Judgment and Decision Making, Vol. 15, No. 1, January 2020, pp. 74-81

Robust consistency of choice switching in decisions from experience

Eldad Yechiam*

Abstract

Decision making is a multifaceted process but studies of individual differences in decision behavior typically use only the proportions of choices from different options as behavioral indices. I examine whether the probability of choice switching in decisions from experience, reflecting one's exploration strategy, is consistent across sessions and tasks. In Study 1, I re-analyzed an experiment in which participants performed six decision tasks in two sessions that were 45 days apart. Choice switching rates were highly consistent across sessions and tasks, and their consistency exceeded that of rates of risky choices. In Study 2 I conducted a similar analysis for the Technion Prediction Tournament, and also found higher consistency across tasks in switching rates than in choice rates. Additionally, in both studies, there were moderate to high correlations between switching rates at the beginning and towards the end of the task. The results thus highlight an often overlooked but highly consistent and independent aspect of human behavior.

Keywords: decision making, experience, exploration, individual differences

1 Introduction

In decisions from experience individuals have no explicit information about the incentive structure, and past decision outcomes are used to guide behavior (Erev & Haruvy, 2015). Studies of individual differences in decisions from experience typically examine the rate of choices from different alternatives across multiple trials (e.g., Bechara, Damasio, Damasio & Anderson, 1994; Lejuez et al., 2002; Figner, Mackinlay, Wilkening & Weber, 2009; Weller, Levin, Shiv & Bechara, 2007; Yechiam & Ert, 2011; Frey, Pedroni, Mata, Rieskamp & Hertwig, 2017). This index is then used to infer psychological constructs such as the advantageousness of the choice (i.e., expected value maximization), subjective weighting of positive and negative events, or the degree of relevant biases such as the certainty and reflection effects (Weller et al., 2007). For example, in the Iowa Gambling task (Bechara et al., 1994), which involves repeated selection between four decks of cards, the rate of choices from the two alternatives with the higher (or lower) expected value is typically used as the performance index. The rate of switching between different alternatives, which delineates one's exploration strategy, has rarely been examined in this type of decisions (but see exceptions in Hills & Hertwig, 2017; Yechiam, Arshavski, Shamay- Tsoory, Yaniv & Aharon, 2010; Erev & Haruvy, 2015). In the present study I examine the consistency of the switching rate across tasks and sessions, in comparison to that of the often investigated choice rate, which is used here as a benchmark.

In a stable environment, only a few switches between choices (i.e., equal to the number of alternatives minus one) are required to learn about the outcome distribution of each option, as decision makers can thoroughly examine one alternative after the other. Therefore, above a minimal level, the extent of switching seems to reflect a certain style of exploration rather than the level of exploration per se. In the stable phase after learning has occurred, choice switching is no longer necessary for exploration (unless one assumes the environment may change), and likely reflects biases such as a win-stay lose-shift strategy (Worthy, Hawthorne & Otto, 2013) and probability matching (Edwards, 1961; Vulkan, 2000), or erratic choice behavior (Busemeyer & Stout, 2002). These considerations suggest that individual differences in choice switching may depend on the task phase.

Nevertheless, the rate of choice switching across the entire task might also be a meaningful individual differences parameter. For example, studies of autism have found that high functioning autistic individuals have an extremely high rate of switching between choice options in the Iowa Gambling task (Johnson, Yechiam, Murphy, Queller & Stout, 2006; Yechiam et al., 2010; Mussey, Travers, Grofer, Klinger & Klinger, 2014), and this tendency is exhibited throughout the task (Mussey et al., 2014). A similar finding was observed in individuals with complete or partial absence of the corpus callosum (Brown, Anderson, Symington & Paul, 2012). However, no previous study systematically examined whether healthy adults are individually consistent in choice switching across different sessions of the same task, across different trial blocks, and across different tasks.

The main goal of the present study was to evaluate the in-

This work was supported by the I-CORE program of the Planning and Budgeting Committee and the Israel Science Foundation (Grants no. 1821/12 and 207/17).

*Max Wertheimer Minerva Center for Cognitive Studies, Technion – Israel Institute of Technology, Haifa, Israel. Email: yeldad@tx.technion.ac.il.

TABLE 1: Payoff structure for Study 1 tasks and mean choice and switch rates across sessions. Top panel: The five two-option tasks. Bottom panel: The Iowa Gambling task.

Two-option tasks:					
	Option S (safe)	Option R (risky)		P(R)	Switch
Mixed	0 for sure	−200 or 200 with equal probability (0.5)		0.41	0.23
Gain	200 for sure	0 or 400 with equal probability (0.5)		0.35	0.19
Loss	−200 for sure	−400 or 0 with equal probability (0.5)		0.39	0.23
Assym. Gain	320 for sure	400 with 0.8 probability, otherwise 0		0.46	0.19
Assym. Loss	−80 for sure	0 with 0.8 probability, otherwise −400		0.52	0.22
Iowa Gambling task:					
Deck	Gain	Loss	EV	P(Deck)	Switch
A	100 for sure	250 with 0.5 probability	−25	0.12	0.28
B	100 for sure	1,250 with 0.1 probability	−25	0.23	
C	50 for sure	50 with 0.5 probability	25	0.30	
D	50 for sure	250 with 0.1 probability	25	0.34	

Notes: EV = Expected value; P(R), P(deck) = Proportion of option R or deck selections across trials; Switch = Rate of switching choices across trials.

ternal and external consistency of choice switching in comparison to that of the often used choice rates. In addition, I examined the relationship between choice switching and choice rates in different stages of the task in order to gain insight into the implications of high choice switching. Study 1 used data from Yechiam and Telpaz (2011; see also Yechiam, Telpaz, Krupenia & Rafaeli, 2016) to investigate the consistency of these indices across diverse tasks performed in two sessions. Study 2 used data from the Technion Prediction Competition (Erev et al., 2010) to further study the consistency across similar and diverse tasks.

2 Study 1

Yechiam and Telpaz (2013; and see also Yechiam et al., 2016) evaluated the consistency of choice rates in decisions from experience across two sessions that were 45 days apart. Participants performed five two-option decision tasks in each session, in random order, and then also performed the Iowa Gambling task (IGT; Bechara et al., 1994). The payoff structure for these tasks was highly diverse, as shown in Table 1. The two-option tasks encompassed different domains (gain, loss, and mixed gain-loss) as well as different probabilities for relatively positive/negative outcomes (equiprobable positive and negative outcomes vs. rare negative outcomes). The current study focuses on the average switching rate, namely the probability of switching in a given trial. I compare the consistency of switching rates across sessions and tasks to that of choice rates, and also examine the effect of task phase on individual differences in choice switching.

Participants: One hundred and thirty undergraduate students (65 men and 65 women) participated in the study. Their average age was 23.5 (ranging between 18 and 28). They received a basic fee of NIS 50 per session for performing the two-option task, 20 NIS for performing the IGT (as described in Yechiam et al., 2015), and additional payoffs according to the total score in one randomly selected two-option task as well as in the IGT at a rate of NIS 1 per 1000 points.

Decision tasks: The five two-option tasks involved 60 repeated trials each. Upon clicking a button, the participant saw the payoff for the choice, and an accumulating payoff counter was updated. Payoffs were randomly drawn from each task's payoff distribution (Table 1). In order to decrease the transparency of the payoff structure, a noise factor randomly sampled in each trial from the set [−1,0,1] was added to each outcome in each task. As noted above, final fees were based in part on the accumulated payoffs. The instructions are in the Appendix. Importantly, in these decisions-from-experience tasks, participants do not know the incentive structure in advance but instead learn it by making decisions and receiving payoff feedback.

The IGT was administered by computer with 100 trials. In this task participants repeatedly choose among four decks of cards, labeled A, B, C, and D. Payoff feedback is displayed following each choice. Payoffs were predetermined based on the exact outcomes in Bechara et al. (1994). Their general

https://doi.org/10.1017/S1930297500006914 Published online by Cambridge University Press

distribution is shown in Table 1. As in Bechara et al. (1994), a noise factor was added to some of the outcomes (losses in A were compounded by [−100,−50, 0, 50, or 100]; and losses in C by [−25, 0, 25]). The original task instructions were used (see Appendix). All tasks, including the IGT, were performed in two sessions. The average time interval between sessions was 46 days (with a standard error of 0.72 days).[1]

Analysis: I compared the consistency of the mean probability of switching to that of the mean rate of selections from the riskier alternative in each task. To calculate the mean rate of selections from the riskier alternatives in the IGT, I pooled across the two riskier (and disadvantageous) decks as commonly done in the literature for simplicity (the results for each of the four decks appear in the Supplement

For the analysis of consistency across sessions, I calculated the correlation between session 1 and session 2 indices (i.e., switching and choice rates) in each task, and then compared the average correlation across tasks in session 1 with the average in session 2. I also examined the correlation of each task in one session with the same task in the other session.

For the analysis of consistency across tasks, I similarly averaged the correlations between pairs of tasks performed in the same session and also calculated Cronbach's α. For the analysis of consistency within tasks, I used the mean correlation (across tasks) between the first and last 20-trial blocks. Finally, I also examined the average correlations between switching and choice rates, for the entire task and for the first and last blocks.

Differences between correlations were tested by Zou's test, (Zou, 2007) which corrects for the correlation between pairs of variables. Zou's test produces a confidence interval for the correlation difference, which was converted into an exact p-value using Lin's approximation (Lin, 1989)

2.1 Results

The mean switching rates at the first and last blocks of each session are presented in Figure 1. As can be seen, across tasks there was a drop in switching rates between the first and second session ($F(1, 129) = 34.83$, $p < .001$) and from the first to the last block, $F(1, 129) = 268.18$, $p < .001$). There were also large differences in switching rates among the six tasks ($F(5, 129) = 19.46$, $p < .001$). I tested whether, despite these differences, individual differences in choice switching are consistent across tasks and sessions.

[1]After completing the decision tasks, participants also performed several other tests, including the Reading the Mind in the Eyes Test (RMET: Baron-Cohen et al., 2001), a facial affect-recognition test which was developed to study autism. I checked whether similar to the findings in autism (e.g., Johnson et al., 2006), those with poor RMET performance would exhibit more choice switching.

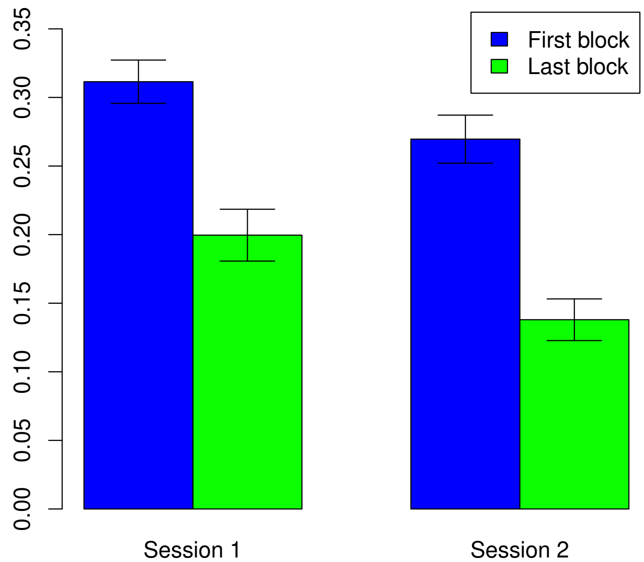

Figure 1: Choice switching rates in Study 1. Mean switching rates across tasks in the first and last blocks of 20 trials, in each session. Error bars denote standard errors.

Consistency across sessions: The correlations between session 1 and session 2 switching rates and choices rates are presented in Figure 2. I first examined the consistency for the same tasks performed in different sessions (the diagonal cells in Figure 2). The mean correlation for switching rates was 0.49 ($p < .001$), compared to 0.28 ($p = .001$) for choice rates. This difference was significant (Zou's test, $p = .04$). The correlation was higher for switching rates than for choice rates in all six tasks. The highest correlation across sessions in both indices was in the IGT. In this task, the correlation for choice switching reached 0.74 compared to 0.43 for choice rates.

I also examined the consistency between tasks performed in session 1 and different tasks performed in session 2. Theses findings are also shown in Figure 2 (the non-diagonal cells). As can be seen, overall, correlations tended to be higher for switching rates. Across tasks, the mean correlation between different tasks performed in session 1 and 2 was 0.38 for switching rates ($p < .001$) and only 0.12 for choice rates ($p = .17$).

Consistency across tasks within a session: The average Cronbach's α across sessions for switching rates was 0.85 (adequate consistency), whereas for choice rates it was 0.45 (very low consistency). The correlations between different tasks in each session are shown in Figure 3. The mean correlation for switching rates in session 1 was 0.43 ($p < .001$) whereas for choice rates it was 0.09 ($p = .91$). The mean correlation for switching rates in session 2 was 0.53 ($p < .001$) compared to 0.16 ($p = .07$) for choice rates. Furthermore, as indicated Figure 3, in both sessions the correlations were

Figure 2: Correlations across the two sessions in Study 1. The top and bottom panels denote the correlations for switching rates and choice rates, respectively.

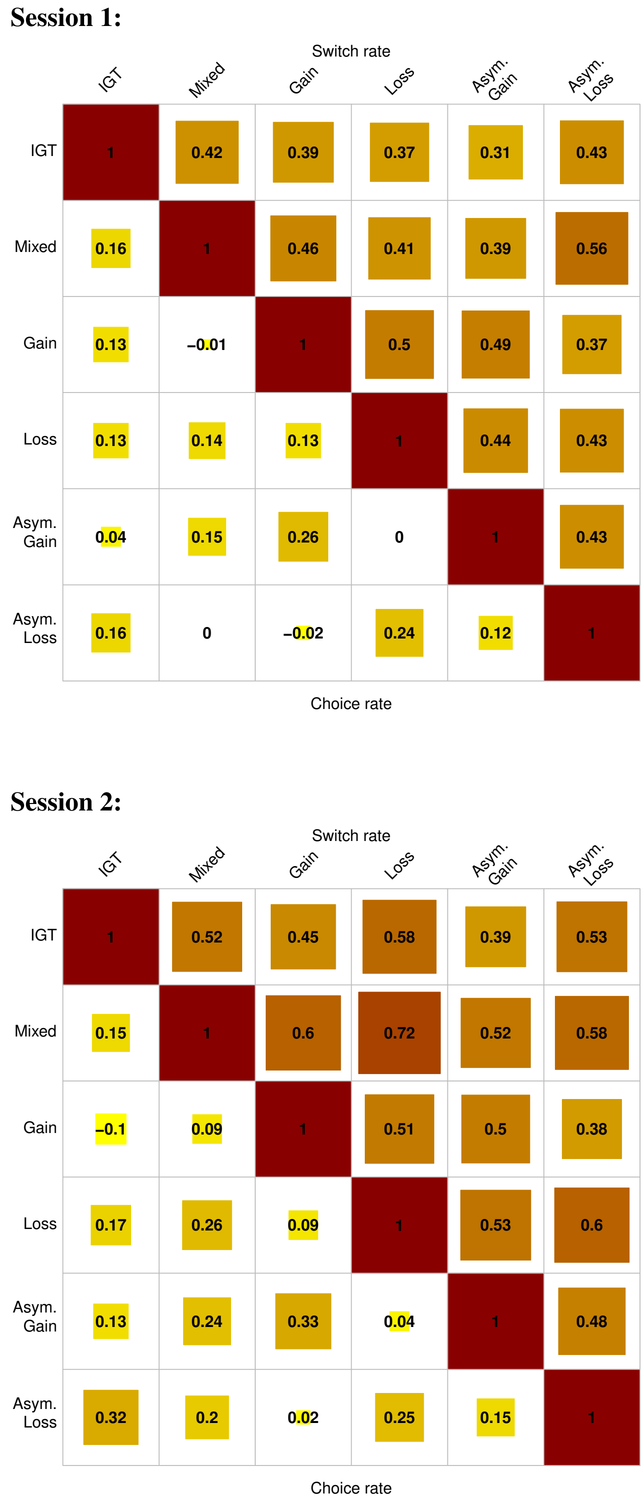

Figure 3: Correlations across tasks within each session in Study 1. Correlations in switching rates appear in the top right triangles, and correlations in choice rates appear in the bottom left triangles. The top and bottom panels denote the correlations for session 1 and 2, respectively.

https://doi.org/10.1017/S1930297500006914 Published online by Cambridge University Press

higher for switching rates than for choice rates in *all* tasks. Thus, whereas individual differences in choice switching were highly consistent in different tasks, individual differences in choice rates were not.

Consistency within the task: An examination of the average correlation between the first and last block switching rates reveals that the two indices were highly correlated, with a mean correlation of 0.48 ($p < .001$) across tasks. This suggests that a common factor affects choice switching in the first and last block, though there is also unshared variance between blocks.[2]

By comparison, the mean correlation for choice rates was 0.33 ($p < .001$; Zou's test, $p = 0.16$, for the difference between switching rates and shoice rates.).

The mean correlation between switching and choice rates in each task was 0.30 ($p < .001$). The correlations were somewhat higher for the IGT ($r = 0.54$, $p < .001$) and somewhat lower for the two-option tasks ($r = 0.25$, $p = .004$). Thus, although those who switched more tended to also take more risk and to select disadvantageously on the IGT, the two constructs exhibited a reasonable degree of independence.[3]

I then examined these correlations in different blocks. In this respect, it was of interest to study the IGT, the only task where choice rates reflect decision performance (given the large differences in expected value between options). The results showed that in the first block of trials, the correlation between switching rates and disadvantegous selections on the IGT was 0.09 ($p = .31$) while in the last block it increased to 0.46 ($p < .001$).[4] Therefore, the correlation between choice switching and disadvantegous selections in the IGT emerged only in advanced phases of the task. A less distinct trend emerged in the two-option tasks, where the initial correlation between switching rates and risk taking in block 1 was 0.15 ($p = .09$) and in the last block it was 0.20 ($p = .02$).

3 Study 2

In this study I tested the replicability of the findings of Study 1 as well as their boundary conditions. Data from the Technion Prediction Tournament (TPT; Erev et al., 2010) were used to examine consistency across tasks of different domains as well as within domains. The TPT was a set of competitions for the best model in terms of predictive ability. It included three separate competitions for decisions from description, decisions from sampling (where one first samples options and then makes a single choice) and decisions from feedback (where one makes repeated choices and receives payoff feedback following each choice). Fourteen research teams participated in the three competitions. Behavioral decisions of experimental participants were examined, and the data were given to the participating researchers for estimating the models in each competition (estimation set). The researchers' goal was to best predict additional behavioral decisions collected after the prediction models were submitted (prediction set).[5] The current study focused on the decisions from feedback competition, in which choice switching could be assessed as in Study 1.

The TPT dataset for the decisions-from-feedback section included 120 decision tasks in which the payoffs were randomly determined. In 40 of the problems all outcomes were losses (loss domain), in 40 they were all gains (gain domain), and in 40 both (mixed domain). Each participant performed 12 randomly chosen tasks in random order. In addition to studying consistency across all tasks, I evaluated whether given the narrower incentive structure within each domain (loss, gain, and mixed), consistency in choice rates within domains would be closer to that of switching rates.

3.1 Method

Participants: Two-hundred undergraduate students participated in the decisions-from- feedback section of the TPT (the authors did not report gender rates). All participants received a basic fee of NIS 40 and additional payoffs according to the score in one randomly chosen trial.

Decision task: The task layout was as in the two-option tasks of Study 1 with 100 trials in each task. The payoff scheme was as follows: In each task a safe option S produced a fixed payoff M and a riskier option R produced a larger payoff H with probability p and otherwise a lower payoff L (where $L < M < H$). The values of L, M, H, and p were randomly set within certain constraints so that the expected values of the two options would be similar (as detailed in Erev et al., 2010). In 40 of the problems S, M, and H were losses; in another 40 they were gains, and in the last 40 S was a loss and H was a gain. The complete task payoffs appear in the Supplement. Each participant performed 12 randomly chosen tasks from the 120 tasks, and groups of twenty participants performed the same set of tasks in random order.

[2]Consistency across sessions in a given task in choice switching was also similar for the first and last blocks (first block: $r = 0.33$, $p < .001$; last block: $r = 0.36$, $p < .001$).

[3]There was also a weak but significant correlation between switching rates across sessions and the RMET score ($r = -0.18$; $p = .04$); participants who switched more frequently had lower scores on the RMET.

[4]High switching in the first block also predicted significantly more disadvantegous selections in the last block ($r = 0.25$, $p = .004$). Possibly, those who switched more in the beginning continued to do so in later task phases, with subsequent effects on IGT performance.

[5]The prediction set followed the same general constraints as the estimation set with randomized probabilities and outcomes within these constraints (see Method section).

TABLE 2: Study 2 results. Correlations across tasks in switching and choice rates, for each payoff domain. (** = $p < .01$)

Domain	Switch rates r	Choice rates r
Mixed	0.34**	0.16
Gain	0.42**	0.11
Loss	0.47**	0.10

Analysis: I compared the consistency in switching rates and choice rates (from option R) across each pair of tasks as in Study 1. The correlations were calculated for each group of twenty individuals who performed the same set of tasks, and were then averaged across these ten groups. I also examined the consistency between the first and last block of each task as in Study 1.

3.2 Results

Consistency across tasks: Mean choice and switching rates for each of the 120 decision tasks are reported in the Supplement. The mean Cronbach's α score for switching rates across all 10 groups was 0.88, while for choice rates it was 0.54. The mean correlation across task pairs for switching rates was 0.39 (p < .001), compared to 0.10 (p = .16) for choice rates. Higher correlation across task pairs was observed for nine of the 10 groups of participants (performing different tasks). I also compared the consistency in each of the three payoff domains; the results are shown in Table 2. The difference between correlations was robust in all three domains.

Consistency within the task: The within-task correlation between switching rates in the first and last blocks was 0.32 (p < .001), compared to 0.41 (p < .001) for choice rates. In this respect, higher correlation for choice rates was observed in 9 out of 10 groups of participants. The mean correlation between switching and choice rates in a given task was 0.17 (p = .053), confirming their distinctiveness.

4 General Discussion

In Study 1 and 2 I found consistent individual differences in choice switching across different tasks and in Study 1 also across different sessions. I also found that the consistency of choice switching across sessions and tasks was higher than the consistency of choice rates. First, in Study 1 the correlation between individuals' switching rates in different sessions was significantly higher than the correlation for choice rates. Secondly, in both Study 1 and Study 2 correlations between switching rates in different tasks were considerably higher than those for choice rates. For example, in Study 1 the mean correlation in switching rates between tasks performed in the same session was 0.48, compared to 0.13 for choice rates. This was also found in Study 2 both across payoff domains and also within domains. Therefore, choice switching in decisions from experience seems to be a robust index of individual differences.

With respect to consistency across different blocks of trials, I found moderate to high consistency between switching rates in the very first block and in the last block. The consistency across blocks was somewhat higher (and similar to that of choice rates) in Study 1 and somewhat lower in Study 2. This suggests that while task phase affects individual differences in choice switching, there is still a common factor driving extensive choice switching in different stages.

Finally, I observed that though there were positive correlations between switching and choice rates, these correlations were not high. Across tasks, choice rates accounted for 9% of the variance in switching rates in Study 1, and 3% in Study 2, affirming the distinctiveness of the two constructs. Higher associations were recorded for the IGT, where particularly in advanced trials, the switching rate was highly correlated with making disadvantegous selections. This finding sheds light on the effect of task phase on the adaptiveness of high choice switching. Specifically, it seems to reflect potential detrimental upshots of high choice switching in the post-learning phase.

Probably the next challenge in the examination of choice switching is better mapping of this construct to different psychological traits. A tremendous mapping effort along similar lines was carried out for choice rates (see e.g., Levin, Hart, Weller & Harshman, 2007; Dohmen, Falk, Huffman, Sunde, Schupp & Wagner, 2011; Frey et al., 2017) and it would be interesting to test this for a decision characteristic – switching rate – which seems to be more consistent than choice rates in decisions from experience.

An important further avenue of research is to examine choice switching in dynamic tasks. Why do some people consistently switch options while others gain information by selecting one alternative and then the other? Extensive switching has the advantage that it can reveal contingencies associated with the timing of one option's outcomes in relation to the others. For example, it can more easily reveal patterns such as "when one is high the other is low". By contrast, low switching has the advantage that it can more easily reveal dynamic changes within an alternative, such as drops and increases in its value that are contingent on its past trials' outcomes. In the later phases of a task, after learning from experience, habitually high switching rates can impair performance, e.g., when the best strategy at this point is to choose the single best option every time.

Examining choice switching in decision problems with dynamically changing outcomes might yield further insight into these adaptive properties of high and low switching rates.

Appendix: Task Instructions

Initial instructions in Study 1

In this experiment you will perform six decision making tasks. Your basic payoff is NIS 50. This payoff will be updated based on the accumulated score in one randomly chosen task. This will be determined after you perform the six tasks by the throw of a die.

Instructions for the two-option tasks in Study 1

This is the first of six tasks that you will perform. In the form on the computer screen there are two buttons, labeled A and B. Your task is to choose between the two buttons by clicking any of them. You can click on a button several times in a row (as much as you want) or switch between buttons (as much as you like). The payment you receive for your choice will appear on the chosen button, and the accumulating payoff will appear below. You will not know the payment for each choice in advance. Some choices might be followed by gains and others by losses. For the task that would be randomly selected, you will gain or lose NIS 1 for every 1000 game points. You will receive a message telling you when the task is ended and a new task begins.

Instructions for the Iowa Gambling task

In front of you on the screen, there are four decks of cards labeled A, B, C, and D. When we begin the game please select one deck at a time, by clicking it. Each time you select a deck, the computer will indicate that you won a certain amount of game money. Sometimes the computer will indicate that you won but also lost money. It is not possible to know in advance how much you may win and how much you might lose. You will find that out as you play.

You are absolutely free to switch from one deck to another at any time. The goal of the game is to win as much money as possible and avoid losing money as much as possible. You begin the game with 2,000 game points which is given to you as a loan. The red bar will remind you how much money was loaned to you at the beginning of the task and which you will need to return at the end. The green bar indicates your current earnings throughout the game. In other words, it keeps score of all your gains and losses and tells you your net gains at any given moment.

The only hint we can give you is this: Out of these four decks of cards, there are some decks that are 'bad' and others that are 'good' decks. To win you should try to avoid the 'bad' decks. No matter how much you find yourself losing, you can still succeed if you avoid the 'bad' decks. Note that the computer does not change the order of the decks once the game begins. It does not make you lose at random, or make you lose money based on your previous choice.

References

Baron-Cohen, S., Wheelwright, S., Hill, J., Raste, Y., & Plumb, I. (2001). The 'Reading the Mind in the Eyes' Test Revised Version: A study with normal adults, and adults with Asperger Syndrome or High-Functioning Autism. *Journal of Child Psychiatry and Psychiatry*, *42*, 241–252.

Bechara, A., Damasio, A.R., Damasio, H., & Anderson, S. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*, 7–15.

Brown, W.S., Anderson, L.B., Symington, M.F., & Paul, L.K. (2012). Decision-making in individuals with agenesis of the corpus callosum: Expectancy-valence in the Iowa Gambling Task. *Archives of Clinical Neuropsychology*, *27*, 532–544.

Busemeyer, J.R., & Stout, J.C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, *14*, 253–262.

Dohmen, T., Huffman, D., Schupp, J., Falk, A., Sunde, U., & Wagner, G.G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association, 9*, 522–550.

Edwards, W. (1961) Probability learning in 1000 trials. *Journal of Experimental Psychology*, *4*, 385–394.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., Hertwig, R., Stewart, T., West,R., & Lebiere, C. (2010). A choice prediction competition for choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15–47.

Erev, I., & Haruvy, E. (2015). Learning and the economics of small decisions. In J.H. Kagel, & A.E. Roth (Eds.), *The Handbook of experimental economics, volume 2* (pp. 638–716). Princeton, NJ: Princeton University Press.

Figner, B., Mackinlay, R.J., Wilkening, F., & Weber, E.U. (2009). Hot and cold cognition in risky decision making: accounting for age and gender differences in risk taking. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*, 709–730.

Frey, R., Pedroni, A., Mata, R., Rieskamp, J. & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances, 3*, 1–13.

Hills, T.T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science, 21*, 1787–1792.

Johnson, S.A., Yechiam E., Murphy, R.M., Queller, S., & Stout, J.C. (2006). Motivational processes and autonomic responsivity in Asperger's disorder: Evidence from the

Iowa gambling task. *Journal of the International Neuropsychological Society*, *12*, 668–676.

Lejuez, C.W., Read, J.P., Kahler, C.W., Richards, J.B., Ramsey, S.E., Stuart, G.L., Strong, D.R., & Brown, R.A. (2002). Evaluation of a behavioral measure of risk-taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, *8*, 75–84.

Levin, I.P., Hart, S.S., Weller, J.A., & Harshman, L.A. (2007). Stability of choices in a risky decision making task: A 3-year longitudinal study. *Journal of Behavioral Decision Making, 20*, 241–252.

Lin, J-T. (1980). Approximating the normal tail probability and its inverse for use on a pocket calculator. *Applied Statistics, 38*, 69–70.

March, J.G. (1996). Learning to be risk averse. *Psychological Review, 103*, 309–319.

Mussey, J.L., Travers, B.G., Klinger, L.G. & Klinger, M.R. (2015). Decision-making skills in ASD: performance on the Iowa Gambling Task. *Autism Research*, *8*, 105–114.

Teodorescu, K., & Erev I. (2014). Learned helplessness and learned prevalence: exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science, 25*, 1861–1869.

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys, 14*, 101–118.

Weller, J.A., Levin, I.P., Shiv, B., & Bechara, A. (2007). Neural correlates of adaptive decision making for risky gains and losses. *Psychological Science*, *18*, 958–964.

Worthy, D.A., Hawthorne, M.J., & Otto, A.R. (2013). Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models. *Psychonomic Bulletin & Review*, *20*, 364–371.

Yechiam, E., Arshavski, O., Shamay-Tsoory, S.G., Yaniv, S., & Aharon, J. (2010). Adapted to explore: Reinforcement learning in autistic spectrum conditions. *Brain and Cognition*, *72*, 317–324.

Yechiam, E., & Ert, E. (2011). Risk attitude in decision making: In search of trait-like constructs. *Topics in Cognitive Science*, *3*, 166–186.

Yechiam, E., & Telpaz, A. (2013). Losses induce consistency in risk taking even without loss aversion. *Journal of Behavioral Decision Making*, *26*, 31–40.

Yechiam, E., Telpaz, A., Krupenia, S., & Rafaeli, A. (2016). Unhappiness intensifies the avoidance of frequent losses while happiness overcomes it. *Frontiers in Psychology*, *7*, 1703.

Zou, G.Y. (2007). Toward using confidence intervals to compare correlations. *Psychological Methods*, *12*, 399–413.