process

theory of intertemporal choice, and subsequently predict choice probabilities and RTs, controlling for underlying monetary amounts and time delays; (2) The use of the DDM allows us to define the term *automatic bias* precisely, through model parameters, and we test for the presence or absence of different types of automatic bias in our dual process implementation, by estimating model parameters from data; and (3) Unlike previous work, which has primarily examined group-level patterns, we permit individual differences by allowing different participants to have different automatic biases.

Drift diffusion model

The DDM and related sequential sampling models (Peters & Büchel, 2011; Ratcliff & Smith, 2004; Townsend & Ashby, 1983) are widely considered to be accurate descriptors of both perceptual and preferential choice processes. They have also been recently shown to account for key behavioral patterns in intertemporal choice (Dai & Busemeyer, 2014; Rodriguez et al., 2014; Rodriguez et al., 2015).

Applications of the DDM to preferential choice assume that decision makers dynamically and stochastically accumulate preferences for the available rewards. The speed with which evidence is accumulated is reflected in the drift rate, v, which corresponds to the relative preference for one reward over the other in the evaluation process. The diffusion process continues until one of the decision boundaries, $+a$ or $-a$, is hit. The specific decision boundary to be hit determines the chosen reward, and time to hit the boundary plus a non-decisional time, t_{ND}, corresponds to the RT in the trial. Finally, the decision maker may favor one of the rewards at time 0 before evaluation begins. We write this starting point as z, which is the ratio of the preference at time 0 to the size of the decision boundary, a. Here v can take any value, a and t_{ND} can take any positive value, whereas z ranges from -1 to 1.

We consider intertemporal choice problems that offer decision makers a choice between an immediate reward R_I and a delayed reward R_D with some time delay T_D. We assume that the positive decision boundary $(+a)$ corresponds to the choice of the delayed reward and that the negative decision boundary $(-a)$ corresponds to the choice of the immediate reward. A schematic of the model is presented in Fig. 1. In this framework, automatic biases influence decisions in two distinct ways. Firstly, it is possible for a decision maker to begin the choice process with a starting point bias favoring the immediate or delayed reward, before seeing the choice options. This is a bias that predisposes the decision maker to respond by selecting one reward or the other (prior to evaluating the attribute values of the two rewards), and thus we will refer to it as a *response bias* (B_R). Secondly, it is possible for the drift rate to favor one or the other reward, independently of the time delays or monetary magnitudes involved in the specific choice

problem. This is a bias that reflects automatic influences on the decision maker's preferences in the evaluation process, and thus we will refer to it as an *evaluation bias* (B_E). Negative values of B_R and B_E correspond to automatic biases in favor of the immediate rewards, whereas positive values correspond to automatic biases in favor of the delayed rewards. In recent work, White and Poldrack, (2014) have shown that these two biases can be disassociated through quantitative model fits of the DDM to choice and RT data.

We consider these two biases to be automatic as they have an exogenous influence on the decision process, that is, they are unaffected by the specific rewards or time delays in the choice problem. Of course, the decision maker's evaluations of these rewards and time delays also play a key role in choice, and need to be accommodated within the model. Recent experimental results have suggested that intertemporal preferences can be modeled as a linear function of the differences in rewards and time delays in the choice problem (González-Vallejo, 2002; Scholten & Read, 2010). Such an approach has also been shown to be better than alternate ways of studying intertemporal choice within the drift diffusion framework (e.g., approaches based on discounted utility) (Dai & Busemeyer, 2014). In a linear approach, the drift rate can be written as: $v = B_E + w_{\Delta R} (R_D - R_I) + w_{\Delta T} T_D$, where $w_{\Delta R}$ and $w_{\Delta T}$ are free parameters that capture attribute weights for monetary amounts and time delays. Positive $w_{\Delta R}$ and negative $w_{\Delta T}$ indicate preferences for larger rewards and shorter delays in the evaluation process. In the supplemental materials we also present fits for an exponential model for the drift rate, with $v = B_E + \alpha (R_D e^{-kT_D} - R_I)$. Here k is a discounting factor, α is a scale factor, and B_E is again the evaluation bias. We find that the use of an exponential model for our drift rate does not alter any of our conclusions regarding automatic biases in intertemporal choice.

Note that by using the word "bias," we do not imply that the resulting behavioral tendency is irrational. In line with the terminology used in the evidence accumulation literature (e.g., White & Poldrack, 2014), this tendency merely captures decision components that are insensitive to stimuli values (the monetary amount or delayed duration). This naming is also consistent with the way in which dual process models of intertemporal choice conceptualize automatic behavioral tendencies (though unlike these models, we do allow the biases to favor the delayed reward). This is also why we specify the evaluation bias as a shift in the drift rate that adds to, and potentially distorts, the decision maker's preferences, which are $w_{\Delta R} (R_D - R_I) + w_{\Delta T} T_D$ or $\alpha (R_D e^{-kT_D} - R_I)$. Unlike preferences, the evaluation bias does not depend on the stimuli in the choice problem. That said, the evaluation bias could equivalently be seen as a component of the decision maker's preferences (i.e., as a feature of the utility function that favors the immediate or delayed reward), without altering the underlying
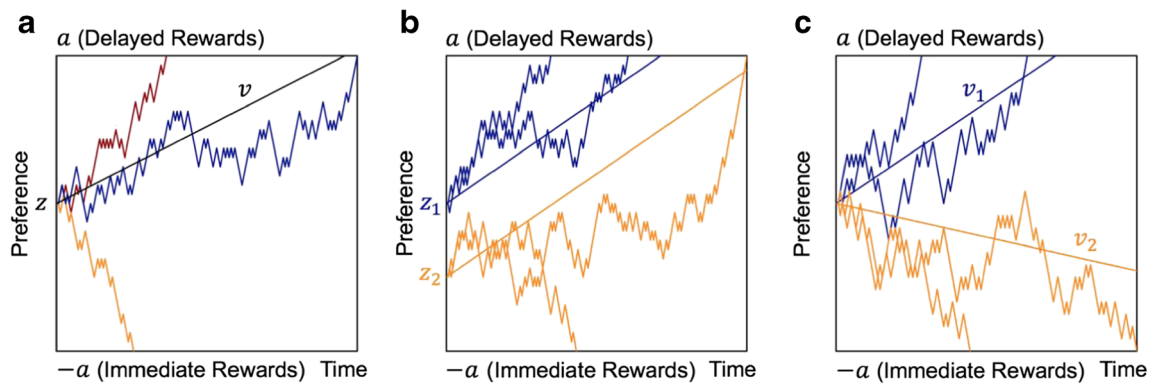
Fig. 1 (**a, b, & c**) Drift diffusion model for intertemporal choices fit in this paper. The x-axis represents time and the y-axis represents the accumulated preference state. The slopes of the solid lines represent the expected accumulation speed, i.e. the drift rate v. Each trajectory represents a hypothetical accumulation process in a single trial. An accumulation process starts with z and choice is made when a trajectory hits one of the boundaries for the first time. Here, the lower boundary ($-a$) corresponds to the choice of the immediate reward, and the upper boundary ($+a$) corresponds to the choice of the delayed reward. We assume that automatic biases can influence the starting point and/or the drift rate. A response bias shifts the starting point (panel b), and an evaluation bias shifts the drift rate (panel c). Note that even though both the evaluation bias and the resulting drift rate are negative in panel c, generally it is possible to observe a positive (negative) drift rate with a negative (positive) evaluation bias, depending on the decision weights and the stimuli

functional form of the decision model (see, e.g., Loewenstein et al., 2015; McClure et al., 2004; O'Donoghue & Rabin, 1999 for similar models). It may also be the case that a single system model that mimics the effect of an evaluation bias provides a better explanation of behavior than an evaluation bias in our dual systems implementation (we consider this possibility in detail in the discussion section and the supplemental materials). Finally, although the automatic biases specified in our model may be able to explain prominent findings such as the immediacy effect (O'Donoghue & Rabin, 1999), our focus here is not to propose a new model to account for these findings, but rather to formally instantiate and test theories of automatic processes in intertemporal choice.

Study 1: New experiment

We wished to examine automatic response and evaluation biases in intertemporal choice. Thus, we conducted an incentivized experiment offering individuals choices between immediate and delayed rewards. We collected choices and RTs for multiple choice problems from each participant, which allowed us to fit the DDMs, and infer parameters B_R and B_E, on an individual level.

Methods

We designed our experiment to have approximately 50 participants, which, when combined with extensive within-participant data provides sufficient power to test our hypotheses. Fifty-one subjects (18 females; mean age = 21.92 years, SD age = 3.57) from a paid psychology experimental participant pool eventually took part in the experiment. In each trial, participants were presented with two monetary choices, an immediate reward, R_I, that was available on the day of the experiment, and a delayed reward, R_D, with some delay, T_D. Any choice problem of this kind can be uniquely represented using three factors: R_I, T_D and R_D/R_I (a multiplier for the delayed reward). We manipulated all the three factors and chose four levels for each factor, which generated 64 unique choice problems. The resulting reward amounts ranged between \$3 and \$27.50, and the delay times ranged between 3 and 30 days. Each choice problem was repeated 10 times for each participant. Participants indicated their choice with key presses, and we were able to collect both choice and RT data. See Supplemental Materials for details about experimental stimuli and procedure.

Results

Behavioral results

We excluded one participant who admitted to have intentionally timed-out on many trials, and one participant who always chose the immediate reward. Trials with RTs less than 0.4 s, which accounted for roughly 1% of all trials, were also excluded from our analysis.

Overall, the average probability of choosing the delayed reward across participants was 52.1%. The probability of choosing the delayed reward varied significantly across participants, ranging from 1% to 99% (SD = 0.33). Twenty-five out of the 49 participants chose the immediate rewards more often than the delayed rewards. Likewise, the mean RT was 1.29 seconds. For the 49 participants, mean RTs ranged from 0.53 s to 2.17 s (SD = 0.36). Twenty-five out of the 49 participants chose immediate rewards quicker than delayed rewards.

Additional details regarding choice and RT patterns are provided in supplemental materials.

Model fits Individual-level DDM fits can help determine whether these individual differences are a product of varying attribute weights, and/or whether they are caused by diverging automatic biases. DDMs were fit to each participant's dataset separately using Bayesian parameter estimation implemented in a Python package called HDDM (Wiecki, Sofer, & Frank, 2013). To assess the fit quality, we simulated 500 samples from the posterior of the fitted model for each participant, and computed the summary statistics (probability for choosing the delayed reward and mean RTs associated with delayed and immediate rewards) over each simulated dataset for each of the participants (additional details regarding model fitting are provided in the Supplemental Materials). The summary statistics from the simulated datasets were compared to the summary statistics of the observed datasets, and the results are shown in Fig. 2(a) and 2(b). The correlations between the mean simulated statistics and the observed statistics on the participant-level were very high for choice probability (*corr* > 0.99, $t(47) = 168.96$, $p < 0.001$), RTs for delayed rewards (*corr* = 0.93, $t(47) = 17.05$, $p < 0.001$), and RTs for immediate rewards (*corr* = 0.82, $t(47) = 9.71$, $p < 0.001$), suggesting successful model fits. More results on model fits and predictions can be found in the supplemental materials.

Model parameters The distributions of posterior means of the parameters for the 49 participants are summarized in Table 1. Unsurprisingly, 42/49 participants had positive posterior means for $w_{\Delta R}$ and negative posterior means for $w_{\Delta T}$, indicating a preference for larger rewards with shorter delays in the evaluation process. Only one participant had a negative posterior mean for $w_{\Delta R}$ and seven participants had a positive posterior mean for $w_{\Delta T}$, and none of these were significant as assessed by 95% credible intervals for the parameters.

Unlike the weights for reward magnitudes and time delays, response bias (B_R) and evaluation bias (B_E) varied substantially across participants. Posterior means of B_R were negative for 26 participants and positive for 23 participants. Likewise, posterior means of B_E were negative for 32 participants and positive for 17 participants. Overall, 33 participants had a response bias that was significantly different from zero (20 negative and 13 positive), and 45 participants had an evaluation bias that was significantly different from zero (30 negative and 15 positive), as assessed by 95% credible intervals. Figure 3(a) demonstrates how the two biases varied across participants. The Supplemental Materials provide additional details regarding best-fit parameter values.

We also fit four constrained models to the individual datasets. In each of the constrained models, the effect of one key parameter was eliminated: $B_R = 0$, $B_E = 0$, $w_{\Delta R} = 0$ or $w_{\Delta T} = 0$. These models were compared using the deviance
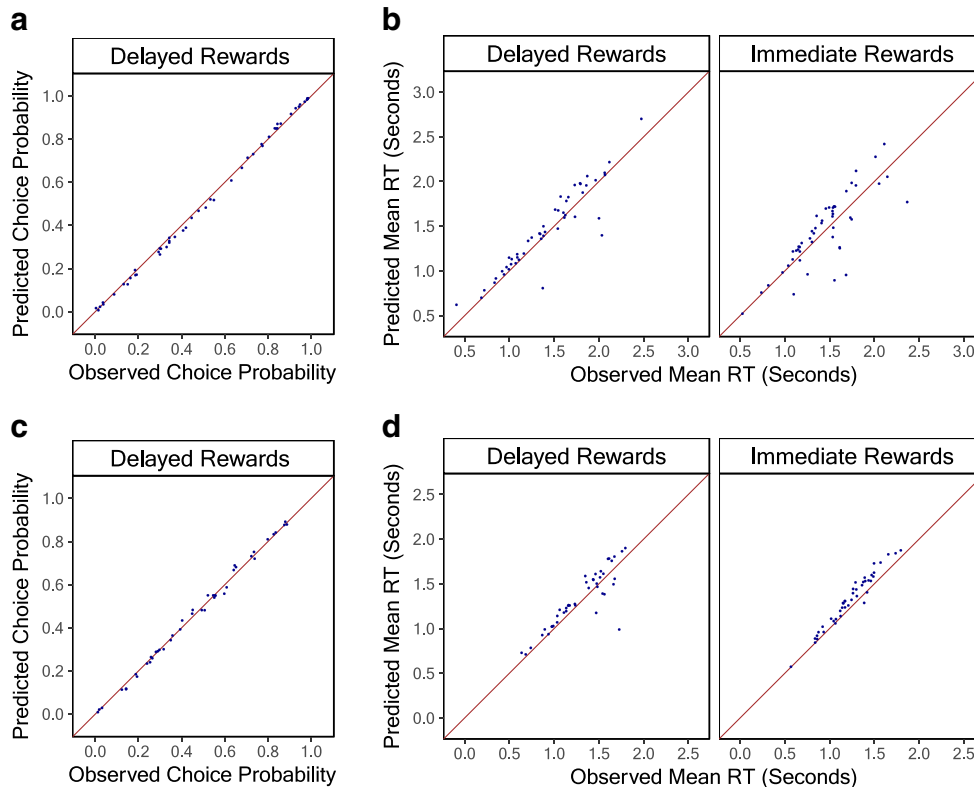

Fig. 2 Model fits for Study 1 (**a & b**) and for Study 2 (**c & d**): Plots of observed data versus predicted data of each participant. The predicted data were computed from 500 simulated samples from the posterior distributions

Table 1 Distributions of parameter posterior means across participants for Studies 1 and 2. Most participants had positive posterior means for $w_{\Delta R}$ and negative posterior means for $w_{\Delta T}$; however, the direction and magnitude of B_R and B_E varied greatly

Dataset		a	B_R	B_E	$w_{\Delta R}$	$w_{\Delta T}$	t_{ND}
Study 1	Mean	1.16	-0.03	-0.29	0.16	-0.020	0.49
	1st quantile	1.03	-0.15	-1.18	0.10	-0.034	0.40
	Median	1.15	-0.02	-0.42	0.18	-0.013	0.49
	3rd quantile	1.30	0.10	0.45	0.23	-0.004	0.57
	SD	0.20	0.17	1.18	0.08	0.021	0.12
Study 2	Mean	0.99	0.01	-0.15	0.07	-0.012	0.55
	1st quantile	0.88	-0.08	-0.60	0.02	-0.016	0.48
	Median	1.03	0.01	-0.22	0.07	-0.013	0.55
	3rd quantile	1.09	0.08	0.38	0.12	-0.008	0.62
	SD	0.15	0.15	0.93	0.06	0.006	0.11

information criterion (DIC; Spiegelhalter, Best, Carlin, & van der Linde, 2002), which measures the quality of model fits while penalizing model complexity to avoid over-fitting. Smaller DICs indicate better model fits. The mean and median DICs of the full and constrained models are summarized in Table 2. The DICs of all the constrained models were larger than those of the full models (all $p < 0.001$, as assessed by Wilcoxon matched-pair signed-rank test applied across the 49 participants), which showed that all four parameters are necessary for describing participants' choices and RTs. The model that eliminated the effect of evaluation biases ($B_E = 0$) had a larger DIC, hence a worse fit, than the model that eliminated the effect of response biases ($B_R = 0$). The median of the difference was 60.40 ($Z = 4.69$, $p < 0.001$), indicating that evaluation biases had a somewhat larger role in explaining behavior relative to response biases in the full model.

Study 2: Dataset of Hare, Hakimi, and Rangel (2014)

We also tested the robustness and generalizability of the above results by reanalyzing Hare, Hakimi, and Rangel's (2014) data, with the DDM fit on an individual level. This data consists of 43 participants, each completing 216 binary choices between $25 now and some larger amount of money available in the future (see Hare, Hakimi, & Rangel (2014) for a detailed description of methods, as well as overview of data). The structure of Hare, Hakimi, and Rangel's (2014) experiment is very similar to ours, and thus we were able to directly apply the techniques described in the above section.

We found that the DDMs described the observed choices and RTs quite well (Fig. 2(c) and 2(d)), with very high correlations between the mean simulated statistics and the observed choice probabilities ($corr > 0.99$, $t(41) = $
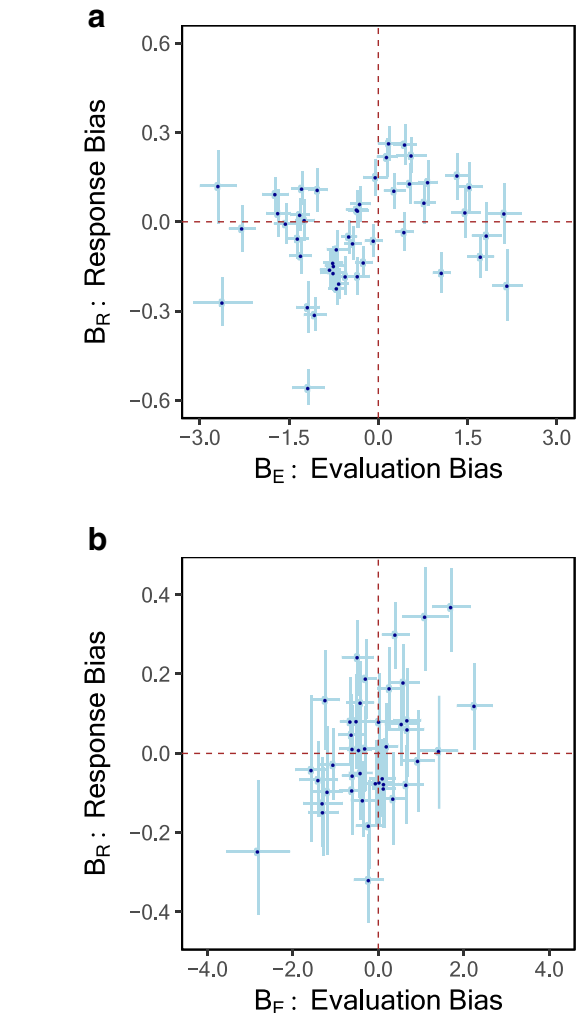
Fig. 3 (**a**) Response and evaluation biases for all participants in Study 1. All participants had either a response bias or an evaluation bias with a 95% credible interval not containing 0. However, the type, direction and magnitude of the biases varied across participants. (**b**) Response and evaluation biases for all participants in Study 2. Thirty-six of 43 participants had either a response bias or an evaluation bias with a 95% credible interval not containing 0. Solid dots indicate posterior means and error bars indicate 95% credible intervals in both plots

96.08, $p < 0.001$), RTs associated with the delayed reward ($corr = 0.87$, $t(41) = 11.46$, $p < 0.001$), and RTs associated with the immediate reward ($corr = 0.98$, $t(41) = 30.96$, $p < 0.001$).

The posterior means of the parameters recovered for each of the 43 participants are summarized in Table 1. Again, a majority of participants (37 out of 43) had $w_{\Delta R} > 0$ and $w_{\Delta T} < 0$. Only five participants had a negative $w_{\Delta R}$ and one participant had a positive $w_{\Delta T}$ (none of these were significant). In contrast, the two automatic biases varied greatly across participants. The posterior means of B_R were negative for 21 participants and B_E were negative for 24 participants. Overall, 16 participants had a response bias that was significantly different from zero (six

Table 2 Model comparisons for Studies 1 and 2: Mean and median DICs of the full and constrained models. All constrained models have significant larger DICs and thus worse fits than the full model. Eliminating B_E (the evaluation bias) increases DICs more than eliminating B_R (the response bias)

Dataset		Full	Constrained			
Study 1	Parameter Restriction	-	$B_R = 0$	$B_E = 0$	$w_{\Delta R} = 0$	$w_{\Delta T} = 0$
	Mean DIC	954.30	973.45	1067.22	1208.71	997.61
	Median DIC	1012.58	1037.82	1159.40	1317.29	1015.32
Study 2	Parameter restriction	-	$B_R = 0$	$B_E = 0$	$w_{\Delta R} = 0$	$w_{\Delta T} = 0$
	Mean DIC	293.74	298.45	311.07	347.70	348.58
	Median DIC	309.85	309.88	321.16	357.82	383.19

negative, 10 positive), and 33 participants had an evaluation bias that was significantly different from zero (21 negative, 12 positive), as assessed by 95% credible intervals (Fig. 3(b)).

As in the prior section, we fit four constrained models to the individual-level data (Table 2). The DICs of all the constrained models were larger than those of the full models (all $p < 0.02$ as assessed by Wilcoxon matched-pair signed-rank test applied across the 49 participants), indicating that all four parameters are necessary for describing participant behavior. The model with $B_E = 0$ had a larger DIC (worse fit) than the model with $B_R = 0$. The median of the difference was 7.74 ($Z = 3.06$, $p = 0.002$), again indicating that overall the evaluation biases had a more important role in explaining choice outcome variation than the response biases.

Discussion

We used the drift diffusion model (DDM), fit on an individual level, to implement dual process theories, and formally examine automatic biases, in intertemporal choice. The use of the DDM allowed us to distinguish between an automatic response bias and an automatic evaluation bias, while also controlling for preferences over monetary payoffs and time delays. In an analysis of both novel experimental data (Study 1) as well as existing data from Hare, Hakimi, and Rangel (2014; Study 2), we found that most participants demonstrated automatic biases when making intertemporal choices. However, the type, direction, and magnitude of these biases varied across participants. Additionally, model comparisons suggested that the evaluation bias played a larger explanatory role than the response bias. In the Supplemental Materials we present the results of a third experiment (Study 3) involving hypothetical payoffs, which implemented additional experimental controls. The Supplemental Materials also show model fits for a DDM with drift rates generated by an exponential discounting model. We find that all of the results presented above also emerge in Study 3, and with the exponential discounting DDM.

It is valuable to note that what we consider to be an automatic evaluation bias could instead be interpreted as a feature of the decision makers' preferences. More specifically, we are currently unable to disentangle shifts in the drift rate intercept that emerge due to biases in preference from those that emerge due to biases that exist independently of preference, as such biases are mathematically equivalent within our model (doing so would likely require other types of process data or neural data). We are agnostic about the source of the drift rate shift, and use the term evaluation bias to indicate that such a shift has an evaluative component.

Relatedly, some of our findings regarding the evaluation bias may vary with different functional specifications of the decision maker's preferences. In this paper we have used a linear model and (in the Supplemental Materials) an exponential discounting model, and have found that they yield nearly identical results. However, other models are possible, and if these models include additive shifts as part of the utility function itself, as with the quasi-hyperbolic model (O'Donoghue & Rabin, 1999), it is likely that the measured magnitude of any additional evaluation bias would weaken (of course, many dual process theories consider such additive shifts to be automatic in nature). It is useful to note that such models cannot accommodate an automatic additive shift favoring the delayed reward, as observed for many participants in our data. Additionally, the use of other utility models in the drift rate would likely not alter our inferences regarding the response bias, which corresponds to a shift in the starting point, rather than a shift in the drift rate.

Although our paper attempted to implement dual process theories of intertemporal choice, we did not provide any evidence that such theories are in fact the right way to model behavior. Our fits of the exponential model in the supplemental materials suggest that our conclusions regarding automatic biases in intertemporal choice are robust to different utility functions in dual process DDMs. Nonetheless, they do not rule out the possibility that there are single process models that could also explain our results. For this reason, what we argued were "automatic" evaluation biases, could be components of single process

theories. To examine this issue, we have attempted to fit a variant of our proposed model with a nonlinear power transformation for time delays. This model, which we outline in detail in the supplemental materials, does not include an evaluation bias, but can implicitly capture the effect of such a bias by assuming non-linearity in time. We find that there are no systematic differences in the performance of this single process model compared to our proposed model (with an automatic evaluation bias). Thus, even though dual process implementations using the DDM give rise to biases that appear automatic, there may be other explanations for these biases that do not rely on a dual process framework. Such explanations may be able to mathematically approximate the predictions of these dual process models, despite being single process models. There is need for further research that better disentangles these two explanations for our observed behavior.

Our results illustrate the value of quantitative model fitting for studying automatic biases in preferential choice. Such approaches are not only able to rigorously describe choice and RT data, they are also useful for formally representing the decision process, and thus conceptualizing the effect of different types of automatic biases on the decision process. Our results also highlight the need for individual-level analysis. Different individuals not only have different preferences; they are also likely to have different automatic tendencies. By fitting our models to each participant separately we were able to pick up important differences across our participants, which would have been obscured in group-level analysis.

Theoretically, our results resolve an outstanding question regarding automatic bias in intertemporal choice, while also raising new questions regarding individual-level differences in these biases. Although we show that many participants display automatic biases (as measured by our dual process implementation), these biases frequently favor the delayed reward, which is contradictory to the predictions of existing dual process theories of intertemporal choices (but consistent with Lindner & Rose, 2017 who find that time pressure increases patience in a large number of subjects). Further understanding the causes of these differences is an important topic for future work. Such work could benefit from newer computational models of dual processes in decision making (e.g., Diederich & Trueblood, 2018).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

Burks, S. V., Carpenter, J. P., Goette, L., & Rustichini, A. (2009). Cognitive skills affect economic preferences, strategic behavior, and job attachment. Proceedings of the National Academy of Sciences, 106(19), 7745-7750.

Dai, J., & Busemeyer, J. R. (2014). A probabilistic, dynamic, and attribute-wise model of intertemporal choice. Journal of Experimental Psychology: General, 143(4), 1489.

Diederich, A., & Trueblood, J. S. (2018). A dynamic dual process model of risky decision making. Psychological review, 125(2), 270.

Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. Perspectives on Psychological Science, 8(3), 223-241.

Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., & Weber, E. U. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. Nature Neuroscience, 13(5), 538-539.

Franco-Watkins, A.M., Rickard, T.C., & Pashler, H. (2010) Taxing executive processes does not necessarily increase impulsive decision making. Experimental Psychology, 57(3), 193-201.

Gonzalez-Vallejo, C. (2002). Making trade-offs: A probabilistic and context-sensitive model of choice behavior. Psychological Review, 109(1), 137.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. Science, 293(5537), 2105-2108.

Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. Science, 324(5927), 646-648.

Hare, T. A., Hakimi, S., & Rangel, A. (2014). Activity in dlPFC and its effective connectivity to vmPFC are associated with temporal discounting. Frontiers in neuroscience, 8, 50.

Lindner, F., & Rose, J. (2017). No need for more time: Intertemporal allocation decisions under time pressure. Journal of Economic Psychology, 60, 53–70.

Loewenstein, G., O'Donoghue, T., & Bhatia, S. (2015). Modeling the interplay between affect and deliberation. Decision, 2(2), 55.

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. Science, 306(5695), 503-507.

O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. American Economic Review, 89(1), 103-124.

Peters, J., & Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: understanding variability. Trends in Cognitive Sciences, 15(5), 227-239.

Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. Nature, 489(7416), 427-430.

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

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Computation, 20(4), 873-922.

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.

Rodriguez, C. A., Turner, B. M., & McClure, S. M. (2014). Intertemporal choice as discounted value accumulation. PloS one, 9(2), e90138.

Rodriguez, C. A., Turner, B. M., Van Zandt, T., & McClure, S. M. (2015). The neural basis of value accumulation in intertemporal choice. European Journal of Neuroscience, 42(5), 2179-2189.

Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. The Economic Journal, 117(523), 1243-1259.

Scholten, M., & Read, D. (2010). The psychology of intertemporal tradeoffs. Psychological Review, 117(3), 925.

Shamosh, N. A., DeYoung, C. G., Green, A. E., Reis, D. L., Johnson, M. R., Conway, A. R., ... Gray, J. R. (2008). Individual differences in delay discounting: relation to intelligence, working memory, and anterior prefrontal cortex. Psychological Science, 19(9), 904-911.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64(4), 583-639. Chicago

Townsend, J. T., & Ashby, F. G. (1983). Stochastic modeling of elementary psychological processes. CUP Archive.

White, C. N., & Poldrack, R. A. (2014). Decomposing bias in different types of simple decisions. Journal of Experimental Psychology: Learning, Memory, and Cognition, 40(2), 385.

Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in Python. Frontiers in Neuroinformatics, 7