MAX PLANCK INSTITUTE
FOR RESEARCH ON
COLLECTIVE GOODS

Silvia Angerer
E. Glenn Dutcher
Daniela Glätzle-Rützler
Philipp Lergetporer
Matthias Sutter

Discussion Paper
2021/16

THE FORMATION OF RISK PREFERENCES THROUGH SMALL-SCALE EVENTS

The formation of risk preferences through small-scale events*

Silvia Angerer[1], E. Glenn Dutcher[2], Daniela Glätzle-Rützler[3], Philipp Lergetporer[4] and Matthias Sutter[5,3]

[1] UMIT TIROL – Private University for Health Sciences and Health Technology

[2] Ohio University

[3] University of Innsbruck

[4] Technical University of Munich

[6] Max Planck Institute for Research on Collective Goods Bonn, University of Cologne and IZA Bonn

Abstract

Understanding how humans deal with risk is important, as it prevails in practically all dimensions of life. Large, macroeconomic shocks have been shown to influence risk-taking. We study in an experiment with 759 subjects whether small-scale, seemingly negligible, events also affect the formation of risk preferences. In line with a reinforcement learning framework, we find that subjects who won a random lottery took significantly more risk in a second lottery almost a year later. So, small-scale, random, events significantly affect the formation of risk preferences, while we also show that memories don't matter.

JEL-Code: C91, D01, D83

Keywords: Reinforcement learning, risk preferences, preference formation, memories, experiment

6 December 2024

* We are particularly grateful for very helpful comments from Zvonimir Basic, Roland Bénabou, Johannes Haushofer, Matthias Heinz, Lukas Kiessling, Nathan Maddix, Pietro Ortoleva, Dimtry Ryvkin, Sebastian Schneider, Peter Schwardmann, Paul Smeets, Florian Zimmermann and seminar participants at Princeton University, University of Copenhagen, Max Planck Institute Bonn, and 3rd LABEL – IEPR Conference "Understanding cognition and decision making by children" in Los Angeles. We thank Rudolf Meraner from the South Tyrolean State Board of Education (*Deutsches Bildungsressort Bereich Innovation und Beratung*), the schools' headmasters (Gabriella Kustatscher, Maria Angela Madera, Eva Dora Oberleiter, Brigitte Öttl, Ursula Pulyer, Vally Valbonesi), the parents of the involved children for making this study possible, and the children for participation. Financial support from the government of the autonomous province South Tyrol through grant 315/40.3 is gratefully acknowledged.

1 Introduction

Our past experiences shape who we are and what we do. For instance, large-scale events like negative macroeconomic shocks in the economy (Malmendier and Nagel, 2011) or in the financial sector (Dohmen *et al.*, 2016; Guiso *et al.*, 2018) have been shown to reduce risk-taking even many years after these events. However, shocks on a large national and supranational scale are rare. Smaller-scale events with almost negligible consequences for single individuals occur much more frequently, make up a large portion of an individual's overall set of experiences, and may play a significant role in shaping the formation of risk preferences. Yet, little is known about the influence of such small-scale events on present-day economic decision-making, and whether such seemingly negligible events can have long-lasting effects. Given that small-scale events are more likely to be forgotten than large-scale (macroeconomic) shocks, it is also an open question whether and to what extent economic decisions are shaped by actual experiences or memories thereof (Chew *et al.*, 2020; Zimmermann, 2020; Huffman *et al*., 2022; Enke *et al*., 2024).

In this paper, we examine if risk-taking is causally affected by a randomly determined outcome of a small-scale lottery. In our experiment with 759 children, two lotteries were set apart by almost one year, and we can show a causal effect of winning the first lottery on risk-taking in the second lottery. In contrast to actual experience, (faulty) memories of the first lottery's outcome do not matter.

We study risk preferences because they play an important role in a broad set of economic domains, including health-related and environmental behavior, financial decision-making or entrepreneurship, to name just a few (Bonin *et al*., 2007; Caliendo *et al*., 2007; Dohmen *et al*., 2011; Sutter *et al.*, 2013; Alan *et al.*, 2017; Castillo *et al*., 2018; O'Donoghue and Somerville, 2018). While genetic factors provide a partial explanation for risk attitudes (Karlsson Linnér *et al*., 2019), the environment (of family, peers, and institutions) a child grows up in is also influential (Heckman, 2006; Cunha and Heckman, 2007; Dohmen et al., 2012; Kosse et al., 2020; Falk et al., 2021; Chowdhury et al., 2022). Yet, random, small-scale factors may also impact the formation of risk preferences. In fact, risk preference formation, like other skill accumulation processes, may work through a reinforcement mechanism where seemingly unimportant events could have noticeable effects on risky decision-making in the longer run. However, it is not immediately obvious what may serve as the reinforcement. Small-scale outcomes may easily be forgotten and replaced by a faulty memory and – given our current understanding of the systematic biases in memories (Bénabou and Tirole, 2016) – choices which rely on memories may yield different outcomes than those which rely on the actual experience. Therefore, we not only lack knowledge about whether risk preferences exhibit a path dependence contingent on past small-scale outcomes, but we also have a limited understanding of the influence of actual outcomes and potentially biased memories on risky choices.

This paper empirically tests if (positive or negative) outcomes of a small-scale risky decision *causally* affects subsequent risky investments. We present a lab-in-the-field experiment with 759 children (aged 6-11 years) who played a small-stakes lottery task in two different periods, set apart by ten months. We ran the experiment with children and teenagers (instead of adults) for several reasons. First, childhood is a formative period for non-cognitive skills like economic preferences (e.g., Fehr *et al*., 2008; Sutter *et al*., 2019; Cappelen *et al.*, 2020; Kosse *et al*., 2020; List *et al*., 2023), permitting a unique opportunity to study preference formation. Second, preference formation in childhood and adolescence has long-term consequences for lifetime outcomes with respect to education, wealth, or health (Sutter *et al.*, 2013; Golsteyn *et al*., 2014; Castillo *et al.*, 2018; Epper *et al*., 2020), yet little is known about the role of random, small-scale events on early preference formation. Third, our sample represents almost 90% of all primary school children of a middle-sized town, which implies that the sample is very comprehensive and has hardly any attrition (actually, only 26 subjects from wave 1 did not participate in wave 2, implying an attrition rate of 2%). This is difficult to achieve outside a school context.

We find two main results. First, and most importantly, subjects who won the first lottery take significantly more risk in the second lottery than those who lost in the first lottery. This suggests that small-scale events (which occurred ten months earlier) influence future risky choices. Specifically, the updating process looks similar to reinforcement learning where actual outcomes are used as the reinforcing mechanism (Roth and Erev, 1995). Second, good outcomes (winning the lottery in the first time period) are more likely to be remembered correctly than bad outcomes. Conversely, bad outcomes are more likely to be misremembered as good than good outcomes as bad. This supports the idea that memories can be used as a consumption good (Gilboa *et al.*, 2016) and that they often deviate from actual outcomes (Chew *et al.*, 2020). Moreover, memories don't affect subsequent risky choices over and above actual outcomes (of winning or losing the first lottery).

Our study contributes to the literature on several dimensions. First, we contribute to the economics literature on preference formation (Cunha and Heckman, 2007; Heckman, 2006, 2008). So far, this literature has mainly investigated how important factors, such as the socio-economic status of parents, the social environment more generally, or participation in education programs affect the formation of non-cognitive skills (Dohmen *et al.*, 2012; Kosse *et al.*, 2020; Castillo *et al*., 2024). Importantly, existing studies focus on the effects of events that either constitute major interventions in an individual's life, such as participation in intense education programs (Cappelen *et al.*, 2020; Kosse *et al.*, 2020) or large-scale macroeconomic shocks (Malmendier and Nagel, 2011; Dohmen *et al.*, 2016). We contribute to this literature by studying how the outcomes of related small-scale events with almost negligible consequences *causally* affect the formation of risk preferences.

Second, many of the prior empirical tests identifying a link between outcomes and changes in risky choices focused on investment-type settings where the returns on the investment were unknown and thus beliefs about the return had to be formed (e.g., Malmendier and Nagel, 2011; Gödker *et al.*, 2024). As such, these studies highlight changes in the beliefs about the probabilities and magnitudes of future returns. In our study, the return on the second lottery, as well as the underlying objective probabilities, are fixed and easy to understand. This means that changes in beliefs about the probability of future returns are ruled out as channels for changes in risk-taking. Rather, we show that the random determination of the outcome has a causal effect on subsequent risk-taking. While a few previous laboratory experiments on risky investment choices with known returns have shown that lottery outcomes can influence subsequent risk-taking immediately after their realization (e.g., Imas, 2016), we extend this evidence by adopting a preference-formation perspective. Specifically, we show that the effect of the lottery outcome is remarkably persistent, manifesting almost a year after the first lottery. Additionally, our focus on children as subjects, rather than the typical sample of university students, offers novel insights into early preference formation.

Third, given the role memories play in our setting, our paper also adds to the growing literature in economics on memory formation and motivated beliefs in which beliefs are allowed to be formed incorrectly via an internal supply and demand mechanism (Bénabou and Tirole, 2016; Saucet and Villeval, 2019; Zimmermann, 2020). Central to these studies is the notion that a memory may be different from the realization of an event *if* the demand for a misinterpretation of the past is high enough and the cost of changing it is low enough. The conditions under which this altered memory is used in a decision process are thought to be tied to some forward-looking rationale. For instance, one of the proposed motivated reasons for forming an incorrect memory relies on anticipatory utility (Bénabou, 2015) or future utility gained from a given identity (Alcott *et al*., 2020). In these cases, a biased memory is formed to meet some future goals. In our study, forward-looking utility is practically eliminated because the outcome of the lottery task does not uncover anything about some future outcome (nor did our participants know that there would be a future lottery). In line with the reasoning that memories can be used as a consumption good (Gilboa *et al.*, 2016), we argue that biased memories can still be formed in our setting, but the motivated reason for doing so is likely due to 'reflective utility', i.e., the utility derived from reflecting upon some past event. Our findings indicate that the anticipatory component may not be a necessary condition for the formation of biased memories. Comparing our finding that faulty memories are not used in the absence of an anticipatory component to studies that have confirmed the use of biased memories when the anticipatory component is present (e.g., Zimmermann, 2020; Gödker *et al.*, 2024) suggests that this component may be necessary for the actual use of faulty memories when making decisions; a finding that provides support for a core assumption in the motivated beliefs literature.

In the following, we introduce our conceptual framework in Section 2. The experimental design is described in Section 3. The results are presented in Section 4. Section 5 concludes the paper.

2 Conceptual framework

We now provide a conceptual framework for the experimental design, which uses concepts of anticipatory utility from Loewenstein (1987), supply and demand of faulty memories from Bénabou and Tirole (2002) and Bénabou (2015), and uses an updating process summarized by Selton and Stoecker (1986) and Erev and Haruvy (2016).

2.1 Setup

We observe an individual over a four-period time horizon, denoted by *t=0,1,2,3*. At *t=0*, the individual chooses to play a risky lottery or accept a safe option. If the risky lottery is played, the outcome $\theta_1 \in \{L_1, W_1\}$ is observed at *t=1*, with corresponding payoffs $P_{W1} > P_{L1}$. If an individual chooses the safe option, he or she receives a secure payoff P_{S1}, with $P_{W1} > P_{S1} > P_{L1}$. Approximately ten months later, at *t=2*, the individual chooses how many tokens, an integer $j \in \{0,1,\dots,K\}$, to invest in a separate risky lottery where the outcome, $\theta_3 \in \{L_3, W_3\}$, is observed at *t=3*.

At *t=1,* individuals can spend lottery winnings on a consumption good and derive utility from consuming this good. Reflective utility is derived from reflecting on the outcome of the lottery and can be represented as a flow of utility, $U(\theta_1)$, which starts at *t=1*.[1] Good outcomes (i.e., winning) generate a positive flow and bad outcomes (i.e., losing) generate a negative flow.

2.2 Formation and consumption of memories

Reconstructive memory (Bartlett and Bartlett, 1995) claims that when a past event is recalled, it may not be remembered perfectly because the individual must reconstruct this past event. The process of reconstruction can lead to systematic errors and inaccurate memories (Roediger, 2001). Bénabou and Tirole (2002) introduce an internal supply and demand mechanism that determines what is remembered, which may affect reconstruction. In our reflective utility framework, individuals wish to reflect on positive outcomes (Gilboa *et al.*, 2016), giving a motivated reason to misremember.[2] That is, a utility-maximizing agent is motivated to either forget bad outcomes ("the event never occurred") or alter the memory of a bad outcome and "choose" to remember it as good by reconstructing the memory in a more favorable light.

Conjecture 1: Correct memory recall is more likely for good outcomes than bad outcomes, and bad outcomes are more likely to be misremembered as good.

2.3 Memories and Choices

Reflective utility provides a motivated reason to hold incorrect memories. Incorrect memories formed from motivated reasoning have been found to lead to motivated beliefs (e.g., Bénabou and Tirole, 2016), which in turn affect choices (e.g., Zimmerman, 2020). A key difference between those settings and ours is that the motivated reason to form an incorrect memory in those studies is forward-looking, while in ours, biases are formed because of a motivation to "consume" good memories. Thus, the bias is not formed to serve some future goal. If a biased memory does not serve a useful purpose in guiding choices, however, then it is natural to assume the actual outcome will be used. Studies examining links between past experiences and subsequent choices typically rely on features of prominent learning models (Thorndike, 1898), which predict that actions that were reinforced by positive (negative)

[1] When a choice is made at *t=0,* there is no knowledge of the existence of the lottery at *t=2*, which eliminates anticipatory utility. However, reflective utility can be considered the counterpart to anticipatory utility (Loewenstein, 1987), where reflective utility accumulates after, instead of before, the event, but follows the same logical structure of Loewenstein (1987).

[2] Psychophysical studies provide a rationale for the demand side. A review by Zoccola and Dickerson (2012) highlights the negative health effects of recalling a past negative event. Chadwick *et al.* (2016) found that subjects who focused on good events had more positive health effects.

outcomes *increase* (*decrease*) the likelihood that this action will be chosen again in the future.[3] However, for many small-scale choices where reinforcement learning is meant to apply, an individual may misremember a relevant outcome from the past that may be used as the reinforcement.

Both neuroscience (Dudai, 2004) and reinforcement-learning studies suggest that correct memories are not a necessary condition for past experiences to affect choices. It is assumed that feedback can be internalized and generates an intuitive response to similar stimuli where the intuitive response does not use cognitively taxing memory recall. In neuroscience, it has been shown that tasks can be learned without having remembered learning them (Corkin, 2002). Because of the path-dependent nature of these learning models, whether an actual outcome ten months prior is misremembered may be immaterial to the decision maker. This leads to our second conjecture.

Conjecture 2: The actual outcome (experience) of the first lottery will affect the choice in the second lottery in a reinforcing manner, and, after controlling for this outcome, the (correct or incorrect) memory will have no additional effect.

Even though this conjecture is founded on empirical studies, it remains largely untested in our domain. It is certainly reasonable to imagine no effect from memories or outcomes of small-scale events on choices, which leaves this as an empirical exercise.

This framework can be extended to uncover how the feedback may have an asymmetric effect on subjects, conditional on heterogeneous initial risk preferences. Roth and Erev (1995) introduce an experimentation parameter where reinforcement alters the propensity to play the chosen *and* adjacent strategies. However, if the initial strategy is close to the boundary of the strategy space, then the updating of adjacent strategies will also be bounded. In our specific setting, individuals who are either the least risk-averse or the most risk-averse are at the boundaries. For the least risk-averse, there is an upper bound on how much positive feedback can be expected to alter their behavior given the limited adjacent strategies to update. The opposite is true for the most risk-averse subjects. This intuition is reflected in our final conjecture.

[3]This does not imply that humans always have perfect recall. A key feature of many learning models is a "forgetting" parameter where recent *experiences* are more influential. However, memories are not allowed to differ from the experience. See Erev and Haruvy (2016).

Conjecture 3: The actual outcome (experience) of the first lottery will have an asymmetric effect on choices in the second lottery in the sense that positive experiences will have a stronger impact on the relatively more risk-averse subjects than the least risk-averse subjects.

3 Experimental design and implementation

Our study was conducted in two waves (in December 2011 and September/October 2012), set apart by 10 months. It was run in all primary schools in the city of Meran (population 38,000) in the province of South Tyrol, Italy.[4] The large majority of parents (87%) of all primary school children permitted their children to participate in the experiment, and only a single child opted out. In total, 1,122 children participated in both waves. In the first wave, they attended grades one to four and were 6-10 years old. In the second wave, they were in grades two to five (Italian primary school has five grades). The children's decisions were always incentivized with tokens that could be exchanged for small presents (e.g., stickers, wristbands, sweets, pencils).

In the first wave, risk preferences were elicited with three binary decision tasks. In each task, a subject could either receive 2 tokens with certainty or play a lottery with a 50% probability of winning 3, 4, or 5 tokens and a 50% probability of losing all tokens (experimental instructions are reprinted in the appendix and the decision sheet shown in Figure A1). The three tasks were first explained – either in ascending order (starting with the 3 tokens priced lottery) or descending order (starting with the 5 tokens priced lottery) – and afterward, a subject had to decide for each task whether to play the lottery or take the safe payoff. One task was randomly selected and implemented for payment. In the following, we analyze primarily the 759 subjects who either chose to play the lottery once or twice. For this subset, it is randomly (and thus exogenously) determined whether the actual payment depends on a lottery outcome and, if so, the outcome of the lottery. We thus exclude in our main analysis all subjects who either always chose the safe amount of 2 tokens or chose the lottery every time (in both cases, it is no longer random whether they are going to play a lottery). For the 759 subjects considered here, we can distinguish between the exogenously created groups of those who did not play the lottery

[4] The experiment reported here was part of a larger research project, which investigated economic decision making of children. In total, we ran a series of six experimental sessions in the academic years 2011–12 and 2012–13. The first risk game was conducted in the first experimental session and the second risk game in the fourth experimental session. Besides risk attitudes, we also elicited time preferences (twice), altruistic preferences and cooperation (twice) within the six data collection waves. Note that the second risk game was the first measure we repeated in the series of experimental sessions we conducted. Adding the behavior in those other games (which were typically set apart by about three months) as further controls in our regressions does not change our main results. A review of outcomes in the other games is provided in Sutter *et al*. (2019).

(because in the randomly selected decision task, the subject chose not to play the lottery), those who experienced a win in the lottery, and those who experienced a loss.

In the second wave, risk preferences were measured with a simple investment task (Charness and Gneezy, 2010). Each subject was endowed with five tokens and had to decide how many of these tokens to invest in a lottery that doubled the number of invested tokens with a 50% probability, while with a 50% probability, the subject lost the investment. Non-invested tokens were safe earnings for the subject (see Figure A2 in the appendix for the decision sheet).

After the decision in wave two, we also elicited the memory of the outcome of wave 1 (where the risk-elicitation task had been called the SMILEY game) by asking the subjects the following question (answered orally).

"*In the last school year, we played the SMILEY-game. In this game, you had to decide three times whether to play the SMILEY-game or whether to take two tokens for sure. Can you remember whether you won, lost, or chose the safe option in the part that was chosen in the end?*

- *Won*
- *Lost*
- *Chose to play the safe option*
- *Don't know*"

By asking this question, we can examine whether memories affected subsequent choices, whether memories were correct, and whether actual outcomes of wave 1 or retrieved memories thereof mattered more for choices in wave 2.[5]

Table 1 presents some descriptive statistics and a balance check. Of the 759 subjects, 45.19% were female, 144 were 7/8-year-olds, 182 were 8/9-year-olds, 208 were 9/10-year-olds and 225 were 10/11-year-olds. From the table, we see that the outcome – winning or losing the lottery conditional on playing the lottery – was random (221 wins vs. 220 losses; *p-value* =

[5] The elicitation of a subject's memory was not incentivized. Since we focus on the role of small-scale events that may also be easily forgotten (in contrast to large macro-economic shocks, for example, that one can hardly forget), we believe that adding incentives for recalling memories would convert the small-scale situation into a more significant event, something which we wanted to avoid deliberately. We also see no reason why subjects might have incentives to misreport their memories when not paying for correct recall. For instance, any concerns of social desirability effects do not apply in our setting. Note also that our results on memory (as shown below) are in line with prior results and our conceptual framework presented in Section 2. Zimmermann (2020) shows that memories become more accurate when correct recall is incentivized, which is attributed to the individual engaging in (cognitively) more costly behavior to recall the actual outcome. In our setting, this implies that unincentivized elicitation gives us the best chance to observe if faulty memories are used in the updating process and thus relevant for choices, as we can expect more incorrect memories in the absence of incentives for correct recall. In the limit, if memories were 100% accurate, it would actually be impossible to disentangle whether and to what extent outcomes and memories influence risk preference formation, as both outcomes and memories would be identical.

1.000 in a binomial test) and that there is practically no difference in the demographic characteristics across all three groups.

Table 1: Descriptive statistics and balancing tests

Background characteristics	**WON**	***N***	**LOST**	***N***	**SAFE OPTION**	***N***	***p-value***
Age in years	8.62 [1.18]	221	8.65 [1.23]	220	8.62 [1.16]	318	0.944
Female (=1)	50.22 %	221	43.18 %	220	43.08%	318	0.204
German school (=1)	46.61 %	221	46.36 %	220	47.48%	318	0.962
No. of siblings	1.28 [0.96]	213	1.35 [0.94]	215	1.31 [1.09]	310	0.734
No. of risky choices Game 1[a]	1.75 [0.44]	221	1.79 [0.41]	220	1.43 [0.50]	318	0.000
IQ[b]	21.01 [4.00]	220	20.98 [4.44]	220	20.84 [4.45]	312	0.889
Income father[c]	1,852 [442]	198	1,839 [427]	185	1,795 [433]	276	0.319
Income mother[c]	1,656 [381]	170	1,700 [432]	151	1,646 [374]	237	0.398
Father self-employed (=1)[c]	21.10 %	199	15.59 %	186	14.13%	276	0.126
Mother self-employed (=1)[c]	5.88 %	170	9.15 %	153	6.3%	238	0.467

Notes: Columns 2,4 and 6: sample means for winners, losers, and those who did not play the lottery (but chose the safe amount); standard deviations in brackets (for non-dummy variables); Columns 3, 5, and 7: number of observations winners, losers and those who chose the safe amount; Column 8: *p*-values from multinomial logistic regressions with the outcome (won, lost, safe amount) as categorical dependent variable.
The numbers of observations sometimes differ between rows because of missing values.
[a] Number of risky choices in the three binary decision tasks in wave 1. This is significantly different between those who won or lost the lottery (a total of 441 subjects) and those who chose the safe option (318 subjects). Those who played the lottery less often are more likely to be randomly selected for payment of the safe amount, and hence those with the safe amount are on average less risk-taking. Yet, note that all actual outcomes are randomly determined and hence identification of the causal effects of luck works well.
[b] The IQ was measured with a modified version of Raven's Colored Progressive Matrices.
[c] We did not gather information about parents' professions for all children.

4 Results

In Figure 1, we display the average level of risk investment in wave 2 on the vertical axis. On the horizontal axis, we distinguish subjects by the actual (randomly determined) outcome they experienced in wave 1. Winners in wave 1 invested, on average, 2.56 tokens in wave 2. However, losers in wave 1 and those who chose the safe option invested 11%-12% less, with an average of only 2.25 tokens, respectively 2.27 tokens (*p-values* < 0.01; two-sided Mann-Whitney U-tests). This result is consistent with conjecture 1 above. There is no significant difference in the average investment between losers and those who chose the safe option (*p-value* = 0.982; two-sided Mann-Whitney U-test).

Figure 1: Average number of tokens invested in wave 2, conditional on the outcome of wave 1. Error bars indicate mean ± standard error.

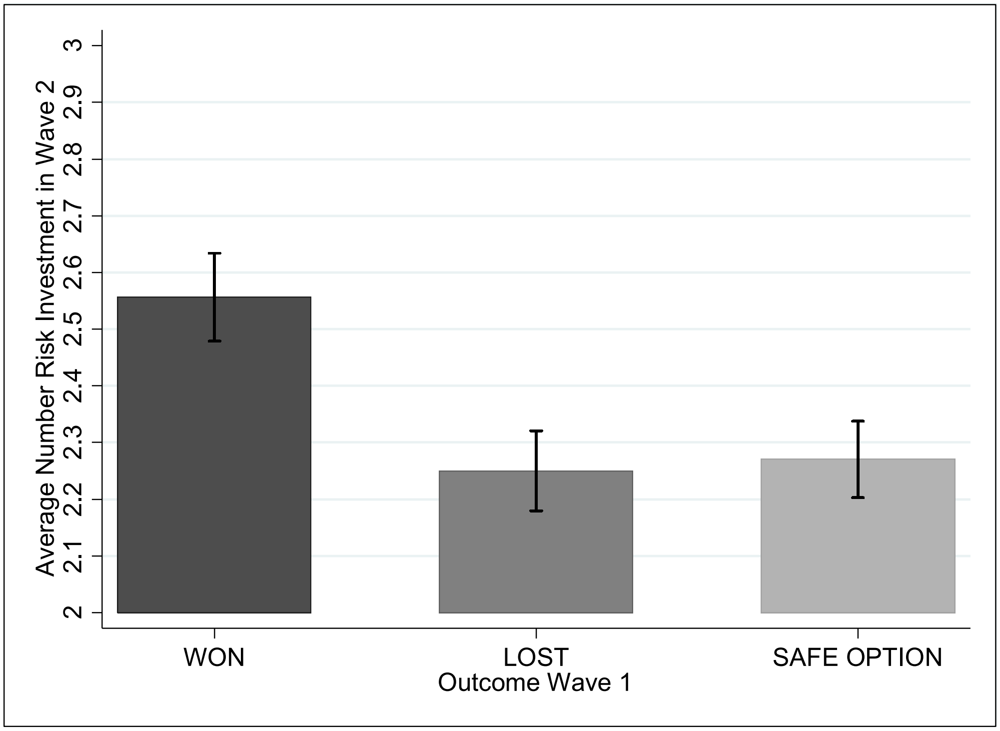

In Table 2, we confirm the insights from Figure 1. The table is based on an OLS regression (that yields very similar results to an ordered probit model; see Table A1 in the appendix). It shows that if subjects experienced a win in wave 1, they made significantly larger risk investments in wave 2 (see positive and statistically significant coefficients on *Won Wave 1* in columns (1), (3), and (4) as well as the *p*-values in the post-estimation Wald tests). The estimated coefficients (around 0.25 to 0.3 in the top row) represent about 10-12% of the average investment in wave 2 (see Figure 1).

Table 2: OLS regression with level of risk investment in wave 2 as the dependent variable

VARIABLES	(1)	(2)	(3)	(4)
Won Wave 1 (=1)	0.262**		0.243**	0.323***
	(0.104)		(0.103)	(0.118)
Lost Wave 1 (=1)	-0.047		-0.039	0.026
	(0.114)		(0.116)	(0.117)
# risky choices Wave 1	0.076	0.105	0.076	0.072
	(0.086)	(0.081)	(0.086)	(0.091)
Memory WON (=1)		0.140	0.082	0.090
		(0.110)	(0.111)	(0.109)
Memory LOST (=1)		-0.072	-0.014	-0.027
		(0.116)	(0.113)	(0.112)
Stake size Wave 1				-0.091*
				(0.047)
Age				-0.006
				(0.042)
Female (=1)				-0.274***
				(0.094)
No. of siblings				0.018
				(0.044)
German school (=1)				-0.089
				(0.094)
Relative IQ				-0.172
				(0.260)
Average # risky choices Wave 1 in class				-0.240
				(0.159)
Fraction of winners of Wave 1 in class				0.043
				(0.323)
Constant	2.162***	2.149***	2.145***	3.099***
	(0.139)	(0.145)	(0.144)	(0.612)
Observations	759	759	759	733
R-squared	0.014	0.006	0.015	0.038
Wald tests (p-values)				
H_0: No diff. between Won and Lost	0.004		0.014	0.014
H_0: No diff. between memory WON and memory LOST		0.106	0.490	0.409

Notes: Robust standard errors clustered on the class level in parentheses, with 97 classes in total. ***Significance at 1% level. **Significance at 5% level. *Significance at 10% level. The number of observations in column (4) is smaller than in the other specifications because some data on the additional controls are missing.

Our dependent variable is the level of risk investments made in wave 2 (ranging from 0 to 5). Independent variables are the following: *Won Wave 1* (=1 if they won in wave 1, 0 otherwise), *Lost Wave 1* (=1 if they lost in wave 1, 0 otherwise), *# of risky choices Wave 1* (= the number of risky choices made in Wave 1, either 1 or 2), *Memory WON* (=1 if they believed they won, 0 otherwise), and *Memory LOST* (=1 if they believed they lost, 0 otherwise). This means that the omitted outcome-category is “Safe option played” and the omitted memory-category is "Memory OTHER". We also control for *Stake size Wave 1,* ranging from 1=Low (3 tokens to be won) to 3=High (5 tokens to be won), depending on which lottery was actually paid out. Further control variables are *Age*, *Female, No. of siblings* (number of siblings)*, German school* (=1 if subject is in a German-speaking school, 0 if Italian-speaking school), *Relative IQ* (number of correct answers in Raven's test relative to the average number of correct answers within each grade), and the average number of risky choices as well as the fraction of winners within a subject's class to account for potential peer effects. In column (1), we only include experience variables, while in column (2) we only include memory variables. Column (3) includes both and column (4) adds the further controls.

We also see from Table 2 that the effects of memory are never found to be significant – see the coefficients on *Memory WON* and *Memory LOST* in columns (2), (3), and (4), as well as the post-estimation Wald tests. So, it is actual outcomes that have a formative effect on risk-taking almost one year later, not memories.

Memory and outcomes may be so strongly correlated that no additional explanatory power comes from including both in the model. To check if memories are biased, Figure 2 presents the share of subjects who remembered the outcome of wave 1 correctly, broken up by the actual outcome of wave 1. The figure shows that 44% of the subjects who won, but only 30% of those who lost, remembered the outcome correctly (*p-value* = 0.003, χ^2-test). Out of those who did not play the lottery (because they had chosen the safe amount in the task randomly selected for payment), only 5% remembered correctly choosing the safe option. Thus, having played the lottery results in a more accurate memory. Among those who played the lottery, winners are much better at recalling the outcome correctly than losers, which is consistent with conjecture 1 from our conceptual framework. Note, however, that even in the group of children playing the lottery, more than half don't remember the lottery outcome, which may explain why the memory of the outcome of wave 1 does not correlate significantly with risk-taking in wave 2. This provides additional evidence that the event was perceived as small-scale and didn't stick in many children's memories.

Figure 2: Share of subjects remembering the outcome correctly, conditional on the outcome of wave 1. Error bars indicate mean $\pm$ standard error.

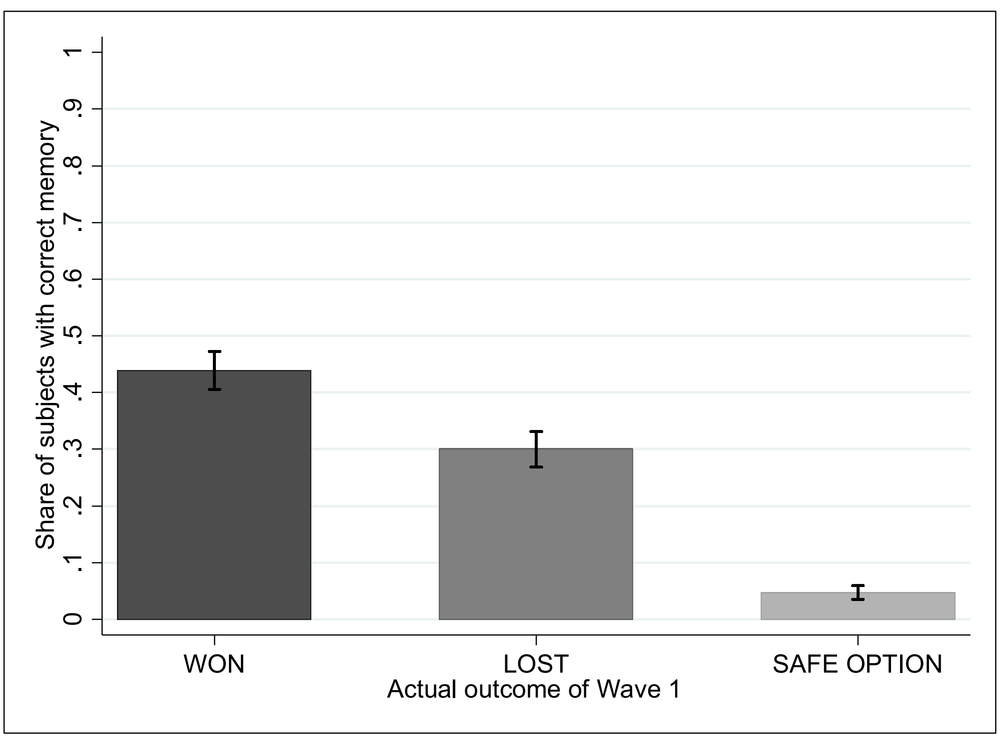

In Figures A3 and A4 in the appendix, we show that the general difference in memory conditional on outcomes is present in each age group – except for the youngest cohort – and for the medium and high stake sizes of lotteries played in wave 1. The figures also show that, overall, there is a tendency for older children to be better able to recall the outcome (*p-value* = 0.066, Cuzick's Wilcoxon-type test for trend) and that the stake size increases the accuracy of memory for those who won (*p-value* = 0.006, Cuzick's Wilcoxon-type test for trend). In Figure A5 in the appendix, we show that the type of misremembering is also related to actual outcomes. Out of those who actually lost the lottery in wave 1, 15% thought they had won (panel A, middle bar). This is different from those who won in wave 1, where only 4% thought they had lost (panel B, left bar) (*p-value* = 0.000, χ^2-Test for difference in memories between those having won and those having lost). This suggests a positivity bias in children's memories and a potential demand for good memories. Despite these noteworthy patterns of memories, our analysis shows that they are not instrumental in determining subsequent risk-taking (see Table 2).

Before concluding, we address conjecture 3 from our conceptual framework. To do so, we consider now the 302 children who chose the lottery in all three choices of wave 1. Recall that we had excluded these children in our preceding analysis because for them it was not random whether they played the lottery or not (as they always played it). Yet, for conjecture 3 it is useful to look at these children. In Figure 3 we show the risky investment in wave 2 for children who chose once, twice, or three times, the lottery in wave 1 *and* who had played the lottery (thus we exclude here subjects who picked the lottery once or twice and the random draw assigned them the safe amount in one of the three tasks). In line with our conjecture 3, we note from Figure 3 that the least risk-averse subjects (those picking always the lottery) do not react significantly to winning or losing the lottery. Their investment in wave 2 does not differ significantly between those experiencing a loss or a win in wave 1. The contrary is true for the two more risk-averse groups who picked the lottery once or twice. For them, those winning in wave 1 invest more in wave 2 than those losing in wave 1 (*p-value* < 0.1 and *p-value* < 0.05 for those who chose one or two risky choices in wave 1 respectively; *p-value* < 0.01 if we pool both subsets).

Figure 3: Average number of tokens invested in wave 2, conditional on the outcome of wave 1 and conditional on having played the lottery in wave 1. Error bars indicate mean ± standard error.

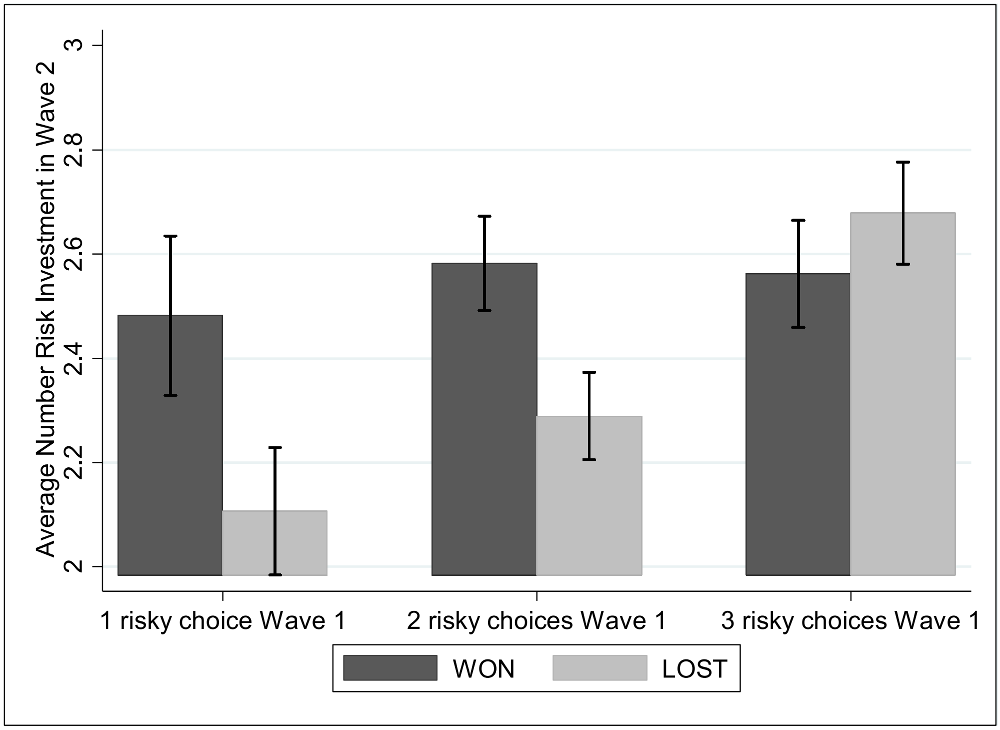

5 Conclusion

In this paper, we have studied whether seemingly negligible outcomes from the past can have a causal effect on current behavior. Small-scale experiences that occur on a daily basis may shape future behavior through reinforcement, which might create a path-dependency. It is unclear, however, whether small-scale events actually have a causal impact on subsequent behavior and what serves as the potentially reinforcing mechanism. Events from the past might affect behavior through the actual outcomes, but also through memories of these events. Given that individuals may perceive small-scale events as insignificant, they may easily be forgotten. Even if not, memories may be incorrect (Chew *et al.*, 2020; Zimmermann, 2020) or systematically biased due to motivated reasoning (Bénabou, 2015; Bénabou and Tirole, 2016). For macroeconomic shocks, memories have been found to influence risky choices – in particular, because they shift beliefs about future realizations of risk – but it was an open question if small-scale events (that are much more frequent than large shocks) have any influence at all and if so, whether and to what extent actual outcomes and memories matter.

Based on experimental data from 759 subjects, and in line with our conceptual framework, we have found that subjects who won in the first wave of a small-stakes lottery take significantly more risks in the second wave 10 months later. Since the outcome of the first wave was random, this suggests that small-scale outcomes have a measurable, *causal* effect on subsequent behavior. This is independent of whether subjects remember the lottery and its outcome and – if they do – which outcome they remember. The manner in which small-scale experiences are used is consistent with the idea that choices can take the form of intuitive (or automatic) decisions that are formed by previous experience even in the absence of remembering them (e.g., see Epstein, 1994). If outcomes are remembered at all, we have found that good outcomes are more likely to be remembered correctly than bad outcomes (and we provide the first evidence with children about this pattern). This finding on its own is not novel, however, it is often assumed that the effects of memories are due to their anticipatory utility. In our study, we eliminate anticipatory components to memory formation almost completely by design and investigate memory formation due to "reflective utility", where memories serve as a consumption good. Because the anticipatory component of memory formation is practically absent in our case and we find faulty memories are not significant when making choices, comparing our results with those testing motivated beliefs suggests that the anticipatory component may be necessary for the use of faulty memories and their influence on subsequent choices when the event was actually experienced.

References

Alan, S., N. Baydar, T. Boneva, T. F. Crossley and S. Ertac (2017). Transmission of Risk Preferences from Mothers to Daughters. *Journal of Economic Behavior & Organization* **134**, 60-77.

Allcott, H., L. Boxell, J. Conway, M. Gentzkow, M. Thaler and D. Yang (2020). Polarization and Public Health: Partisan Differences in Social Distancing During the Coronavirus Pandemic. *Journal of Public Economics* **191**, 104254.

Bartlett, F. C. and F. C. Bartlett (1995). *Remembering: A Study in Experimental and Social Psychology*: Cambridge University Press.

Bénabou, R. (2015). The Economics of Motivated Beliefs. *Revue d'Ȅconomie Politique* **125**(**5**), 665-685.

Bénabou, R. and J. Tirole (2002). Self-Confidence and Personal Motivation. *Quarterly Journal of Economics* **117**(**3**), 871-915.

Bénabou, R. and J. Tirole (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives* **30**(**3**), 141-164.

Bonin, H., T. Dohmen, A. Falk, D. Huffman, and U. Sunde (2007). Cross-Sectional Earnings Risk and Occupational Sorting: The Role of Risk Attitudes. *Labour Economics*, **14**(**6**), 926-937.

Caliendo, M., F. M. Fossen, and A. S. Kritikos (2007). Risk Attitudes of Nascent Entrepreneurs: New Evidence from an Experimentally Validated Survey. *Small Business Economics* **32**(**2**), 153-167.

Cappelen, A., J. List, A. Samek and B. Tungodden (2020). The Effect of Early-Childhood Education on Social Preferences. *Journal of Political Economy* **128**(**7**), 2739-2758.

Castillo, M., J. L. Jordan and R. Petrie (2018). Children's Rationality, Risk Attitudes and Field Behavior. *European Economic Review* **102**, 62-81.

Castillo, M., J. A. List, R. Petrie and A. Samek (2024). Detecting Drivers of Behavior at an Early Age: Evidence from a Longitudinal Field Experiment. *Journal of Political Economy*, forthcoming.

Chadwick, A. E., P. M. Zoccola, W. S. Figueroa and E. M. Rabideau (2016). Communication and Stress: Effects of Hope Evocation and Rumination Messages on Heart Rate, Anxiety, and Emotions after a Stressor. *Health Commun* **31**(**12**), 1447-1459.

Charness, G. and U. Gneezy (2010). Portfolio Choice and Risk Attitudes: An Experiment. *Economic Inquiry* **48**(**1**), 133-146.

Chew, S., W. Huang and X. Zhao (2020). Motivated False Memory. *Journal of Political Economy* **128(10)**, 3913-3939.

Chowdhury, W., M. Sutter and K. Zimmermann (2022). Economic Preferences across Generations and Famliy Clusters: A Large-Scale Experiment in a Developing Country. *Journal of Political Economy* **130**(**9**), 2361-2410.

Corkin, S. (2002). What's New with the Amnesic Patient H.M.?. *Nature Reviews Neuroscience* **3**(**2**), 153-160.

Cunha, F. and J. Heckman (2007). The Technology of Skill Formation. *American Economic Review* **97**(**2**), 31-47.

Dohmen, T., A. Falk, D. Huffman and U. Sunde (2012). 'The Intergenerational Transmission of Risk and Trust Attitudes. *Review of Economic Studies* **79**(**2**), 645-677.

Dohmen T., D. Huffman, U. Sunde, J. Schupp and G. G. Wagner (2011). Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences. *Journal of the European Economic Association* **9**, 522-550.

Dohmen, T., H. Lehmann and N. Pignatti (2016). Time-Varying Individual Risk Attitudes over the Great Recession: A Comparison of Germany and Ukraine. *Journal of Comparative Economics* **44**(**1**), 182-200.

Dudai, Y. (2004). The Neurobiology of Consolidations, or, How Stable Is the Engram?. *Annual Review of Psychology* **55**, 51-86.

Enke, B., F. Schwerter and F. Zimmermann (2024). Associative Memory, Beliefs and Market Interactions. *Journal of Financial Economics*, forthcoming.

Epper, T., E. Fehr, H. Fehr-Duda, C. T. Kreiner, D. D. Lassen, S. Leth-Petersen and G. N. Rasmussen (2020). Time Discounting and Wealth Inequality. *American Economic Review* **110**(**4**), 1177-1205.

Epstein, S. (1994). Integration of the Cognitive and the Psychodynamic Unconscious. *American Psychologist* **49**, 709-724.

Erev, I. and E. Haruvy (2016). Learning and the Economics of Small Decisions. in (Kagel, J. H. and A. E. Roth Eds.), *The Handbook of Experimental Economics,* 638-700, Princeton: Princeton University Press.

Falk, A., F. Kosse, P. R. Pinger, H. Schildberg-Hörisch and T. Deckers (2021). Socio-Economic Status and Inequalities in Children's Iq and Economic Preferences. *Journal of Political Economy*, **129**(**9**), 2504-2545.

Fehr, E., H. Bernhard and B. Rockenbach (2008). Egalitarianism in Young Children. *Nature* **454**, 1079-1083.Garcia, J. (1990). Learning without Memory. *Journal of Cognitive Neuroscience* **2**(**4**), 287-305.

Gilboa, I., A. Postlewaite and L. Samuelson (2016). Memorable Consumption. *Journal of Economic Theory* **165**, 414-455.

Gödker, K., P. Jiao and P. Smeets (2024). Investor Memory. *Review of Financial Studies*, forthcoming.

Golsteyn, B. H., H. Grönqvist and L. Lindahl (2014) Adolescent Time Preferences Predict Lifetime Outcomes. *Economic Journal* **124**, 739-761.

Guiso, L., P. Sapienza and L. Zingales (2018). Time Varying Risk Aversion. *Journal of Financial Economics* **128**(**3**), 403-421.

Heckman, J. J. (2006). Skill Formation and the Economics of Investing in Disadvantaged Children. *Science* **312**(**5782**), 1900-1902.

Heckman, J. J. (2008). Schools, Skills, and Synapses. *Economic Inquiry* **46**(**3**), 289-324.

Huffman, D., C. Raymond and J. Shvets (2022). Persistent Overconfidence and Biased Memory: Evidence from Managers. *American Economic Review* **112(10)**, 3141-3475.

Imas, A. (2016). The Realization Effect: Risk-Taking after Realized versus Paper Losses. *American Economic Review* **106(8)**, 2086-2109.

Karlsson Linnér, R., P. Biroli, E. Kong, … and J. P. Beauchamp (2019). Genome-Wide Association Analyses of Risk Tolerance and Risky Behaviors in over 1 Million Individuals Identify Hundreds of Loci and Shared Genetic Influences, *Nature Genetics* **51**, 245-257.

Kosse, F., T. Deckers, P. Pinger, H. Schildberg-Hörisch and A. Falk (2020). The Formation of Prosociality: Causal Evidence on the Role of Social Environment. *Journal of Political Economy* **128**(**2**), 434-467.

List, J. A., R. Petrie, and A. Samek (2023). How Experiments with Children Can Inform Economics. *Journal of Economic Literature* **61(2)**, 504-564.

Loewenstein, G. (1987). Anticipation and the Valuation of Delayed Consumption. *Economic Journal* **97**(**387**), 666-684.

Malmendier, U. and S. Nagel (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?. *Quarterly Journal of Economics* **126**(**1**), 373-416.,

O'Donoghue, T. and J. Somerville (2018). Modeling Risk Aversion in Economics. *Journal of Economic Perspectives* **32**(**2**), 91-114.

Roediger, H. (2001). Reconstructive Memory, Psychology Of. In Smelser, N. J. and P. B. Baltes (Eds.), *International Encyclopedia of the Social & Behavioral Sciences,* 12844-12849, Pergamon.

Roth, A. E. and I. Erev (1995). Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Run. *Games and Economic Behavior* **8**(**1**), 164-212.

Saucet, C. and M. C. Villeval (2019). Motivated Memory in Dictator Games. *Games and Economic Behavior* **117**, 250-275.

Selten, R. and R. Stoecker (1986). End Behavior in Sequences of Finite Prisoner's Dilemma Supergames. A Learning Theory Approach. *Journal of Economic Behavior and Organization* **7**, 47-70.

Sutter, M., M. G. Kocher, D. Glätzle-Rüetzler and S. T. Trautmann (2013). Impatience and Uncertainty: Experimental Decisions Predict Adolescents' Field Behavior. *American Economic Review* **103**(**1**), 510-531.

Sutter, M., C. Zoller and D. Glätzle-Rützler (2019). Economic Behavior of Children and Adolescents – a First Survey of Experimental Economics Results. *European Economic Review* **111**, 98-121.

Thorndike, E. L. (1898). Animal Intelligence: An Experimental Study of the Associative Processes in Animals. *Psychological Review Monograph Supplements* **2**(**4**).

Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review* **110**(**2**), 337-361.

Zoccola, P. and S. Dickerson (2012). Assessing the Relationship between Rumination and Cortisol: A Review. *Journal of Psychosomatic Research* **73(1)**, 1-9.

Online Appendix

Experimental instructions (translated from German/Italian)

Note: Italic font is used for the instructions to the experimenter.

Risk preferences Wave 1

Register the order of explanation (blue first or green first) in the computer.

Good morning. My name is … Today I have prepared a game for you. In this game you can earn tokens. With these tokens you can buy some presents in our shop. The game consists of 3 parts. The blue part, the yellow part, and the green part *(when mentioning the parts please point at the respective decision sheets)*.

The game works as follows:

In the blue part you have to decide whether you prefer receiving 2 tokens (*please point at the tokens on the decision sheet*) for sure, in this case please tick THIS box (*point at the respective box*), or whether you prefer playing the SMILEY-game, in that case please tick THAT box (*point at the respective box*). The SMILEY-game works as follows: I have here a card with a smiley face and a card with a sad face (*show cards*). I will mingle the two cards under the table, place them face down at the table and you can draw one card *(demonstrate mingling and put cards on the table – IMPORTANT: don't let the child draw one card)*. If you draw the smiley face, you will receive 3 tokens (*point at the box with the three tokens next to the smiley face on the decision sheet*). However, if you draw the sad face, you won't receive any tokens (*point at the box with no tokens in it next to the sad face on the decision sheet*). This is the blue part. Could you please repeat the rules of the game? (*If the child is unable to repeat, please explain the game again; the child has to be able to repeat the correct procedure of the game autonomously.*)

The yellow part is very similar to the blue part. Here you see the decision sheet for the yellow part. Again, there are 2 tokens on the left-hand side, but on the right-hand side there are 4 tokens now. What do you think will happen if you tick THIS box? (*please point at the box with the 2 tokens for sure*) What do you think will happen if you tick THAT box? (*please point at the box with the SMILEY-game and the four tokens*; *the child has to answer the questions correctly, otherwise the experimenter has to repeat the explanation*).

The green part is very similar to the blue and yellow part. Here you see the decision sheet for the green part. Again, there are 2 tokens on the left-hand side, but on the right-hand side there are 5 tokens now. What do you think will happen if you tick THIS box? (*please point at the box with the 2 tokens for sure*) What do you think will happen if you tick THAT box? (*please* *point at the box with the SMILEY-game and the five tokens*; *the child has to answer the questions correctly, otherwise the experimenter has to repeat the explanation*).

It is important to note that at the end only one of the three parts counts. That means that you will receive the tokens for one of the three parts only. After your decisions I will mingle the three decision sheets under the table (*please demonstrate*; *Attention: you have to handle the sheets such that the child is not able to see the color of the respective sheet! You need to cover the three parts with an additional large-format sheet when placing the sheets on the table for drawing*) and then you can draw one of the three parts. (*In what follows, adapt the explanation to the order in which you draw the sheets:*) If you draw the blue part (*demonstrate the drawing of the first sheet*), only the blue part counts and you will receive the tokens for this part only. The other two parts do not count in this case. If you, for example, ticked THIS box (*please point at the box with the 2 tokens for sure*), what happens? If you, for example, ticked THAT box ((*please point at the box with the SMILEY-game*), what happens (*child must answer both questions correctly; IMPORTANT: give both examples*!)? If you however draw the yellow part (*demonstrate the drawing of the second sheet*), only the yellow part counts and you will receive the tokens for the yellow part only. The other two parts do not count in this case. If you draw the green part (*demonstrate the drawing of the third sheet*), only the green part counts and you will receive the tokens for the green part only. The other two parts do not count in this case. However, you need to make a decision for each of the three parts because you don't know yet which part will be drawn at the end of the game. Could you please repeat the last part? Will you receive the tokens for all three parts? Do you need to make a decision for each of the three parts? (*If the child answers incorrectly the experimenter has to repeat the explanation of this part.*)

Please take your decision for each of the three parts now (*place the decision sheets side by side on the table; the child should fill out the decision sheets from left to right)*. Start with this part (*point at the first decision sheet (blue or green, depending on the order of explanation)*) and continue with this part (*point at the second decision sheet*) and finally make your decision in this part (*point at the third decision sheet*). Take as much time as you need. In the meantime, I will turn around so that I don't disturb you. Just call me when you are done.

Figure A1: Decision Sheets "Risk Preferences Wave 1"

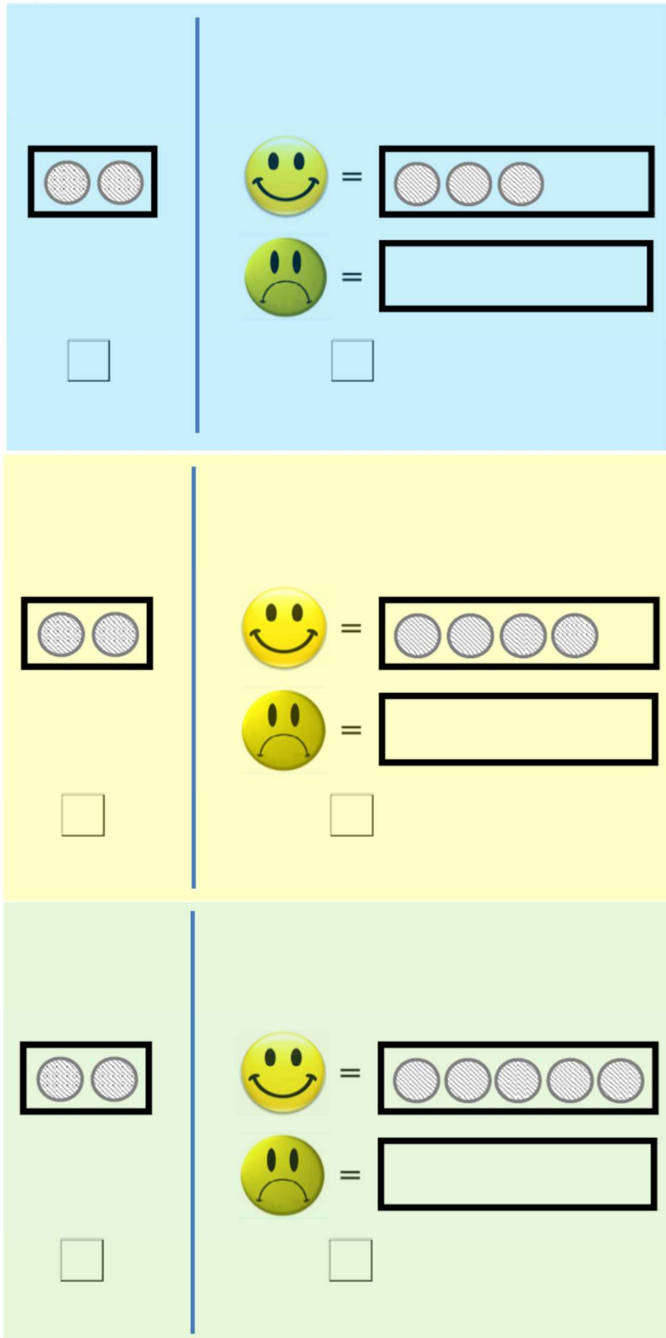

Risk preferences Wave 2

Good morning. My name is … Today's game works as follows:

At the beginning you will receive 5 tokens *(please place the 5 tokens in front of the child)*. You have to decide how many of these 5 tokens you want to keep for sure and with how many of these tokens you want to play the "treasure"-game. You have to put the tokens you keep for sure in this box (*point at the left box*). Likewise, you must put the tokens with which you want to play the treasure-game in that box (*point at the right box*). Each token that you put in the treasure-game will be doubled. The rules of the treasure-game are as follows: Here I have two cards. On this card you see a full treasure chest and on the other card there is an empty treasure chest (*show the respective cards*). I will mingle the two cards under the table and then I will put the cards on the table upside down (*please demonstrate*; *Attention: you have to mingle the cards, such that the child is not able to see the picture on the respective card)*. Then you can draw one of the cards. If you, for example, draw the full treasure chest, (*point at the full treasure chest on the decision sheet*), then you will receive all the tokens from this box. On the other hand, if you draw the empty treasure chest (*point at the empty treasure chest on the decision sheet*) then you will lose all the tokens from this box. At the end you will receive the tokens that you keep for sure (*point at the left box*) and the tokens that you win in the treasure game (*point at the right box*).

Let's consider an example: If you, for instance, want to keep one token for sure and play the treasure-game with the other 4 tokens, what do you have to do? (*Answer of the child: "I have to put 1 token in the left box and 4 tokens in the right box"; please let the child demonstrate this*) How many tokens will be added to this box? (*point at the right box*; *answer of the child: "4"; please demonstrate!*) What happens next? How does the treasure-game work? (*Child has to repeat the rules of the game*). How many tokens will you win if you draw the full treasure chest? (*Answer of the child: "8 tokens"*). And how many tokens will you receive in total? (*Answer of the child: "9"*). Exactly. You will receive 8 tokens from the treasure-game plus 1 additional token which you kept for sure. What happens if you draw the empty treasure chest? (*Answer of the child: "I lose all the tokens of the treasure-game"*) Exactly. How many tokens will you receive in total? (*Answer of the child: "1"*) Exactly. This was only an example. Let's consider another example: Could you please explain the rules of the game if you want to keep 4 tokens for sure and play the treasure-game with 1 token? (*The child has to recapitulate the game with the new example*). What happens if you, for instance, put all your 5 tokens in this box? (*point at the right box; let the child recapitulate the game*) What happens if you, for instance, put all your 5 tokens in this box? (*point at the left box; let the child recapitulate the game*). Could you please repeat the rules of the game?

Please take your decision now. You have to put the tokens which you want to keep for sure in this box (*point at the left box*) and the tokens with which you want to play the treasure-game have to be put in that box (*point at the right box*). Take as much time as you need for your decision. In the meantime, I will turn around so I don't disturb you. Just call me when you are done.

Figure A2: Decision Sheet "Risk Preferences Wave 2"

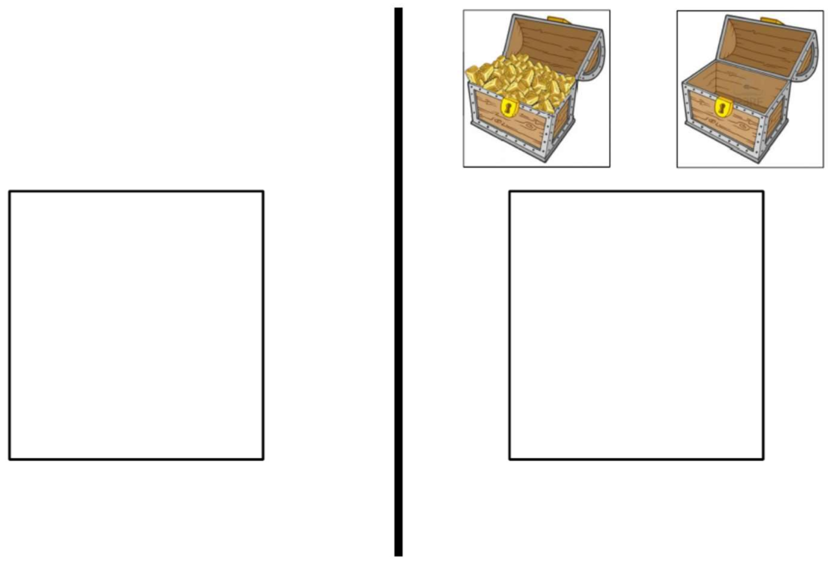

Additional Tables and Figures

Table A1: Ordered probit regression with the number of risk investment in wave 2 as dependent variable

VARIABLES	(1)	(2)	(3)	(4)
Won Wave 1 (=1)	0.250***		0.235**	0.310***
	(0.095)		(0.094)	(0.110)
Lost Wave 1 (=1)	-0.021		-0.016	0.041
	(0.102)		(0.105)	(0.108)
# risky choices Wave 1	0.054	0.087	0.054	0.057
	(0.078)	(0.073)	(0.078)	(0.083)
Memory WON (=1)		0.122	0.068	-0.078
		(0.099)	(0.100)	(0.100)
Memory LOST (=1)		-0.052	-0.004	-0.016
		(0.104)	(0.102)	(0.103)
Stake size Wave 1				-0.087**
				(0.043)
Age				0.004
				(0.039)
Female (=1)				-0.236***
				(0.085)
No. of siblings				0.019
				(0.041)
German school (=1)				-0.084
				(0.086)
Relative IQ				-0.109
				(0.241)
Average # risky choices Wave 1 in class				-0.238*
				(0.143)
Fraction of winners of Wave 1 in class				0.013
				(0.296)
Constant cut1	-1.796***	-1.773***	-1.781***	-2.575***
	(0.143)	(0.151)	(0.150)	(0.570)
Constant cut2	-0.525***	-0.509***	-0.510***	-1.296**
	(0.127)	(0.132)	(0.131)	(0.564)
Constant cut3	0.354***	0.363***	0.368***	-0.405
	(0.129)	(0.133)	(0.133)	(0.569)
Constant cut4	1.171***	1.177***	1.186***	0.439
	(0.141)	(0.145)	(0.146)	(0.562)
Constant cut5	1.851***	1.858***	1.868***	1.134**
	(0.144)	(0.148)	(0.148)	(0.564)
Observations	759	759	759	733
Pseudo R-squared	0.005	0.002	0.005	0.012
Wald tests (p-values)				
H_0: No diff. between Won and Lost	0.004		0.014	0.013
H_0: No diff. between memory WON and memory LOST		0.139	0.569	0.470

Notes: Robust standard errors clustered on the class level in parentheses, with 97 classes in total. ***Significance at 1% level. **Significance at 5% level. *Significance at 10% level. The number of observations in column (4) is smaller than in the other specifications because some data on the additional controls are missing.

Figure A3: Share of subjects remembering the outcome correctly conditional on the outcome of wave 1, by age groups. Error bars indicate mean ± standard error.

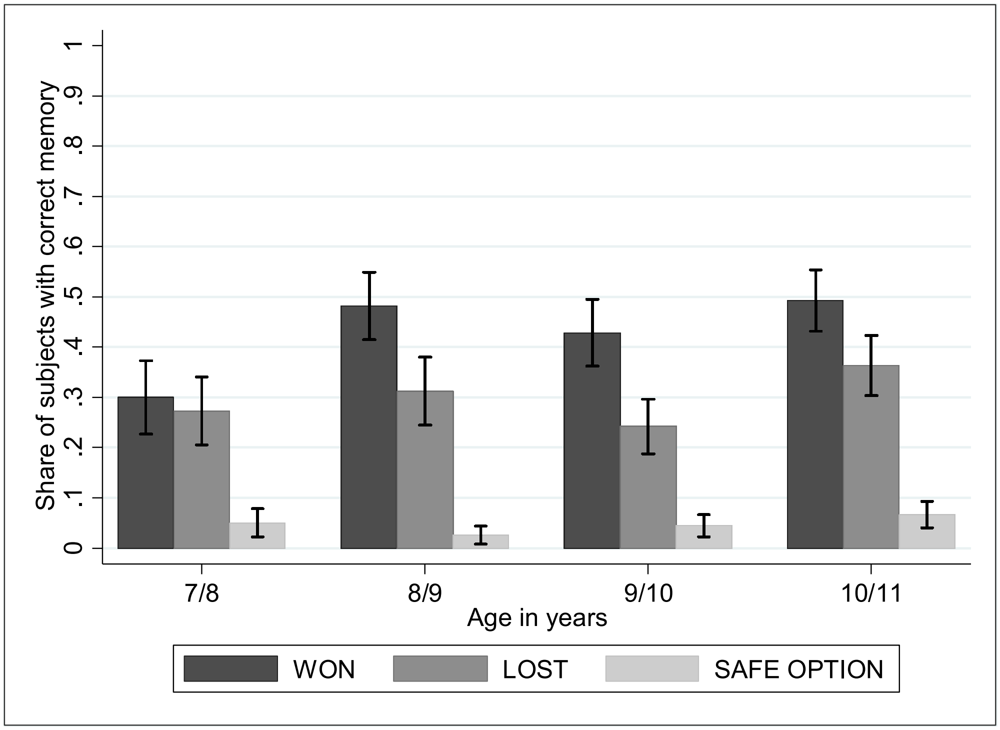

Figure A4: Share of subjects remembering the outcome correctly conditional on the outcome of wave 1, by the stake size of the lottery played in wave 1. Error bars indicate mean ± standard error.

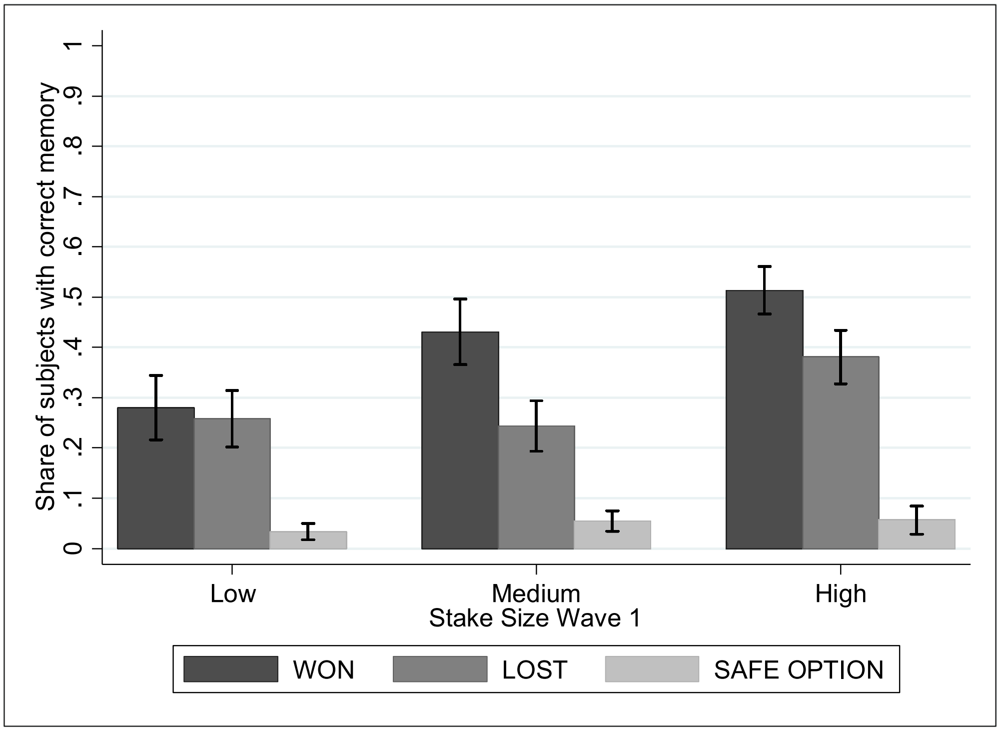

Notes:
In the sample that we analyze, 232 played the "Low", 259 the "Medium" and 268 the "High" stake size lottery in wave 1.

Figure A5: Memory of the outcome of wave 1, conditional on the actual outcome of wave 1 on the horizontal axis. Error bars indicate mean ± standard error.