To Stay or Not to Stay: The Stability-Flexibility Trade-Off in Value-Based Decision Making

Ulrike Senftleben*, Martin Schoemann, & Stefan Scherbaum

Technische Universität Dresden

Author Note

This research was supported by the German Research Council (DFG) (grant SFB 940/2 2016, project A8).

Manuscript version 1.0, 6/19/19. This paper has not been peer reviewed (it is currently under review). Please do not copy or cite without author's permission.

Author information:

Ulrike Senftleben, Department of Psychology, Technische Universität Dresden, Germany, ulrike.senftleben@tu-dresden.de; Martin Schoemann, Department of Psychology, Technische Universität Dresden, Germany, martin.schoemann@tu-dresden.de; Stefan Scherbaum, Department of Psychology, Technische Universität Dresden, Germany, stefan.scherbaum@tu-dresden.de; * corresponding author

Abstract

In order to successfully pursue a goal, humans need to maintain their behavior in the face of distractions. However, to avoid pursuing goals that became undesirable, humans also need to adjust their behavior flexibly to respective changes in the environment. The trade-off between these two forms of human functioning, cognitive stability and cognitive flexibility, is often investigated in cognitive psychology in task shielding and task switching paradigms. Here, we show that cognitive stability and flexibility also play a role in value-based decision making, as indicated by choice perseveration. We combine two experimental manipulations typical for task switching/shielding paradigms, i.e., varying the inter-trial interval and the stimulus onset asynchrony, and implement them in the context of value-based decision making in a binary choice paradigm. We predict how these manipulations will affect choice perseveration using a computational attractor model. We then test these predictions in a value-based decision game in two experiments using a sequential manipulation. Our results show that both the inter-trial interval and the stimulus onset asynchrony modulate choice perseveration as predicted by the model. We discuss how our findings extend research on cognitive stability and flexibility and their underlying mechanisms by adapting it to the domain of decision making.

Keywords: cognitive flexibility, cognitive stability, decision making, choice perseveration, attractor dynamics

To Stay or Not to Stay: The Stability-Flexibility Trade-Off in Value-Based Decision Making

The ability to maintain our behavior in the face of distractions is crucial for reaching our goals. This demand for cognitive stability, however, is complemented by the need to be able to adapt our behavior flexibly to changes in our environment – a demand for cognitive flexibility. Understanding these abilities and their underlying neural mechanisms is hence crucial to understand human behavior, which is why a large body of research has been devoted to study them, often under the term cognitive control.

Emerging from this research is the notion that the maintenance of behavior—stability—and the adaptation of behavior—flexibility—are related, such that increased stability comes at the cost of decreased flexibility, and vice versa (Goschke, 2000, 2003; Hommel & Colzato, 2017; Miller & Cohen, 2001; Monsell, 2003). This relation has been studied intensively based on tasks that use instructed task switching or task shielding (Armbruster, Ueltzhöffer, Basten, & Fiebach, 2012; Dreisbach & Goschke, 2004; Dreisbach & Haider, 2008; Fischer & Hommel, 2012; Goschke, 2000, 2013; Logan, 2003; Meiran, 2010; Ueltzhöffer, Armbruster-Genç, & Fiebach, 2015).

The assumed interrelation of stability and flexibility implies a shared underlying neural mechanism that can account for the trade-off between them. In neural network models, a key component in modeling this trade-off are attractor dynamics (Durstewitz & Seamans, 2008; Miller & Cohen, 2001). Attractors are stable states within the network that correspond to task sets or goal representations. Switching from one task set to another one thus means transitioning from one attractor state to another one. The deeper the basin of attraction of such an attractor state is, the more stable is this state and the more activation is necessary in order to transition into a different state. Thus, in terms of attractor dynamics, increased stability is represented by

deeper basins of attraction that make it less likely that the system is disturbed by noise or distractors; in turn, this would lead to inflexible behavior – it would be more difficult to switch to a different state. In contrast, increased flexibility is represented by shallower basins of attraction that make it easier to transition to a different state; in turn, this would lead to unstable behavior – states could be switched easily and, in the extreme case, randomly (see Figure 1). Thus, different configurations of a network model's attractor dynamics allow for either highly flexible or highly stable behavior; empirically, such attractor dynamics were shown to explain inter- and intra-individual variability in task switching and task shielding performance (Armbruster-Genç, Ueltzhöffer, & Fiebach, 2016; Armbruster et al., 2012; Ueltzhöffer et al., 2015).

Figure 1: Attractor landscapes for two potential attractor states. If the basin of attraction is deep (dotted black line), maintenance of the attractor state is increased and more input is necessary for a switch to the other state (i.e., increased stability at the cost of decreased flexibility). If the basin of attraction is shallow (dashed grey line), maintenance of the attractor state is decreased and small changes in input are sufficient for a switch to the other state (i.e. increased flexibility at the cost of decreased stability). Hence, the depth of the basin of attraction directly affects the trade-off between stability (maintaining the attractor state) and flexibility (switching to the other attractor state). Figure adapted from Armbruster et al. (2012).

We argue that if attractor dynamics are indeed a key mechanism that governs how stable or flexible our behavior is, this should generalize to a variety of cognitive processes that govern our behavior, not just instructed task switching or task shielding. Indeed, attractor dynamics have been proposed to be involved in a variety of cognitive processes, such as perceptual decision making (Bonaiuto, de Berker, & Bestmann, 2016; Deco & Rolls, 2003), working memory (Balaguer-Ballester, Lapish, Seamans, & Durstewitz, 2011; Durstewitz, Seamans, & Sejnowski, 2000), and language (Spivey, Grosjean, & Knoblich, 2005).

Here, we add to this by applying this reasoning from the attractor dynamics framework to the field of decision making. We have previously shown that attractor dynamics can account for intra- and inter-trial dynamics in value-based decision making (Scherbaum et al., 2016). The current study builds on our previous findings and asks if attractor dynamics can predict the modulation of stability in value-based decision making, as indicated by choice perseveration.

In our previous work, we developed a model of attractor dynamics in binary value-based decisions (Scherbaum et al., 2016). The model consists of two neural nodes that represent two options of which it has to choose one. The input to each node represents the subjective value of each of these options¹. The two nodes inhibit each other, and increase their own activity through self-excitation. The described network dynamics (lateral inhibition and self-excitation) lead to a winner-takes-all setup so that the system eventually settles into one of two potential attractor states (option 1 or option 2), at which point the decision for one of the two options is made (see Figure 2).

¹ Please note that the purpose of our model is not to explain how subjective values are generated, but rather explain how attractor dynamics shape the unfolding decision making process.

Figure 2: Attractor network model of binary choice. Upper panel: Each unit (circle) represents a population of neurons that gets the subjective value from option 1 or option 2 as input. Through the mechanisms of self-excitation and lateral inhibition, stable states emerge – attractors. Lower panel: Attractors represent the relative attractiveness of each option. If one option has a higher subjective value, input into this unit is higher and a deeper attractor is formed, making it most likely that the system will settle into that decision state (left and right panels). If both options are equally attractive (equal subjective value, middle panel), both attractors are equally deep; random noise or slight differences in input cause the system to settle into either state.

Crucially, once a decision is made and the options are withdrawn, the nodes' activation decays slowly and the two attractors disappear because the inputs (the decision options) are no longer present (see Figure 3). The residual decaying activation leads to inertia in the system. Thus, if a new decision has to be made, the residual activity of the last choice biases the new decision towards the previously chosen option, as the activation of this option's unit has not fully subsided yet. These inter-trial dynamics of our attractor model explain how choice perseveration emerges. In our previous work, we have shown that people consistently show choice perseveration as predicted by the model, to the point of sticking to the same choice even when the other option becomes gradually more attractive (Scherbaum et al., 2016)—a phenomenon

called path-dependence or hysteresis (Kelso & Schöner, 1988; Rączaszek, Tuller, Shapiro, Case, & Kelso, 1999; Tuller, Case, Ding, & Kelso, 1994).

Figure 3: Inter-trial attractor dynamics. At the beginning of the decision process, the system is in a neutral state (pink dot). When the system settles into an attractor state (red dot), a decision for the corresponding option is made (e.g. choice of option 1 in left panel). After the choice is made, the attractors disappear (middle panel). However, due to the system's inertia, residual activity of the choice that was just made remains and only slowly decreases. Thus, the system is slow to relax to the neutral start point (transition from pink dot to red dot). When a new decision trial starts before the system has completely relaxed to the neutral start point (right panel), the residual activity of the previous choice biases the system towards the previously chosen option (pink dot is closer to attractor of option 1). This process increases the likelihood of repeating the previous choice, even though the other option might be more attractive (i.e., have a higher subjective value) in the current trial.

Here, we take the next step and ask how choice perseveration can be modulated. In task switching and multitasking paradigms, common manipulations are varying the response-cue-interval (Grange, Houghton, Grange, & Houghton, 2009; Mayr & Keele, 2000; for a review, see Koch, Gade, Schuch, & Philipp, 2010) or varying the stimulus onset asynchrony (e.g., Arrington, 2008; Brisson & Jolicœur, 2007; Fischer & Hommel, 2012; Mittelstädt, Miller, & Kiesel, 2018). We adapt these manipulations to value-based decision making and simulate how they modulate choice perseveration in our attractor model. We then test the model's predictions empirically in our previously published value-based decision task (Scherbaum et al., 2016).

The first mechanism to modulate choice perseveration is the inter-trial interval (ITI) (similar to the response-cue-interval; due to the absence of a cue, we manipulate the time interval

between response and stimulus onset). According to the attractor model, choice perseveration emerges due to residual activity from the previous decision. If we introduce more time between decisions (i.e., a larger ITI), this residual activity has more time to subside and thus choice perseveration decreases. Hence, according to our attractor model, we expect choice perseveration to decrease with increasing ITI (please see the Appendix for an implementation of this effect into a computational model).

The second mechanism to modulate choice perseveration is to introduce a stimulus onset asynchrony (SOA), so that one choice option is presented earlier than the other one. We use this SOA to bias the system towards the alternative, previously unchosen option. If both options are presented at the same time, the respective attractors emerge simultaneously. However, if the previously unchosen option is presented first, then the corresponding attractor also emerges first. This creates a bias towards that option, thus counteracting the bias towards the previously chosen option (i.e., counteracting choice perseveration). According to our attractor model, a larger delay between the stimulus onsets results in a larger bias against the previously chosen option. Thus, we expect choice perseveration to decrease with increasing SOA.

Computational Modeling

For the sake of this introduction, we derived our hypotheses verbally from the attractor model. In addition, we implemented both the ITI and the SOA manipulation in our previously published computational model (Scherbaum et al., 2016). Our model simulations showed that choice perseveration decreases with increasing ITI and increasing SOA. Thus, our hypotheses can be mathematically derived from the model. For reasons of comprehensibility, please see the Appendix for a description of the computational modeling.

Experiment 1

In Experiment 1, we used a non-verbal value-based decision game (Scherbaum et al., 2016; Scherbaum, Dshemuchadse, Leiberg, & Goschke, 2013; Scherbaum, Haber, Morley, Underhill, & Moustafa, 2018) in which participants navigated an avatar in a two-dimensional virtual environment to collect rewards of varying size. Participants performed two different variations of the decision game (ITI in Experiment 1a; SOA in Experiment 1b) in two separate sessions. The order of the two sessions was balanced across participants. Since both versions of the decision game were measured independently and due to the differences in the timing manipulation, we will report the results as of two sub-experiments 1a (ITI) and 1b (SOA).

Data Statement

The data for all experiments of this study are openly available and can be downloaded at the Open Science Framework². Data processing was carried out in Matlab R2015a and statistical testing was carried out using JASP 0.9.2 (JASP Team, 2018); all respective scripts can also be downloaded at the link above.

Experiment 1a: ITI

In Experiment 1a, we aimed to validate the prediction that longer time intervals between trials reduce choice perseveration.

Methods.

Participants. We planned to measure 80 participants to achieve a power of .85 with an effect size of $d = .35$ (as calculated with G*Power [Faul, Erdfelder, Lang, & Buchner, 2007]). We recruited 87 participants via the ORSEE-based database (Greiner, 2015) of the Department of Psychology of the TU Dresden, Dresden, Germany. Seven participants had to be excluded due to

² osf.io/hg8m2 (private link, will be made public upon acceptance in a peer-reviewed journal)

technical issues that led to incomplete data collection ($n = 3$) or choice behavior that prohibited our sequential manipulation (i.e., participants that always chose the same option; $n = 4$). Of the remaining 80 participants, 56 were female and the mean age was 22.01 years ($SD = 3.22$ years). All participants had normal or corrected-to-normal vision and color vision. Participants gave informed consent and received reimbursement of 5€ per hour as well as the money they collected within the decision game. The study was approved by the local ethics committee and was performed in accordance with the Declaration of Helsinki.

Figure 4: Value-based decision game. Participants chose between two coins in each trial. One coin was smaller in value, but nearer to the avatar (small/near option); the other coin was larger in value, but farer away from the avatar (large/far option). To make a choice, they could move the avatar field-by-field to the desired coin by clicking into neighboring fields (outlined in white). Trees (in dark green) were included for better spatial orientation and did not restrict movement. The remaining time within each was displayed above the avatar.

Apparatus and Stimuli. The decision game was presented on a 17-inch screen (1280 x 1024 pixels, 85 Hz). The experiment was controlled using the Psychtoolbox version 3 (Brainard, 1997; Pelli, 1997) in Matlab 2006b (the Mathworks, Inc.) on a Windows XP SP2 personal computer. Participants used a Logitech USB mouse to make their responses. The decision game consisted of a two-dimensional world of 20 x 20 fields, with each field consisting of 50 x 50 pixels. The avatar that participants controlled was a red circle that could be moved freely from field to field by clicking with the computer mouse into either vertical or horizontal adjacent fields outlined in white (see Figure 4). The reward stimuli used within this paradigm were circular coins with a diameter matching that of a single field, i.e. 50 pixels. The coins were gold with the value written inside each coin in red. Throughout the whole task, the remaining time within each block was displayed above the avatar.

Procedure. The decision task consisted of four blocks of 8 min. Participants could determine the length of the break in between two blocks themselves. Participants were instructed to collect as much credit as possible; they were informed that they would receive the money they collected in the game at the end of the experiment. In each trial of the decision game, two coins appeared at the same time. Participants had to choose between these two coins by moving the avatar field-by-field until it reached the desired coin. Upon reaching that coin, both coins disappeared and the accumulated credit collected so far appeared. The next trial started with the appearance of two new coins. The ITI ranged from 0 to 2 s (see Design). During the ITI, the mouse cursor was locked to the position of the avatar; it could only be moved again when the next trial started and new coins appeared.

Design. The decision task was split into two different parts: a measurement part intended to measure participants' individual subjective values (one block of 8 min), and an experimental part intended to capture choice perseveration (three blocks of 8 min). For both parts, trials were constructed so that one coin had a smaller value but was nearer to the avatar—the small/near (SN) option—and the other coin had a larger value but was further away from the avatar—the large/far (LF) option. The difference in distance between both options (distance from the avatar to the LF option minus distance from the avatar to the SN option; called *distance* from here on) varied between one to twelve fields. Coins were situated in a way that moving towards one coin increased the distance to the other one. One credit equaled 0.01 cent.

The measurement part of the task consisted of one block of trials with a wide range of combinations of distances and values. The SN option was either two or three fields away from the avatar, the LF option was either one, four, eight, or twelve fields further away from the avatar than the SN option. In each trial, the value of the LF option was drawn randomly from a range of 65 to 85 credits, whereas the value of the SN option randomly varied between 20% to 95% of the value of the LF option. During the measurement block, the ITI was kept constant at 1 s.

We used the choice data of this measurement block to estimate participants' subjective values in the form of indifference points. Indifference points describe the specific value ratio where both options are equally attractive to the participant, i.e. the probability of choosing either the SN option or the LF option is 50% respectively. For each participant, we calculated indifference points for each distance between the SN and LF option (one, four, eight, twelve fields) by fitting a logistic function to the choice data and taking the point of inflection (see Scherbaum et al., 2013). Based on these four indifference points for the four distances used in the measurement block, we interpolated the indifference points for the remaining distances between

one and twelve³. As a result, we had estimates of indifference points for each distance between one and twelve, which we then used to construct trials in the subsequent experimental part.

The experimental part consisted of three blocks of 8 min. We used a sequential manipulation to capture choice perseveration. In each sequence, twelve trials were presented where the distance between the SN and LF option was incrementally increased (ascending sequence) or decreased (descending sequence). Thus, for ascending sequences the first trial entailed a distance of one field between the SN and LF option and the twelfth trial entailed a distance of twelve fields between the SN and LF option (vice versa for descending sequences). The order of the sequences was randomized throughout the experiment. Within each sequence, the values of the coins were kept constant. The values were chosen based on participants' subjective values from the measurement block, so that they represented participants' indifference point at the middle of the sequence (distance of six fields).

This implied that the lowest and highest distances led to choice scenarios with a high distance to the indifference point. That was important because it ensured that all participants could show choice perseveration. Specifically, on the example of ascending sequences, this manipulation caused participants to initially prefer the LF option (due to its larger value and only one field of additional distance); participants then found both options approximately equally attractive in the middle of the sequence, and would prefer the SN option at the end of the sequence (when the distance to the LF option was maximal at twelve fields). We thus expected participants to prefer one option in the beginning, stick to that option due to choice perseveration

³ We interpolated in the following ways: If the measured indifference points were monotonically decreasing, we interpolated cubically using the “pchip” method within the function `interp1` in Matlab. If the measured indifference points were not monotonically decreasing, we fit them to an exponential and a hyperbolic model. The model that provided the better fit (as quantified by standard errors) was used to interpolate the indifference points.

even as the other option becomes more attractive, and finally switch to the other option towards the end of the sequence.

In order to modulate choice perseveration, we manipulated the ITI. The ITI was selected to be either 0, 1, or 2 s. These different ITIs were balanced between sequences, but within each sequence, the ITI was kept constant.

Results. In the experimental block, participants completed 490.44 trials on average ($SD = 80.48$ trials), hence they repeated each scenario (2 sequence directions x 3 ITIs x 12 trials per sequence) approximately seven times. Data from the measurement block can be found in the supplementary materials.

First, in order to investigate modulations of choice perseveration we needed to ensure that our sequential manipulation indeed led to choice perseveration. As our model predicted maximal choice perseveration for the ITI of 0 s, we analyzed choice behavior in that condition. We ran a repeated measures ANOVA on participants' percentage of LF choices with the factors Direction and Distance (for ITI = 0 s; see Figure 5, left panel). As expected, we found a significant main effect of Direction, $F(1,79) = 75.77, p < .001, \eta^2 = .49$, indicating that choice behavior depended on the sequence direction, indicating choice perseveration. The significant interaction Direction x Distance, $F(11,869) = 25.34, p < .001, \eta^2 = .24$, showed choice perseveration was larger in medium distances (i.e., the middle of the sequences) than at small or large distances (start or end of the sequences, where either the LF or SN option has a higher subjective value). The analysis also revealed a significant main effect of Distance, $F(11,869) = 248.73, p < .001, \eta^2 = .76$, indicating that participants discounted the rewards by their distance, as has been shown before (Scherbaum et al., 2013). Taken together, our sequential manipulation successfully led to choice perseveration.

Figure 5: Mean large/far (LF) choice percentages for each inter-trial interval (ITI). Choice perseveration (calculated as the difference between ascending (orange) and descending (blue) sequences) decreases with increasing ITI. Note: *** $p_{\text{holm}} < .001$; error bars indicate standard errors.

Next, we analyzed if and how ITI impacted choice perseveration. We expected longer ITIs to reduce choice perseveration. We looked at the choice behavior across ascending and descending sequences for each ITI (see Figure 5). For each participant, we calculated a perseveration value for each ITI. For this perseveration value, we first calculated the mean percentage of LF choices for ascending sequences by averaging across all distances for each ascending sequence and then taking the overall mean (separately for each ITI). We repeated this procedure for descending sequences. Then, we subtracted the resulting mean percentage of LF choices for descending sequences from the mean percentage of LF choices of ascending sequences, separately for each ITI. The resulting three differences represented the mean perseveration values for each ITI for a participant, with larger values indicating larger choice perseveration. We used this mean perseveration value in a repeated measures ANOVA with the

factor ITI to test for modulation of choice perseveration. As expected, the analysis revealed a significant main effect of ITI on perseveration, $F(2,158) = 10.68$, $p < .001$, $\eta^2 = .12$, showing that ITI impacts perseveration. Thus, in line with the attractor model's predictions, longer ITIs indeed reduced choice perseveration (see Figure 5). Post-hoc tests (Holm corrected) revealed significant differences in choice perseveration between no ITI (0 s) and the longest ITI (2 s), $p_{\text{holm}} < .001$, $d = .53$, and between the medium (1 s) and longest ITI, $p_{\text{holm}} < .001$, $d = .42$. This further indicates that the ITI modulates perseveration.

Discussion. Experiment 1 established the choice perseveration effect and showed that choice perseveration decreases with increasing ITI, as predicted by our attractor model. According to the attractor model, longer time in between trials gives the system more time to relax between trials, thus reducing the influence of the previous choice. This leads to reduced choice perseveration, as found in experiment 1.

Experiment 1b: SOA

In Experiment 1b, we aimed to validate predictions regarding the biasing of the attractor state by introducing a stimulus onset asynchrony (SOA). Stimulus onset asynchrony means that the reward options were not presented at the same time, but with a temporal delay. This temporal delay, or SOA, should create a bias towards the option that was shown first, because that option's attractor started to emerge before the second option was shown. Hence, presenting the previously unchosen option first should create an advantage for that option, therefore decreasing the likelihood of repeating the previous choice and in turn decreasing choice perseveration.

Methods. Participants, apparatus and stimuli, procedure and design were the same as in Experiment 1a, except for changes in design that are described in the following.

Design. The task was again split into a measurement part to measure participants' subjective values, and an experimental part capturing choice perseveration (see Experiment 1a). In divergence from Experiment 1a, the ITI was kept constant throughout the whole task at 1 s. We manipulated the SOA of the two coins in each trial so that one coin appeared either at the same time (0 ms delay), 100 ms, or 200 ms before the other coin. Which coin appeared first (SN or LF option) depended on the direction of the sequence. In ascending sequences, where the LF option is preferred in the beginning, the SN option appeared first; in descending sequences, where the SN is preferred in the beginning, the LF option appeared first. Hence, the SOA was intended to create a bias against the initially more attractive option and thus counteract the repeated choice of the initially more attractive option, leading to reduced choice perseveration.

Results. In the experimental block, participants completed 487.15 trials on average ($SD = 81.45$ trials), hence they repeated each scenario (2 sequence directions x 3 SOAs x 12 trials per sequence) approximately seven times. Data from the measurement block can be found in the supplementary materials.

First, we checked if our manipulation successfully led to choice perseveration. We looked at participants' choices in the condition without SOA, as our model predicted maximal perseveration for this condition—this also serves as a replication for the same condition in Experiment 1a. Similarly to data analysis in Experiment 1a, we performed a repeated-measures ANOVA on the percentage of LF choices with the factors Direction and Distance (for $SOA = 0$ ms, see Figure 6, left panel). As predicted, we found a significant main effect of Direction, $F(1,79) = 46.78, p < .001, \eta^2 = .37$, indicating choice perseveration. Like in

Experiment 1a, we found a significant main effect of Distance, $F(11,869) = 303.38, p < .001, \eta^2 = .79$, indicating discounting. We also again found a significant interaction effect Direction x Distance, $F(11,869) = 16.62, p < .001, \eta^2 = .17$; perseveration was larger for medium distances than for smaller or larger distances. Hence, our manipulation again successfully induced choice perseveration.

Next, we investigated if this perseveration effect was modulated by variations in SOA. We expected longer SOAs to reduce choice perseveration. In order to test this, we calculated a perseveration measure following the same steps described in Experiment 1a. The resulting measure represents the difference of mean LF choice percentages between the ascending and descending sequence; the larger this measure, the larger the perseveration effect. We performed a repeated-measures ANOVA with the factor SOA. As expected, the main effect of SOA on perseveration was significant, $F(2,158) = 44.22, p < .001, \eta^2 = .35$. Post-hoc comparisons (Holm corrected) revealed that there were significant differences in choice perseveration between no SOA (0 ms) and 100 ms SOA, $p_{\text{holm}} < .001, d = .79$, as well as between no SOA and 200 ms SOA, $p_{\text{holm}} < .001, d = .84$. There were no significant differences in perseveration between the 100 ms SOA and the 200 ms SOA, $p_{\text{holm}} = .30$. This implies that perseveration was completely disrupted by SOA. As Figure 6 shows, perseveration was present when no SOA was introduced; for both the 100 ms and 200 ms SOA, perseveration was not present. Thus, while SOA did affect perseveration, the impact seemed to be stronger than we expected. Additionally, Figure 6 surprisingly shows higher LF choice percentages for descending than ascending sequences for the smallest distance (one field), a pattern which did not occur in Experiment 1a.

Figure 6: Mean large/far (LF) choice percentages for each stimulus onset asynchrony (SOA). Choice perseveration (calculated as difference between ascending (orange) and descending (blue) sequences) occurred in the absence of SOA (left panel), but not for SOAs of 100 and 200 ms. Note: *** $p_{holm} < .001$; error bars indicate standard errors.

Discussion. In Experiment 1b, we investigated the impact of SOA on choice perseveration. We expected a greater SOA, that is, a greater time delay between the presentation of the SN and LF option, to lead to weaker choice perseveration.

Experiment 1b yielded two main findings: the replication of the choice perseveration effect in absence of SOA, and the disruption of choice perseveration due to SOA. The replication of the perseveration effect in the condition without SOA further supports the validity of the effect and the attractor model, in line with our expectations and the model simulation. The disruption of perseveration due to SOA, however, is somewhat surprising. While increasing SOA did reduce perseveration, just as in our simulation, this reduction was not a modulation but a complete eradication of the perseveration effect. In our model simulation, small SOAs led to a small reduction in perseveration and large SOAs led to a large reduction in perseveration. Thus,

because SOA reduced perseveration more strongly than expected, we concluded that our SOAs might have been too large to allow for a true modulation of perseveration. Further, we hypothesized that the strong reversal effects at the beginning of a sequence might be caused by a bias from the previous sequence in combination with the rather strong SOA manipulation, which should also be addressed in an adapted study design.

Experiment 2

With Experiment 2, we aimed to investigate the effects of smaller SOAs on choice perseveration. In Experiment 1b, the perseveration effect was completely eradicated by SOAs of 100 ms or 200 ms. We hypothesized that our chosen time delays of 100 and 200 ms might have been too long to capture this modulation. Thus, in Experiment 2 we used shorter time delays to allow for a more subtle modulation of the perseveration effect. Furthermore, we aimed to remove the strong reversal effects at the beginning of a sequence. We reasoned that these effects might be biasing effects from previous sequences. In fact, in the computational simulation, we had already included empty trials between sequences to avoid such biases (see Appendix). We hence included such buffer trials also into the experiment, to isolate each sequence from the preceding one.

Methods

Procedure, apparatus and stimuli were the same as in Experiment 1b. The design was kept mostly similar to the design of Experiment 1b; the changes we made are discussed below under Design.

Participants. We based our sample size estimation on the effect size of the perseveration effect difference between the 0 ms SOA and the 100 ms SOA condition, $d = .79$; with a power of .85, this yielded a sample size of 17 participants. We aimed for a larger sample size than that because we reasoned the effect size might be smaller due to the smaller SOAs. We thus recruited 25 participants via the ORSEE-based database of the Department of Psychology of the TU Dresden, Dresden, Germany. Eighteen participants were female, the mean age was 25.2 years ($SD = 6.87$ years). All participants had normal or corrected-to-normal vision and color vision. Participants gave informed consent and received reimbursement of 5€ per hour as well as the money they collected within the decision game. The study was approved by the local ethics committee and was performed in accordance with the Declaration of Helsinki.

Design. The design closely resembled the design of Experiment 1b. The task was again split into a measurement and an experimental part. However, in Experiment 2 we added an additional block in the experimental part, so that participants performed four blocks à 8 min in the experimental part (as opposed to three blocks in Experiment 1b). This was necessary because we used four instead of three SOAs. The SOAs in Experiment 2 were smaller than in Experiment 1b, in order to create more subtle effects on participants' choice behavior. The first coin appeared either 0, 33, 66, or 100 ms before the second coin. As in experiment 1b, SOAs were randomized across sequences and kept constant within each sequence.

Furthermore, we introduced two buffer trials between sequences. The purpose of the buffer trials was to prevent the first choice of the new sequence to be biased by the last choice of the former sequence. We constructed the buffer trials so that either the SN or the LF option was clearly more attractive, based on participants' individual indifference points. The direction of the bias (SN/LF) of these buffer trials was consistent with the initially more attractive option in the

new sequence. Thus, the two buffer trials biased participants towards choosing the initially more attractive option at the beginning of each sequence, hence maximizing the chance of observing perseveration effects.

As another difference to Experiment 1b, we shortened the choice sequences from twelve trials to nine, encompassing the distances 2 to 10 (instead of 1 to 12 in Experiment 1b). The shortened sequences and the additional block of eight min were implemented to compensate for the additional trials needed due to the extra SOA and the two buffer trials between sequences. Overall, participants completed 732.44 trials on average ($SD = 211.92$ trials). This allowed them to run through every possible condition approximately eight times, as they needed 88 trials to complete every condition once (72 sequence trials: $9 [\text{distance}] \times 4 [\text{SOA}] \times 2 [\text{direction}]$; plus 16 bias trials: $2 [\text{bias trials}] \times 8 [\text{sequence conditions}]$).

Results

In the experimental block, participants completed 732.44 trials on average ($SD = 211.92$ trials), hence they repeated each scenario (2 sequence directions $\times$ 4 SOAs $\times$ 11 trials per sequence) approximately eight times. Data from the measurement block can be found in the supplementary materials.

First, we again checked for choice perseveration. Similarly to data analysis in Experiment 1b, we performed a repeated-measures ANOVA on the percentage of LF choices with the factors Direction and Distance (for $\text{SOA} = 0$ ms, see Figure 7, left panel). As expected, we found evidence for choice perseveration, as indicated by the significant main effect of Direction, $F(1,24) = 21.06, p < .001, \eta^2 = .47$, and by the significant interaction Direction $\times$ Distance, $F(8,192) = 3.25, p = .002, \eta^2 = .12$. We also replicated the discounting effect of previous experiments, as indicated by the main effect of Distance, $F(8,192) = 29.17, p < .001, \eta^2 = .55$.

Next, we wanted to investigate if increasing SOA reduced the perseveration effect. We calculated the perseveration measure similarly to analysis in Experiments 1a and 1b. This perseveration measure contains the difference between the mean percentage of LF choices in ascending and descending sequences. We performed a repeated-measures ANOVA on the perseveration measure with the factor SOA. As expected, SOA had a significant main effect on perseveration, $F(3,72) = 7.47, p < .001, \eta^2 = .24$ (see). Post-hoc comparison (Holm-corrected) of the four SOAs revealed a significant difference in perseveration between the smallest (0 ms) and biggest SOA (100 ms), $p_{holm} = .006, d = .74$, as well as trends towards significance for differences between several SOAs (see Table 1).

Figure 7: Mean large/far (LF) choice percentages for each stimulus onset asynchrony (SOA). Choice perseveration (calculated as the difference between ascending (orange) and descending (blue) sequences) decreases with increasing SOA. Note: Error bars indicate standard errors.

*Table 1**Post Hoc Comparisons – Stimulus Onset Asynchrony*

		Mean Difference	SE	<i>t</i>	<i>p</i> _{holm}	Cohen's <i>d</i>
0 ms	33 ms	.049	.019	2.595	.069	.52
	66 ms	.059	.024	2.500	.069	.50
	100 ms	.109	.029	3.714	.006	.74
33 ms	66 ms	.010	.020	0.493	.627	.10
	100 ms	.060	.023	2.656	.069	.53
66 ms	100 ms	.050	.023	2.198	.076	.44

Discussion

Experiment 2 shows that choice perseveration is reduced by increasing SOA, as predicted by the model. These results augment the findings from Experiment 1b, where SOA completely disrupted choice perseveration instead of reducing it. We hypothesized that the time delays used for the SOAs in Experiment 1b were too long (100 ms and 200 ms) and led to a ceiling effect. We thus used smaller SOAs in Experiment 2 in order to allow for a more subtle modulation of choice perseveration. We indeed found evidence for a more gradual reduction of perseveration for SOAs of 33 ms and 66 ms. Taken together with the eradication of the perseveration effect for longer SOAs of 100 ms and 200 ms in Experiment 1b, we conclude that SOA reduces choice perseveration as expected, but only for SOAs in a small time window (< 100 ms). Additionally, the buffer trials in between sequences successfully reduced carry-over effects from one sequence to another, as we expected based on the attractor model.

General Discussion

In this study, we investigated the trade-off between cognitive stability and cognitive flexibility which is usually studied in the cognitive domain by using cognitive tasks, e.g., task switching. Here, we investigated this trade-off in value-based decision making, using a binary choice task. We implemented a sequential manipulation in this binary choice task that allowed us

to measure choice perseveration: How long do participants stick to an initial choice before switching to a more attractive alternative option? We used two manipulations that are typically studied in the cognitive domain, namely the length of the inter-trial interval (ITI, i.e., the time between a choice and the presentation of the new options), and the stimulus onset asynchrony (SOA, i.e., introducing a time delay between presenting the first and the second option). We found that choice perseveration increased with longer ITIs and with larger SOA. These effects match well with predictions from computational models of attractor dynamics.

Choice perseveration in value-based decision making

We found stable choice perseveration effects in all experiments. Specifically, we designed our experiments to include one condition that was directly comparable across all experiments, thus enabling us to directly replicate choice perseveration within this study and in comparison to the original findings of perseveration in the same binary choice paradigm (Scherbaum et al., 2016). Further, the two manipulations we adapted from task switching paradigms—ITI and SOA—modulated choice perseveration as expected. Overall, this speaks to the robustness of this effect.

Our findings on choice perseveration show that decision making is influenced by the decision history and not only by the properties of the current choice situation. Notably, such effects of the decision history seem to stem from the inertia of the cognitive system, intrinsic dynamics that lead to carry-over effects from trial to trial. Such dynamics are neglected in linear models of decision making, for example when modeling the accumulation of evidence for a choice option *within* one trial (e.g., in linear-ballistic accumulator [Brown & Heathcote, 2008] or drift diffusion models [Ratcliff, 1978]). While such models include parameters that could be used to fit biases across trials, they do not incorporate the dynamics that allow these models to predict

such effects. In contrast, these dynamics across trials naturally emerge from models incorporating attractor dynamics, for example the leaky competing accumulator model (e.g., Usher & McClelland, 2001) or biophysical attractor models (e.g., Hämmerer, Bonaiuto, Klein-Flügge, Bikson, & Bestmann, 2016; Wang, 2008). In fact, such models have been successful in predicting inter-trial dynamics and sequence effects in the cognitive domain, such as working memory maintenance (Balaguer-Ballester et al., 2011; Durstewitz et al., 2000) or task switching (Armbruster-Genç et al., 2016; Armbruster et al., 2012; Gilbert & Shallice, 2002; Ueltzhöffer et al., 2015). Efforts on modeling inter-trial effects are so plentiful in research of cognitive control processes because inter-trial dynamics are often an integral part of answering these research questions. How do people adjust their behavior after an error or a conflict trial? What makes people successful at switching between tasks? With these research questions, the trial sequence or decision history is crucial and a large body of research has been dedicated to investigate different types of sequence effects; for example, the Gratton effect (Gratton, Coles, & Donchin, 1992), the finding that congruency effects are smaller after an incongruent trial, has sparked a debate that is still ongoing decades later (for reviews, see Braem, Abrahamse, Duthoo, & Notebaert, 2014; Schmidt, 2013).

As this research on cognitive control shows, the choice sequence or history (i.e., what happened in the previous trial) can clearly impact the decision making process in the present trial. In our view, this makes adopting such a focus on inter-trial dynamics and sequential effects to the field of decision making a very worthwhile endeavor. Indeed, the importance of inter-trial dynamics in decision experiments has been pointed out previously (see e.g., Cho et al., 2002; Fründ, Wichmann, & Macke, 2014; Gao, Wong-Lin, Holmes, Simen, & Cohen, 2009), and recently a growing body of research has demonstrated these kind of effects in value-based

decision making (e.g., Alós-Ferrer, Hügelschäfer, & Li, 2016; Hämmerer et al., 2016; Padoa-Schioppa, 2013). This is in line with our own findings on the inter-trial dynamics of choice perseveration; thus, we conclude that the choice history reliably influences decision making. We therefore argue that, instead of neglecting dynamics across trials, research should focus on how such dynamics unfold across trials in order to gain new insights into the underlying decision making process and their presumably non-linear dynamics.

The Stability-Flexibility Trade-Off

The trade-off between stability and flexibility is often investigated in the context of switch costs, such that increased flexibility means it is easier to switch from one task to another, thus decreasing switch costs, while increased stability means it is more difficult to switch from one task to another, thus increasing switch costs. This trade-off has been demonstrated empirically in the context of task switching (Armbruster et al., 2012; Ueltzhöffer et al., 2015) and set shifting (Dreisbach & Goschke, 2004), but also in a variety of other tasks requiring different degrees of stability and flexibility, such as cognitive search tasks (e.g., Hills, Todd, & Goldstone, 2010; Mekern, Sjoerds, & Hommel, 2019), the Simon task (Plessow, Fischer, Kirschbaum, & Goschke, 2011), and dual-task paradigms (Fischer & Hommel, 2012; Zwosta, Hommel, Goschke, & Fischer, 2013).

Here, we investigated similar effects of stability-flexibility trade-offs, but in the context of decision making. Our findings show that the same experimental manipulations used to manipulate stability in instructed task switching paradigms (i.e., cue-stimulus-interval [e.g., Altman, 2005; Koch & Allport, 2006; Meiran, 1996; Meiran, Chorev, & Sapir, 2000] and stimulus onset asynchrony [e.g., Arrington, 2008; Brisson & Jolicoeur, 2007; Fischer & Hommel, 2012]) can be used to manipulate stability in value-based decision making. One possible

interpretation of our findings is that the mechanisms underlying stability of behavior seem to be similar across seemingly distinct higher cognitive functions. Indeed, flexible adjustment of stability is a ubiquitous requirement for almost every cognitive function or process, from controlling our own movement over learning to goal-directed behavior. Thus, it seems plausible that our brain has a common underlying mechanism allowing for this stability-flexibility trade-off. If there is one common mechanism, the individual configuration of the stability-flexibility trade-off should be stable across different tasks and domains. There is some evidence for such a general trait bias towards flexibility or stability (for a review, see Hommel & Colzato, 2017). In that vein, our finding that the stability-flexibility trade-off can be modulated with the same experimental manipulations across different domains could be interpreted in the context of a common underlying mechanism.

However, it is debated if and how well stability-flexibility biases truly generalize over tasks: While some studies report that measures of stability and flexibility are correlated across tasks (e.g., Hills, Todd, & Goldstone, 2008; Hills et al., 2010), other studies find the opposite (e.g., Mekern et al., 2019). This would go against the notion of a general stability-flexibility trait bias. However, there is evidence for a state-dependent stability-flexibility bias. For example, the stability-flexibility trade-off has been shown to be modulated by affect (Dreisbach & Goschke, 2004; Goschke & Bolte, 2014) and by task demands (Mekern et al., 2019). As discussed by Mekern and colleagues (2019), people might adjust their individual stability-flexibility bias according to situational constraints and demands, and this bias might differ between tasks even if an underlying trait bias exists.

Thus, while our findings demonstrate that stability-flexibility trade-offs can be found across different domains and seem to behave somewhat similarly, it is unclear if this is due to a

common underlying mechanism (such as a trait bias), or if the stability-flexibility trade-off is a general principle with distinct underlying mechanisms depending on situational and task demands. Future research with a focus on inter-individual differences is necessary in order to determine how stable stability-flexibility biases are over time within the same task, how well such individual biases correlate between tasks within the same domain, and finally how well such individual biases generalize from one task domain to another.

Limitations

Though our findings were robust across experiments and replicated previous findings, three limitations of our study should be mentioned here.

First, we interpret our findings to show that stability in decision making can be manipulated similarly to stability in other domains, namely task-switching. However, in order to support this claim in general, it is necessary to test this effect in a variety of decision making paradigms. Future research is thus needed to investigate the stability effects discussed here in both other value-based decision tasks as well as other domains of decision making, for example risky choice or economic decisions. Only then will we have more certainty about how well the stability-flexibility trade-off generalize across tasks and domains. However, as choice perseveration similar to the stability effects we studied here have been reported in different decision making tasks (e.g., Alós-Ferrer et al., 2016; Bonaiuto et al., 2016; Hämmerer et al., 2016), we expect that replicating our findings across different paradigms should be successful.

A second limitation to our study is the effect of the stimulus onset asynchrony (SOA) manipulation. In a first experiment (Experiment 1b), a SOA of 100 ms was too strong so that this within-trial manipulation overpowered any across-trial effects of choice stability or perseveration. This effect was not predicted by our model or by the literature on SOA in task

switching. However, we argued that this could be due to experimental parameters. Thus, based on the attractor dynamics discussed in the introduction, we chose to try smaller SOAs and buffer trials in between sequences, as that should lead to clear modulation effects. Indeed, these theoretically-informed changes to the task design produced the expected modulation effects in a follow-up experiment. Thus, we conclude that the SOA modulation does work as expected, but only in limited range of under 100 ms. Further, it is not clear how well this manipulation (including this specific time range) can be directly transferred to other tasks.

Third, our task design does not follow the same structure as a classic task switching experiment. In our task, we had sequences of choices where we gradually changed the two choice options over the sequence. The sequences all had the same length. In contrast, there is much more trial-to-trial variability in task switching experiments and often participants cannot fully predict if the next trial will be a switch or a repetition trial. Thus, testing stability effects in a decision making task that is designed to more closely match the setup of a task switching paradigm would allow for a more direct comparison. However, we see strong similarities between our choice sequence approach and voluntary task switching, where participants are presented with a sequence of trials and can decide freely when to switch from one task to another one (e.g., Mittelstädt, Dignath, Schmidt-Ott, & Kiesel, 2018; Mittelstädt, Miller, et al., 2018). In particular, these studies also report a stability-flexibility trade-off that can be modulated by SOA. We take this to suggest that our results are indeed comparable to at least this type of study reported in the field of task switching, and that replicating these findings with a task design that matches the design of a classic task switching paradigm more closely will be a promising endeavor.

Conclusions

In conclusion, our study bridges two fields, namely research on the cognitive stability and flexibility in the field cognitive psychology and the mechanisms of decision making in the field of decision science. We show that the stability-flexibility trade-off, typically investigated in task switching and shielding, generalizes to decision making. Our findings demonstrate that stability-flexibility effects impact decision making across trials similarly to stability-flexibility effects in cognitive control tasks. While this could indicate a common mechanism, there is also some evidence that individual configurations of this trade-off do not generalize across tasks. This raises the question if the stability-flexibility trade-off is grounded in a common mechanism across domains, or if there are distinct domain- or task-specific mechanisms that follow the principle of a stability-flexibility trade-off.

References

- Alós-Ferrer, C., Hügelschäfer, S., & Li, J. (2016). Inertia and decision making. *Frontiers in Psychology*, 7(24), 1–9. <https://doi.org/10.3389/fpsyg.2016.00169>
- Altman, E. M. (2005). Repetition priming in task switching: Do the benefits dissipate? *Psychonomic Bulletin & Review*, 12(3), 535–540.
- Armbruster-Genç, D. J. N., Ueltzhöffer, K., & Fiebach, C. J. (2016). Brain signal variability differentially affects cognitive flexibility and cognitive stability. *Journal of Neuroscience*, 36(14), 3978–3987. <https://doi.org/10.1523/JNEUROSCI.2517-14.2016>
- Armbruster, D. J. N., Ueltzhöffer, K., Basten, U., & Fiebach, C. J. (2012). Prefrontal cortical mechanisms underlying individual differences in cognitive flexibility and stability. *Journal of Cognitive Neuroscience*, 24(12), 2385–2399. https://doi.org/10.1162/jocn_a_00286
- Arrington, C. M. (2008). The effect of stimulus availability on task choice in voluntary task switching. *Memory and Cognition*, 36(5), 991–997. <https://doi.org/10.3758/MC.36.5.991>
- Balaguer-Ballester, E., Lapish, C. C., Seamans, J. K., & Durstewitz, D. (2011). Attracting dynamics of frontal cortex ensembles during memory-guided decision-making. *PLoS Computational Biology*, 7(5). <https://doi.org/10.1371/journal.pcbi.1002057>
- Bonaiuto, J. J., de Berker, A. O., & Bestmann, S. (2016). Response repetition biases in human perceptual decisions are explained by activity decay in competitive attractor models. *eLife*, 5(e20047), 1–28. <https://doi.org/10.7554/eLife.20047>
- Braem, S., Abrahamse, E. L., Duthoo, W., & Notebaert, W. (2014). What determines the specificity of conflict adaptation? A review, critical analysis, and proposed synthesis. *Frontiers in Psychology*, 5, 1–13. <https://doi.org/10.3389/fpsyg.2014.01134>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.

<https://doi.org/10.1163/156856897X00357>

Brisson, B., & Jolicoeur, P. (2007). A psychological refractory period in access to visual short-term memory and the deployment of visual-spatial attention: Multitasking processing deficits revealed by event-related potentials. *Psychophysiology*, 44(2), 323–333.

<https://doi.org/10.1111/j.1469-8986.2007.00503.x>

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

<https://doi.org/10.1016/j.cogpsych.2007.12.002>

Cho, R. Y., Nystrom, L. E., Brown, E. T., Jones, A. D., Braver, T. S., Holmes, P. J., & Cohen, J. D. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive, Affective, & Behavioral Neuroscience*, 2(4), 283–299.

Deco, G., & Rolls, E. T. (2003). Attention and working memory: A dynamical model of neuronal activity in the prefrontal cortex. *European Journal of Neuroscience*, 18(8), 2374–2390.

<https://doi.org/10.1046/j.1460-9568.2003.02956.x>

Dreisbach, G., & Goschke, T. (2004). How positive affect modulates cognitive control: reduced perseveration at the cost of increased distractibility. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 343–353. <https://doi.org/10.1037/0278-7393.30.2.343>

<https://doi.org/10.1037/0278-7393.30.2.343>

Dreisbach, G., & Haider, H. (2008). That's what task sets are for: Shielding against irrelevant information. *Psychological Research*, 72(4), 355–361. <https://doi.org/10.1007/s00426-007-0131-5>

Durstewitz, D., & Seamans, J. K. (2008). The dual-state theory of prefrontal cortex dopamine

- function with relevance to Catechol-O-Methyltransferase genotypes and Schizophrenia. *Biological Psychiatry*, 64(9), 739–749. <https://doi.org/10.1016/j.biopsych.2008.05.015>
- Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience*, 3, 1184–1191. <https://doi.org/10.1016/B978-0-12-236530-0.50021-2>
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, 109(3), 545–572. <https://doi.org/10.1037/0033-295X.109.3.545>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–91. <https://doi.org/10.3758/BF03193146>
- Fischer, R., & Hommel, B. (2012). Deep thinking increases task-set shielding and reduces shifting flexibility in dual-task performance. *Cognition*, 123(2), 303–307. <https://doi.org/10.1016/j.cognition.2011.11.015>
- Fründ, I., Wichmann, F. A., & Macke, J. H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision*, 14(7), 1–16. <https://doi.org/10.1167/14.7.9>
- Gao, J., Wong-Lin, K., Holmes, P., Simen, P., & Cohen, J. D. (2009). Sequential effects in two-choice reaction time tasks: Decomposition and synthesis of mechanisms. *Neural Computation*, 21(9), 2407–2436. <https://doi.org/10.1162/neco.2009.09-08-866>
- Gilbert, S. J., & Shallice, T. (2002). Task switching: A PDP model. *Cognitive Psychology*, 44, 297–337. <https://doi.org/10.1006/cogp.2001.0770>
- Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task set switching. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance*

- XVIII* (pp. 331–355). Cambridge, Massachusetts: MIT Press Ltd.
- Goschke, T. (2003). Voluntary action and cognitive control from a cognitive neuroscience perspective. In S. Maasen, W. Prinz, & G. Roth (Eds.), *Voluntary action: Brains, minds, and sociality* (pp. 49–85). New York, NY: Oxford University Press.
- Goschke, T. (2013). Volition in action: intentions, control dilemmas and the dynamic regulation of intentional control. *Action Science: Foundations of an Emerging Discipline*. MIT Press, Cambridge, MA, 409–434.
- Goschke, T., & Bolte, A. (2014). Emotional modulation of control dilemmas: The role of positive affect, reward, and dopamine in cognitive stability and flexibility. *Neuropsychologia*, 62, 403–423. <https://doi.org/10.1016/j.neuropsychologia.2014.07.015>
- Grange, J. A., & Houghton, G. (2009). Temporal cue – target overlap is not essential for backward inhibition in task switching. *The Quarterly Journal of Experimental Psychology*, 62(10), 2068–2079. <https://doi.org/10.1080/17470210802696096>
- Gratton, G., Coles, M. G., & Donchin, E. (1992). Optimizing the use of information: strategic control of activation of responses. *Journal of Experimental Psychology: General*, 121(4), 480–506.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114–125. <https://doi.org/10.1007/s40881-015-0004-4>
- Hämmerer, D., Bonaiuto, J., Klein-Flügge, M., Bikson, M., & Bestmann, S. (2016). Selective alteration of human value decisions with medial frontal tDCS is predicted by changes in attractor dynamics. *Scientific Reports*, 6(25160), 1–13. <https://doi.org/10.1038/srep25160>
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2008). Search in external and internal spaces.

- Psychological Science*, 19(8), 802–808. <https://doi.org/10.1111/j.1467-9280.2008.02160.x>
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2010). The central executive as a search process: Priming exploration and exploitation across domains. *Journal of Experimental Psychology: General*, 139(4), 590–609. <https://doi.org/10.1037/a0020666>.The
- Hock, H. S., Schöner, G., & Giese, M. (2003). The dynamical foundations of motion pattern formation: Stability, selective adaptation, and perceptual continuity. *Perception & Psychophysics*, 65(3), 429–457. <https://doi.org/10.3758/BF03194574>
- Hommel, B., & Colzato, L. S. (2017). The social transmission of metacontrol policies: Mechanisms underlying the interpersonal transfer of persistence and flexibility. *Neuroscience and Biobehavioral Reviews*, 81, 43–58. <https://doi.org/10.1016/j.neubiorev.2017.01.009>
- JASP Team. (2018). JASP (Version 0.9.2)[Computer software]. Retrieved from <https://jasp-stats.org/>
- Kelso, J. A. S., & Schöner, G. (1988). Self-organization of coordinative movement patterns. *Human Movement Science*, 7(1), 27–46. [https://doi.org/10.1016/0167-9457\(88\)90003-6](https://doi.org/10.1016/0167-9457(88)90003-6)
- Koch, I., & Allport, A. (2006). Cue-based preparation and stimulus-based priming of tasks in task switching. *Memory and Cognition*, 34(2), 433–444. <https://doi.org/10.3758/BF03193420>
- Koch, I., Gade, M., Schuch, S., & Philipp, A. M. (2010). The role of inhibition in task switching: A review. *Psychonomic Bulletin & Review*, 17(1), 1–14. <https://doi.org/10.3758/PBR.17.1.1>
- Logan, G. D. (2003). Executive control of thought and action: In search of the wild homunculus. *Current Directions in Psychological Science*, 12(2), 45–48.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward

inhibition. *Journal of Experimental Psychology: General*, 129(1), 4–26.

<https://doi.org/10.1037/0096-3445.129.1.4>

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning Memory and Cognition*, 22(6), 1423–1442.

<https://doi.org/10.1037/0278-7393.22.6.1423>

Meiran, N. (2010). Task switching and cognitive self control. In R. Hassin, K. N. Ochsner, & Y. Trope (Eds.), *Self control in society, mind and brain* (pp. 202–220). NY: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195391381.001.0001>

Meiran, N., Chorev, Z., & Sapir, A. (2000). Component processes in task switching. *Cognitive Psychology*, 41(3), 211–253. <https://doi.org/10.1006/cogp.2000.0736>

Mekern, V. N., Sjoerds, Z., & Hommel, B. (2019). How metacontrol biases and adaptivity impact performance in cognitive search tasks. *Cognition*, 182(May 2018), 251–259.

<https://doi.org/10.1016/j.cognition.2018.10.001>

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167–202.

Mittelstädt, V., Dignath, D., Schmidt-Ott, M., & Kiesel, A. (2018). Exploring the repetition bias in voluntary task switching. *Psychological Research*, 82(1), 78–91.

<https://doi.org/10.1007/s00426-017-0911-5>

Mittelstädt, V., Miller, J., & Kiesel, A. (2018). Trading off switch costs and stimulus availability benefits : An investigation of voluntary task-switching behavior in a predictable dynamic multitasking environment. *Memory and Cognition*, 46, 699–715.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7(3), 134–140.

[https://doi.org/10.1016/S1364-6613\(03\)00028-7](https://doi.org/10.1016/S1364-6613(03)00028-7)

- Noest, A. J., van Ee, R., Nijs, M. M., & van Wezel, R. J. A. (2007). Percept-choice sequences driven by interrupted ambiguous stimuli: A low-level neural model. *Journal of Vision*, 7(8), 10. <https://doi.org/10.1167/7.8.10>
- Padoa-Schioppa, C. (2013). Neuronal origins of choice variability in economic decisions. *Neuron*, 80(5), 1322–1336. <https://doi.org/10.1016/j.neuron.2013.09.013>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442. <https://doi.org/10.1163/156856897X00366>
- Plessow, F., Fischer, R., Kirschbaum, C., & Goschke, T. (2011). Inflexibly focused under stress: Acute psychosocial stress increases shielding of action goals at the expense of reduced cognitive flexibility with increasing time lag to the stressor. *Journal of Cognitive Neuroscience*, 23(11), 3218–3227.
- Rączaszek, J., Tuller, B., Shapiro, L. P., Case, P., & Kelso, S. (1999). Categorization of ambiguous sentences as a function of a changing prosodic parameter: A dynamical approach. *Journal of Psycholinguistic Research*, 28(4), 367–393. <https://doi.org/10.1023/A:1023289031747>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. <https://doi.org/10.1037/0033-295X.85.2.59>
- Rolls, E. T. (2010). Attractor networks. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 119–134. <https://doi.org/10.1002/wcs.1>
- Scherbaum, S., Dshemuchadse, M., Leiberg, S., & Goschke, T. (2013). Harder than expected: Increased conflict in clearly disadvantageous delayed choices in a computer game. *PLoS ONE*, 8(11), e79310. <https://doi.org/10.1371/journal.pone.0079310>
- Scherbaum, S., Frisch, S., Leiberg, S., Lade, S. J., Goschke, T., & Dshemuchadse, M. (2016).

Process dynamics in delay discounting decisions : An attractor dynamics approach.

Judgement and Decision Making, 11(5), 472–495.

Scherbaum, S., Haber, P., Morley, K., Underhill, D., & Moustafa, A. A. (2018). Biased and less sensitive: A gamified approach to delay discounting in heroin addiction. *Journal of Clinical and Experimental Neuropsychology*, 40(2), 139–150.

<https://doi.org/10.1080/13803395.2017.1324022>

Schmidt, J. R. (2013). Questioning conflict adaptation: Proportion congruent and Gratton effects reconsidered. *Psychonomic Bulletin and Review*, 20(4), 615–630.

<https://doi.org/10.3758/s13423-012-0373-0>

Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). From the cover: Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, 102(29), 10393–10398. <https://doi.org/10.1073/pnas.0503903102>

Tuller, B., Case, P., Ding, M., & Kelso, J. A. S. (1994). The nonlinear dynamics of speech categorization. *Journal of Experimental Psychology. Human Perception and Performance*, 20(1), 3–16. <https://doi.org/10.1037/0096-1523.20.1.3>

Ueltzhöffer, K., Armbruster-Genç, D. J. N., & Fiebach, C. J. (2015). Stochastic dynamics underlying cognitive stability and flexibility. *PLoS Computational Biology*, 11(6), 1–46.

<https://doi.org/10.1371/journal.pcbi.1004331>

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592.

<https://doi.org/10.1037/0033-295X.108.3.550>

Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron*, 60(2), 215–234.

<https://doi.org/10.1016/j.neuron.2008.09.034>

Zwosta, K., Hommel, B., Goschke, T., & Fischer, R. (2013). Mood states determine the degree of task shielding in dual-task performance. *Cognition and Emotion*, 27(6), 1142–1152.

<https://doi.org/10.1080/02699931.2013.772047>

Appendix

Simulations of Attractor Dynamics and Choice Perseveration

We used our previously published attractor model (Scherbaum et al., 2016) to simulate choice perseveration in a binary value-based decision task. We will describe the model architecture and simulation procedure briefly; for an in-depth description, please see Scherbaum et al. (2016).

Model Architecture

The model consists of two neural units with sigmoidal activation functions. Each unit represents one of the decision options, small/near (SN) or large/far (LF), and receives input from that decision option. Both units are potentially stable states, i.e. attractors. When the units receive input, they are activated and form attractors. Through the mechanisms of self-excitation and mutual inhibition, over time one attractor becomes more stable than the other attractor and the system settles into this attractor, at which point the decision is made (Rolls, 2010).

Mathematical Description

Mathematically, the model is based on two coupled differential equations (one for each unit) that represent non-linear neural activation dynamics (Hock, Schöner, & Giese, 2003; Noest, van Ee, Nijs, & van Wezel, 2007).

$$\tau \dot{u}_{SN} = -u_{SN} + h + w_r \cdot \sigma(u_{SN}) + w_i \cdot \sigma(u_{LF}) + I_{SN}$$

$$\tau \dot{u}_{LF} = -u_{LF} + h + w_r \cdot \sigma(u_{LF}) + w_i \cdot \sigma(u_{SN}) + I_{LF}$$

Here, τ denotes the timescale (defining the step size of the Euler solution), h denotes the resting level, w_i the (inhibitory) coupling strength of the two equations, and w_r the recurrent feedback, I_{SN} and I_{LF} the input representing the attractiveness of the two options, and σ a sigmoid non-linearity, mirroring non-linear neural population dynamics: $\sigma(u) = \frac{1}{e^{-\beta(u-a)}}.$

The non-linearity of σ limits interactions between the two units only to the extent that the activation u exceeds a soft threshold (Erlhagen & Schöner, 2002). Note that the β -parameter is also called the gain-parameter, modulating the discreteness of neural activation states. We kept parameter values identical to the original simulation (Scherbaum et al., 2016), see Table A 1.

Table A 1

Basic and varied parameters of the model used for simulation

Parameter	Values
τ	10
h	-5
w_i	-7
w_r	6
a	0
β	1
t_{iii}	50 vs. 57 vs. 64
s_{soa}	.2
t_{soa}	1 vs. 3 vs. 5

The two coupled differential equations constitute a neural system with two units inhibiting each other so that only one unit can win the competition and determine the final choice. This competition unfolds over time. The dynamics of the system are modulated by a control-parameter, c , representing the relative attractiveness of the two options via the strength of the two Inputs I_{SN} and I_{LF} relative to a general input strength $I = 6$ ($c < 0$ favoring the LF option and $c > 0$ favoring the SN option): $I_{SN} = I + c$, and $I_{LF} = I - c$. Hence, the lower c , the lower

the input to the SN option and the higher the input to the LF option, leading to u_{LF} winning the competition over u_{SN} . Hence, c determines the final choice when both units are in the same starting state, i.e., an activation $u_{SN} = u_{LF} = 0$. In contrast, if the starting state differs, i.e., $u_{SN} \neq 0$ and $u_{LF} = 0$, u_{SN} might win the competition due to its initial advantage. The equation of motion that defines the dynamics of this system used for simulation is derived by differentiation. We simulated the behavior of the derived dynamical system by numerical integration with each trial having a maximum length of 200 time steps. In our simulation, the inputs for the choice options were switched on after t_{iti} time steps constituting the system's relaxation time. Results were obtained using Matlab 2015a running under Windows 7.

Simulation 1: Inter-trial dynamics of decision making modulated by relaxation time

In Simulation 1, we wanted to test our prediction that choice perseveration decreases with increasing relaxation time.

Methods. In Simulation 1, one unit represented the SN and the other unit the LF option. We defined the control parameter c such that $c = 0$ represented options of equal attractiveness for a participant. Consequently, $c < 0$ represented a more attractive LF option, and $c > 0$ represented a more attractive SN option. Analogously to the original simulation (Scherbaum et al., 2016), the parameter c was varied, first across different intervals between the small and the large option, and second, across different value differences between the near and the far option. We assumed multi-stability of the system for intermediate intervals between options and mono-stability for extreme intervals between options. Hence, we overall varied c within a parameter window of $[-.12, .12]$ for each option. We built ascending and descending sequences of twelve intervals between the two options: In ascending sequences, c favored the LF option in the beginning ($I_{SN} = I + c$, and $I_{LF} = I - c$, with $c = .12$) and then linearly decreased for the LF option

while linearly increasing for the SN option, thus favoring the SN option at the end of the sequence ($I_{LF} = I + c$, with $c = -.12$; $I_{SN} = I + c$, with $c = .12$); in descending sequences, c was switched around so that it favored the SN option in the beginning and the LF option at the end of the sequence. Within one sequence, we simulated trials continuously, so that activation from the previous trial could carry over to the next trial. As explained in the introduction (see main text), this led to incomplete relaxation to the starting state of the system and hence, a bias in the next trial resulting in different switch points between SN and LF choices depending on previous choice history, i.e., choice perseveration or hysteresis (Scherbaum et al., 2016). To completely reset the system state between sequences, we inserted an empty trial (I_{LF} and I_{SN} set to 0) allowing the model to completely relax back to the neutral starting state.

In order to realize the conceptually derived modulation of perseveration by the relaxation time, we varied the time at which inputs were switched on within trials. In the original simulation, the input for the choice options was switched on after 50 time steps. As an example of the modulation of perseveration by relaxation time, we ran the current simulation three times, with the choice options being switched on after 50, 57, or 64 time steps (t_{iti} ; see Table A 1), respectively, to simulate small, medium, and large time intervals between trials. We expected choice perseveration to decrease with increasing time steps, due to longer relaxation time.

Results. Simulation 1 revealed choice perseveration effects for all variations of relaxation time. Furthermore, simulation 1 showed that, just as predicted, perseveration effects decrease with increasing relaxation time (see Figure A 1).

Figure A 1: Simulated choice percentages of the large/far (LF) option for different inter-trial intervals (ITI). The c parameter comprises difference between the c value for the small/near (SN) option and the c value for the LF option, with $c < 0$ favoring selection of the LF option and $c > 0$ favoring selection of the SN option. The perseveration effect (area between LF choices in ascending (orange) and descending (blue) sequences) decreases with increasing ITI.

Simulation 2: Inter-trial dynamics of decision making modulated by stimulus onset asynchrony

In Simulation 2, we wanted to test our prediction that choice perseveration decreases with increasing stimulus onset asynchrony (SOA).

Methods. Methods for simulation 2 were the same as in simulation 1, except for the following changes. For simulation 2, we kept the ITI constant at 50 time steps. We implemented SOA so that one option received a boost (s_{soa} in Table A 1) one, three, or five time steps before the inputs for both options were switched on (t_{soa} in Table A 1). This boost resembles an attentional bias favoring the option that appears first. In ascending sequences (where the LF option is favored in the beginning of the sequence), the SOA was implemented so that the SN option received the boost for t_{soa} time steps, thus biasing the choice process against the more attractive LF option. In descending sequences (where the SN option is favored in the beginning of the sequence), the SOA was implemented so that the LF option received the boost for t_{soa} time steps, thus biasing the choice process against the more attractive SN option. We varied t_{soa} to model a small SOA ($t_{soa} = 1$), medium SOA ($t_{soa} = 3$) and large SOA ($t_{soa} = 5$). We expected choice perseveration to decrease with increasing SOA because the SOA counteracts the bias towards the initially favored, more likely choice option.

Results. Simulation 2 revealed choice perseveration effects for all variations of SOA. Additionally, perseveration decreased with increasing SOA, as expected (see Figure A 2).

Figure A 2: Simulated choice percentages of the large/far (LF) option for different stimulus onset asynchronies (SOA). The c parameter comprises difference between the c value for the small/near (SN) option and the c value for the LF option, with $c < 0$ favoring selection of the LF option and $c > 0$ favoring selection of the SN option. The perseveration effect (area between LF choices in ascending (orange) and descending (blue) sequences) decreases with increasing SOA.

Conclusion

In summary, we used our previously published attractor model (Scherbaum et al., 2016) to simulate how ITI and SOA manipulations affect the stability or perseveration of decision making. The model simulation demonstrates that choice perseveration decreases with increasing ITI and with increasing SOA. Thus, the simulations show that the hypotheses we put forward in this paper can be mathematically derived from our attractor model.