# Automatic Biases in Intertemporal Choice[1]

**Wenjia Joyce Zhao (zhaowenj@sas.upenn.edu)**
University of Pennsylvania

**Adele Diederich (a.diederich@jacobs-university.de)**
Jacobs University

**Jennifer S. Trueblood (jennifer.s.trueblood@vanderbilt.edu)**
Vanderbilt University

**Sudeep Bhatia (bhatiasu@sas.upenn.edu)**
University of Pennsylvania

**Abstract**

Research on intertemporal choice has suggested that decision processes automatically favor immediate rewards. In this paper, we use a drift diffusion model to conceptualize and empirically investigate the role of these biases. Our model permits automatic biases in the response process, automatic biases in the evaluation process, as well as differential weighting for monetary payoffs and time delays. We fit our model to individual-level choice and response time data, and find that automatic biases 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.

**Keyword:** 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, Carpenter, Goette, & Rustichini, 2009; Figner et al., 2010; Hare, Camerer, & Rangel, 2009; Loewenstein, O'Donoghue, & Bhatia, 2015; McClure, Laibson, Loewenstein, & Cohen, 2004; Peters & Büchel, 2011; Shamosh et al., 2008). Trade-offs between attribute values are assumed to be processed by the control process.

One approach to testing the predictions of such 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 shorter RTs than corresponding trials in which the alternate reward is chosen (e.g. De Neys & Glumicic, 2008; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Kahneman, 2011; 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; Krajbich, Bartling, Hare, & Fehr, 2015 for a discussion). Indeed, in an analysis of intertemporal choice data with controls for option attractiveness, Krajbich et al. (2015), found no difference between RTs associated with the choice of immediate rewards and delayed rewards. This contradicts a number of existing theories and suggests that automatic processes may not systematically bias preferences in intertemporal choice.

In the present paper, we test for the existence of 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 quantitatively predict RTs controlling for attribute differences; 2) The use of the DDM allows us to define the term *automatic bias* precisely (through estimable model parameters), and we test for the presence or absence of different types of automatic bias; 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.

[1] Acknowledgment: A. Diederich received funding from Deutsche Forschungsgemeinschaft DI506/14-1 and DI506/15-1. J. S. Trueblood was supported by National Science Foundation Grant SES-1556325 and by the Alfred P. Sloan Foundation. S. Bhatia received funding from the National Science Foundation grant SES-1626825.

## 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, Turner, & McClure, 2014; Rodriguez, Turner, Van Zandt, & McClure, 2015). Importantly, these models provide a formal theory of both response probability and response time, and are potentially able to disentangle attribute weights from automatic biases that could be at play in intertemporal choice.

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 response time 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 $v$, $a$ 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 Figure 1. In this framework, automatic biases[2] can be seen to influence the decision 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 an *a priori 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* ($B_E$). In recent work, White & 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) This direct-difference attribute-wise instantiation has also been shown to be successful in accounting for intertemporal choice data when implemented in drift diffusion models (Dai & Busemeyer, 2014). Following this assumption, 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 money and time delay. Unlike the drift rate, which is influenced by both the evaluation bias and the choice attributes, the starting point is completely determined by the response bias, i.e. $z = B_R$. 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.

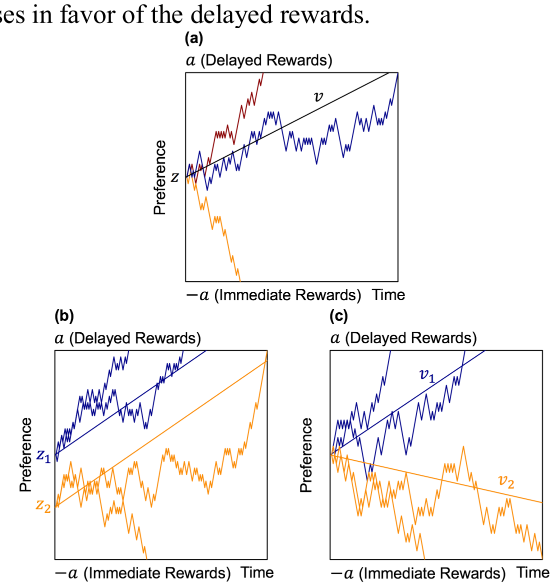


Figure 1. (a, b & c) Drift diffusion model for intertemporal choices. The x-axis represents time and the y-axis represents the preference state. The slope of solid lines represents the expected accumulation speed, the drift rate v. Each trajectory represents a hypothetical accumulation process in a single trial. (-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 v. A response bias shifts the starting point (panel b), and an evaluation bias shifts the drift rate (panel c).

[2] Note that by using the word *bias*, we do not imply that these automatic tendencies are irrational. Instead, they capture decision components that are insensitive to attribute values (the monetary amount or delayed duration). This naming is chosen to be consistent with traditions in intertemporal research, e.g. present bias. The word *attribute weighting* is used to refer to decision components sensitive to exact attribute values, manifested in weights for monetary amounts and time delays.

## Experiment

We wished to test for the presence of 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. 51 subjects (18 females; Mean age = 21.92, 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 variables: $R_I$, $T_D$ and $R_D/R_I$ (a multiplier for the delayed reward). We manipulated all the three factors and chose four levels for each factor (based on a separately collected pilot dataset), 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.

The choices were displayed side-by-side, and the position of the two rewards (left vs. right) was counterbalanced. There was also an automatic time-out after 5 seconds, after which the experiment progressed to the subsequent trial. The time limit was also determined based on the separate pilot. The trials appeared in a randomized order and were separated into 10 blocks. The experiment was incentivized, and participants received a bonus payment either immediately or after a time delay, according to their response to one randomly selected trial. Participants indicated their choice with key presses, and we were able to collect both choice and RT data.

We also measured our participants' abilities to exert deliberative control using three pre-existing questionnaires: 1. Barrat Impulsive Scale (BIS-11) (Patton & Stanford, 1995); 2. Cognitive Reflection Test (CRT) (Frederick, 2005); and 3. The Brief Self-Control Scale (BSCS) (Tangney, Baumeister, & Boone, 2004). The questionnaires were administered after participants completed the intertemporal choice task.

Table 1. Distributions of parameter posterior means across participants for our experiment and for Krajbich et al. (2015). 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 | | $a$ | $B_R$ | $B_E$ | $w_{\Delta R}$ | $w_{\Delta T}$ | $t_{ND}$ |
|---|---|---|---|---|---|---|---|
| | 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 |
| Our experiment | 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 |
| | 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 |
| Krajbich et al. (2015) | 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 |

Table 2. Model comparisons for our experiment and Krajbich et al. (2015): 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 | | | |
|---|---|---|---|---|---|---|
| Our experiment | 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 |
| Krajbich et al. (2015) | 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 |

## Results

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

Overall, the average probability of choosing the delayed reward across participant was 52.1%. This is not statistically different from 50% $(t(48) = 0.45, p = 0.66)$. The probability of choosing the delayed reward varied significantly across participants, ranging from 1% to 99% $(SD = 0.33)$. 25 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 0.53s to 2.17s $(SD = 0.36)$. 25 out of the 49 participants chose immediate rewards quicker than delayed rewards. There are large individual differences in choice patterns: some participants chose the immediate reward more often and more quickly, whereas others chose the delayed reward more often and more quickly. Individual-level DDM fits can help determine whether these individual differences are a product of varying attribute weights, and/or whether they are instead caused by diverging automatic biases.

**Model Fits.** 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 fit the models, 50,000 samples were generated for each participant, where the first 25,000 were burn-ins, and a thinning of 2 was applied. 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. The summary statistics from the simulated datasets were compared to the summary statistics of the observed datasets, and the results are shown in Figure 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.

**Model Parameters.** The posterior means of the parameters for the 49 participants are summarized in Table 1. Unsurprisingly, most 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 7 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 magnitude 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 out of 49 participants and positive for 23 out of 49 participants. Likewise, posterior means of $B_E$ were negative for 32 out of 49 participants and positive for 17 out of 49 participants. Overall, 33 participants had a response bias that was significantly different from zero, and 45 participants had an evaluation bias that was significantly different from zero, as assessed by 95% credible intervals. Figure 3(a) demonstrates how the combination of the two biases varied across participants.

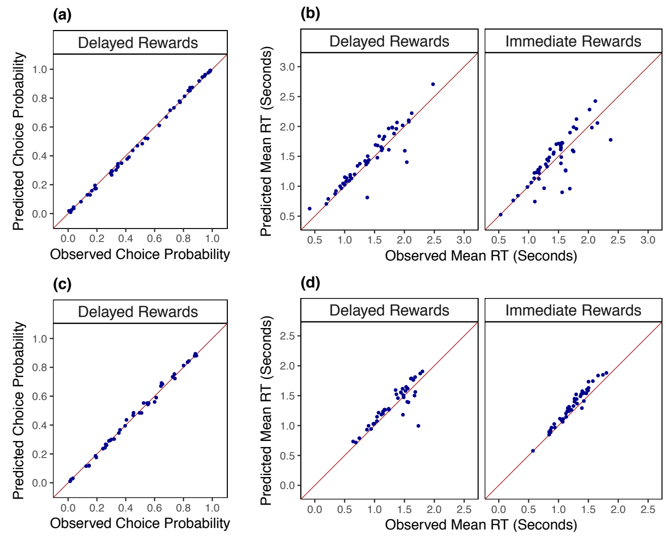


Figure 2. Model fits for our experiment (a & b) and for Krajbich et al. (2015; c & d): Plots of observed data versus predicted data of each participant. The predicted data were computed from 500 simulated samples from the posterior distributions.

We also fit four constrained models to the individual datasets. In each of the constrained models, the effect of one key parameter was eliminated, so that one of the following constraints was applied: $B_R = 0, B_E = 0, w_{\Delta R} = 0$ or $w_{\Delta T} = 0$. The mean and median DICs of the full and constrained models are summarized in Table 2. Here smaller DICs indicate better fits. The DICs of all the constrained models were larger than those of the full models (all $p < 0.001$, as assessed by Wilcoxon matched pairs 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.

**Survey Data.** We also examined how heterogeneity in the best-fit parameters related to the survey-based measures of impulsivity and control. We used our three questionnaires to generate a single composite measure of deliberative control, by performing a principle component analysis on the correlations between BIS-11, CRT and BSCS. The first principle component (PC1) explained over 61% of the variance in our questionnaire datasets. Scores for BIS-11, CRT and BSCS all loaded well to PC1, with loadings of -0.68, 0.38 and 0.63 respectively. This indicated that PC1 could serve as a suitable measure for deliberative control in our analysis. In general, larger PC1 is associated with less impulsivity and more deliberative control.

Overall, PC1 was significantly correlated with the probability for choosing the delayed reward ( $corr = 0.36, t(47) = 2.64, p = 0.011$ ). DDM parameters offer a potential process-level explanation for this correlation. However, we find that out of all the parameters fit to participant data, only $B_E$ is significantly correlated with PC1 ($corr = 0.38, t(47) = 2.85, p = 0.006$). This indicates that the predictive power of PC1 in our experiment was most likely due to its relationship with the evaluation biases.

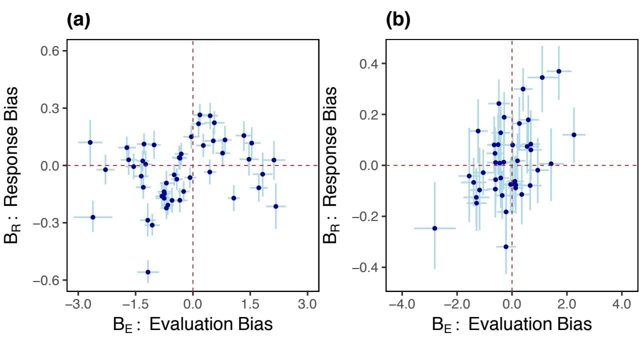


Figure 3. (a) Response and evaluation biases for all participants in our experiment. All participants had either a response bias or an evaluation bias with a 95% credible interval not containing 0. However, the type, direction and magnitude of the biases varied across participants. (b) Response and evaluation biases for all participants in Krajbich et al. (2015). 36/43 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.

## Analysis of Krajbich et al.'s (2015) Data

We also tested the robustness and generalizability of the above results by reanalyzing Krajbich et al.'s (2015) 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. There was no automatic time-out in their experiment design. (see Krajbich et al., 2015 for a detailed description of methods, as well as overview of data). The structure of this 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, with very high correlations between the mean simulated statistics and the observed choice probabilities ($corr > 0.99, t(41) = 96.08, p < 0.001$), 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 relationship between the model simulated data and the observed data is shown in Figure 2(c) and 2(d).

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 $w_{\Delta T} < 0$ when evaluating their parameter posterior means. Only 5 participants had a negative $w_{\Delta R}$ and one participant had a positive $w_{\Delta T}$, and none of these were significant as assessed by 95% credible intervals. In contrast, the two automatic biases varied greatly across participants. The posterior means of $B_R$ were negative for 21 out of 43 participants and $B_E$ were negative for 24 out of 43 participants. Overall, 16 participants had a response bias that was significantly different from zero, and 33 participants had an evaluation bias that was significantly different from zero, as assessed by 95% credible intervals. Figure 3(b) illustrates the direction and magnitude of the two biases across participants.

As in the prior section, we fit four constrained models to the individual-level data. 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 $p < 0.02$ as assessed by Wilcoxon matched pairs signed-rank test applied across the 49 participants), which showed that all four parameters are necessary for describing participant behavior. The model that eliminated the effect of evaluation biases had a larger DIC (worse fit) than the model that eliminated the effect of response biases. 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 variations than the response biases.

## Discussion

We used the drift diffusion model (DDM), fit on an individual level, to formally examine whether automatic biases play a role in intertemporal choices. The use of the DDM allowed us to distinguish between an automatic response bias and an automatic evaluation bias, while also controlling for attribute weights for monetary payoffs and time delays. In both novel experimental data as well as existing data from Krajbich et al. (2015), 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 comparison suggested that the evaluation bias played a larger explanatory role than the response bias. The magnitude of the evaluation bias displayed by participants was also significantly correlated with survey-based measures of deliberative control.

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.

This way, our results resolve an outstanding theoretical question regarding automatic biases in intertemporal choice.

Our results also highlight the need for individual-level analysis. Previous studies have shown that different people assign different weights to monetary amounts and time durations when comparing immediate vs. delayed rewards. Our study further indicates that different people are likely to vary in terms of the type and direction of their automatic biases as well. By fitting our models to each participant's data separately we were able to pick up important differences across our participants, which would have been obscured in group-level analysis. In our study, some participants display automatic biases towards the rewards that are available sooner, as predicted by existing dual process theories of intertemporal choice. However, contradicting these theories, many other participants display biases in favor of rewards with larger monetary amounts, despite their associated delays. Admittedly, these biases could be due to the choice of stimuli and experimental context. Further studies should test how these individual-level biases generalize to natural environments.

## 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.

De Neys, W., & Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. Cognition, 106(3), 1248-1299.

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.

Frederick, S. (2005). Cognitive reflection and decision making. The Journal of Economic Perspectives, 19(4), 25-42.

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. Statistical science, 457-472.

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.

Kahneman, D. (2011). Thinking, fast and slow. Macmillan.

Krajbich, I., Bartling, B., Hare, T., & Fehr, E. (2015). Rethinking fast and slow based on a critique of reaction-time reverse inference. Nature Communications, 6.

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

Matzke, D., & Wagenmakers, E. J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. Psychonomic bulletin & review, 16(5), 798-817.

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.

Patton, J. H., & Stanford, M. S. (1995). Factor structure of the Barratt impulsiveness scale. Journal of Clinical Psychology, 51(6), 768-774.

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.

Tangney, J. P., Baumeister, R. F., & Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. Journal of Personality, 72(2), 271-324.

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