

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

Working Paper

The Formation of Risk Preferences through Small-Scale Events

CESifo Working Paper, No. 9270

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Angerer, Silvia; Dutcher, E. Glenn; Glätzle-Rützler, Daniela; Lergetporer, Philipp; Sutter, Matthias (2021) : The Formation of Risk Preferences through Small-Scale Events, CESifo Working Paper, No. 9270, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/245451>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Formation of Risk Preferences through Small-Scale Events

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

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

The Formation of Risk Preferences through Small-Scale Events

Abstract

Large, macroeconomic shocks in the past have been shown to influence economic decisions in the present. We study in an experiment with 743 subjects whether small-scale, seemingly negligible, events also affect the formation of risk preferences. In line with a reinforcement learning model, we find that subjects who won a random lottery took significantly more risk in a second lottery almost a year later. The same pattern emerges in another experiment with 136 subjects where the second lottery was played more than three years after the first lottery. So, small-scale, random, events affect the formation of risk preferences significantly.

JEL-Codes: C910, D010, D830.

Keywords: reinforcement learning, risk preferences, preference formation, experiment.

Silvia Angerer
UNIT – Private University for Health Sciences,
Medical Informatics and Technology / Austria
silvia.angerer@unit.at

E. Glenn Dutcher
Ohio University
Ohio / USA
dutcher@ohio.edu

Daniela Glätzle-Rützler
University of Innsbruck / Austria
Daniela.Ruetzler@uibk.ac.at

Philipp Lergetporer
ifo Institute / Munich / Germany
lergetporer@ifo.de

Matthias Sutter
Max Planck Institute for Research on Collective Goods
Bonn / Germany
matthias.sutter@coll.mpg.de

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 preferences and thus future choices. This line of reasoning is supported by Heckman’s (2006) observation that non-cognitive skills – like economic preferences – are formed over time with everyday experiences. Yet, little is known about the influence of such past, 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 (Huffman *et al.*, 2019; Chew *et al.*, 2020; Enke *et al.*, 2020; Zimmermann, 2020).

In this paper, we examine if risk taking is causally affected by a randomly determined outcome of a small-scale lottery. In our main experiment, the two lotteries are set apart by almost one year, and in another experiment there are more than three years between playing the first and the second lottery. We find a strong effect of winning or losing the first lottery on the choices made almost one year, respectively more than three years, later.

Risk preferences have been at the core of formal explanations for economic choices since at least Arrow (1965) and have subsequently played a large role in the analysis of insurance, crime and punishment, tax evasion, principal-agent relationships, entrepreneurship and financial decisions, to name just a few (Caliendo *et al.*, 2007; Dohmen *et al.*, 2011; Alan *et al.*, 2017; O’Donoghue and Somerville, 2018). Moreover, risk preferences have also been found to determine individual life outcomes, such as educational outcomes, health status, occupational choices, and wealth (Bonin *et al.*, 2007; Sutter *et al.*, 2013; Castillo *et al.*, 2018). Given the outsized role risk preferences play in economic decision-making, it is imperative that we improve our understanding of the endogenous formation of risk preferences.

While genetic factors provide a partial explanation for an individual’s risk preferences (Karlsson Linnér *et al.*, 2019), empirical regularities – which find that choices in risky domains change over time – indicate that the environment also plays a role (see Sutter *et al.*, 2019, for a review of age-related changes in risk attitudes). That role has been partially verified by the

observance of systematic adjustments individuals make to risky choices after experiencing large-scale shocks such as the past financial crisis or the great depression in the inter-war period (a similar effect may be expected for the current Covid-19 crisis). Malmendier and Nagel (2011), for example, find that individuals who had previously experienced economic depressions and relatively lower returns from stocks were less likely to invest in stocks relative to those who had experienced good economic conditions and better returns on stock investments. Similar patterns of behavior have been reported in Dohmen *et al.* (2016), Necker and Ziegelmeyer (2016) and Guiso *et al.* (2018).

Even though these papers provide evidence about changes in risk-taking behavior as a consequence of different experiences on financial markets, they focus in their explanation on changes in subjects' beliefs about the returns from risky investments (a channel that has also been identified by G dker *et al.*, 2019), but not on an individual's attitude towards taking risks. While belief-updating can be domain specific, a subject's attitudes may be more general and may be shaped by many small-scale events that are much more frequent than large-scale shocks. 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 present in memories (B nabou and Tirole, 2016) – choices which rely on memories may yield different outcomes than those which rely on the actual experience. Therefore, not only do we not know if risk preferences exhibit a path dependence contingent on past small-scale outcomes, we also lack a general understanding of the influence of actual outcomes and potentially biased memories on risky choices.

This paper tests empirically if (positive or negative) feedback on a small-scale risky decision *causally* affects subsequent risky investments. We use two lab-in-the-field experiments for this purpose. In the main dataset, 743 subjects (aged 7-11 years) played a small-stakes lottery task in two different periods, set apart by ten months. In the second dataset, intended to show the robustness of our main result, 136 subjects (aged 11-15 years) played two lotteries where the second one was played more than 3 years after the first one. Both experiments were run with children and teenagers. They serve as an ideal sample for our research question 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.*, 2021), permitting a unique opportunity to study preference

formation. Second, endogenous 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), yet little is known about the role of random, small-scale events on early preference formation. Third, our main sample represents almost 90% of all primary school children of a middle-sized town, which implies that our sample is very comprehensive. This can hardly be achieved outside a school context. Moreover, sampling children and teenagers in a school context allows us to study longer-term treatment effects with almost no attrition.

Our experimental setting consisted of choices in two different time periods. In the first time period of our main experiment, 1,122 subjects could choose to play up to three lotteries with various risks or accept a safe alternative. One of these lotteries was randomly chosen for payment and if the subject had chosen to play that lottery (N=743), they were informed of the outcome (i.e., whether they won or lost) and were given their prizes. Ten months later, the subjects played a similar, but previously unannounced, lottery task. Afterwards, they were also asked about their memories of the outcomes from the first lottery task. We use the outcome from the first lottery task as the actual experience along with their (possibly biased) memory of the outcome and examine the influence of both on risk taking in the second lottery. Three features of this experiment are worth highlighting. First, winning or losing the lottery was random (as it was determined by a 50:50 draw). This random nature is important in establishing a causal link between past lottery outcomes (winning or losing) and decisions in the second lottery.¹ Second, the school context enables matching choices across both lotteries with only 2% attrition that does not depend on the randomly determined outcome of the first lottery. Third, because all outcomes and probabilities of the lottery are easy to understand and known prior to the risky choice, changes in beliefs are eliminated as an explanation for potential changes in behavior.

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. We confirm this core result in a second experiment, run with 136 teenagers who played two lotteries set apart by more than 3 years (exactly 37 months). The lotteries differed slightly from those in the main experiment, but winning and losing was again randomly determined. Those who won the lottery more than 3 years prior take significantly more risk in the second lottery, thus confirming the same main effect in a different sample (from a different country), using a

¹ Levitt (2021) studies how a coin flip influences major life decisions like quitting a job or ending a relationship. These are examples for a macro-type of shock and they are in stark contrast to the small-scale events we are interested in.

different lottery task. These results suggest that even small-scale events (which occurred ten months, respectively 3 years earlier) do influence future risky choices in a predictable way. 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). Even though memories are selectively formed, in the main experiment we find that they have no effect on 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, often large-scale, 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; Castillo *et al.*, 2020; Kosse *et al.*, 2020). 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).² 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.

² Take a simple example. Let there be two states of the world where good states generate above average returns and bad states generate below average returns. These studies test if feedback from a prior risky investment alters one’s belief about the state of the world where prior good (bad) outcomes may lead one to believe in a higher probability of a good (bad) state (see, e.g. Gödker *et al.*, 2019).

Third, our focus on small-scale outcomes can be viewed as testing for the boundaries of the literature on large-scale shocks by establishing whether the behavioral response to large-scale events observed in those studies also prevails in (much more frequent) small-scale events. In Section 2, we provide a conceptual framework for how past outcomes may affect current choices. Our results suggest that previously identified behavior in large-scale domains likely extends to outcomes of any scale along the proposed conceptual framework.

Finally, 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 an individual holds about a prior event may be different from the realization of this event *if* the demand for a misinterpretation of the past is high enough and the cost of changing it low enough. The conditions under which this altered memory is used in a decision process is 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)³ or future utility gained from a given identity (Alcott *et al.*, 2020). In these cases, a biased memory is formed to meet some future goal and is thus instrumental in a future decision. In our study, most aspects of forward-looking utility are 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 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’, or 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., Gödker *et al.*, 2019, Zimmermann, 2020) 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.

³ Bénabou (2015) also identifies instrumental beliefs as another rationale where these (mis)beliefs are held in order to attain some future advantage (such as overconfidence leading to others’ willingness to follow a leader’s orders; Schwarzmann and van der Weele, 2019).

2 Conceptual framework

We now provide a conceptual framework for reinforcement learning, which adds 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 Erev and Haruvy (2016). The framework is intended to guide our empirical strategy in determining if small-scale past events can influence future choices in a predictable way and if the actual outcome, or the (possibly faulty) memory thereof, are more important in determining choices. This outline is not meant to set up a test for what specific psychological underpinnings drive behavior, but rather to provide a logically consistent description of what behavior may look like. The general framework used can encompass many different psychological concepts found to influence behavior in similar settings. We will first outline why memories may be biased, followed by how these memories or past experiences influence subsequent choices in similar situations.

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, at $t=1$, the outcome $\theta_1 \in \{L_1, W_1\}$ is observed. It is determined by a random draw that has a 50:50 chance of winning (W_1) or losing (L_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, $\theta_2 \in \{L_2, W_2\}$, is observed at $t=3$. This second lottery doubles investment j with 50% probability, and with 50% probability, the investment is lost.

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. We assume the individual derives two forms of utility from the outcome: an immediate consumption utility and backward-looking reflective utility. 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).⁴

⁴ The general idea that individuals derive utility from thinking about the future and reflecting upon the past can serve as a general umbrella for many more specific concepts. For instance, motivations such as ego utility (Kőszegi, 2006) can be thought of as altering the anticipatory component in accordance with the instrumental value reasoning from Bénabou and Tirole (2002). Similarly, regret minimization (e.g., Loomes and Sugden, 1982) can be thought to give rise to the demand for altering memories due to the consumption value of reflective utility. We

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 $t=1$. 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

The theory of reconstructive memory (Bartlett and Bartlett, 1995) claims that when a past event is recalled, the individual must reconstruct this past event. The process of reconstruction can lead to systematic errors and inaccurate memories.⁵ Bénabou and Tirole (2002) introduce an internal supply and demand mechanism that determines what is remembered. In our reflective utility framework, individuals wish to reflect on positive outcomes (Gilboa *et al.*, 2016), giving a motivated reason to mis-remember.⁶ That is, even though altering a memory may be costly, 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.

In our setting, this implies that good outcomes – winning the lottery – are more likely to be remembered and bad outcomes – losing the lottery – forgotten or misremembered as good. We can assign some probability of correct recall, $\Pr(r|\theta_1)$, which is dependent upon the outcome (θ_1). Based on the logical framework presented thus far, $\Pr(r|W_1) > \Pr(r|L_1)$. This leads to our first conjecture.

Conjecture 1: Correct memory recall will be more likely for good outcomes than bad outcomes.

2.3 Outcomes and choices

Studies on motivated beliefs require a two-stage process. First, biased memories are formed based on a motivated reason to do so and, second, these memories must be used in a decision process. Our first conjecture indicates we expect biased memories to form for motivated reasons, but different than the motivated beliefs literature, this bias is formed because of a motivation to “consume” good memories. Thus, the bias is not formed to serve some goal in the future and thus there is no reason to believe they will be used, for which reason an alternative

do not discount the importance of these specific concepts, however our goal is to first identify empirically if a general pattern is observed before we concern ourselves with specific factors underlying these patterns.

⁵ See Roediger (2001) for an overview.

⁶ Psychophysical studies provide a rationale for the demand side. A review by Zoccola and Dickerson (2012) highlights the negative health effects stemming from recalling a past negative event while Chadwick *et al.* (2016) found the converse to be true as well: those subjects who focused on good events had more positive health effects.

approach must be pursued. Studies examining links between past experiences and subsequent choices typically rely on features of reinforcement learning (Thorndike, 1898), which predicts that actions that were reinforced by positive (negative) outcomes *increase (decrease)* the likelihood that this action will be chosen again in the future. It is natural to extend this framework to our setting given the nature of the feedback received from the first lottery.⁷ Following Roth and Erev (1995), we allow the probability that investment j (how many tokens to invest in the second risky lottery) will be chosen in time period $t+1$ by an individual to be equal to:

$$p^j(t+1) = \frac{q^j(t+1)}{\sum_{s=0}^K q^s(t+1)}$$

where $q^j(t+1)$ is some propensity of the individual to invest j tokens at time $t+1$. In our setting, $t+1$ is the time period ($t=2$) when the second lottery choice was made. Reinforcement learning predicts changes in propensities according to the dynamic process: $q^j(t+1) = q^j(t) + x$ where x is the payoff gained from choosing a given investment at t , which is the payoff from the choice in the first lottery.⁸ Propensities for all other investments $m \neq j$ (which were not chosen) are not updated, or $q^m(t+1) = q^m(t)$.

If experiences, and not biased memories, are used for reinforcement learning, the number of tokens invested at $t=2$ is increasing if $\theta_1 = W_I$ and non-increasing if $\theta_1 = L_I$ because x is increasing in $\theta_1 = W_I$ and non-increasing if $\theta_1 = L_I$. Because past reinforcement learning studies have been conducted in environments where virtually perfect recall of the reinforcing experience exists, the possibility and potential influence of biased memories have not been considered.⁹ 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.¹⁰

⁷ More complex models, such as I-SAW (Erev *et al.*, 2010) or its variants, have advanced beyond the basic reinforcement learning model presented here. We stay with the simpler framework since our goal is not to lay the framework for structural estimates, but to provide the logical foundations to generate the behavioral predictions.

⁸ Other studies have documented that “transfer” of learning will occur where what is learned in one setting (our first lottery) will be transferred to a related setting (our second lottery). Cooper and Kagel (2003, 2005) highlight such a transfer in economic games.

⁹ For an overview of the many laboratory studies where memory of the most recent event is used, see Erev and Haruvy (2016). This does not imply that humans always have perfect recall. In fact, a key feature of many learning models is a “forgetting” parameter where recent *experiences* are more influential. However, it is still assumed that the reinforcement is based on *recall* of the most recent events and memories are not allowed to differ from the experience.

¹⁰ There is a long history of learning and memory usage. Outside of economics, it has been found that learning can occur without memory in the way we typically define it in humans (Garcia, 1990). In humans, the most prominent explanations in neuroscience arise from patient H.M. who, after a brain surgery, could not remember learning new skills, but nonetheless learned them (Corkin, 2002). Learning is also possible without any feedback about previous choices’ outcomes (Weber, 2003).

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 (positive or negative reinforcement) 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). In reinforcement learning, choices have a path dependence that initializes with an actual experience that guides future choices in similar situations, which lead to further feedback, and so on. Because of the path-dependent nature, whether an actual outcome ten months prior is misremembered or not may be immaterial to the decision maker according to this line of reasoning. 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 we have highlighted logical and empirical support for the second conjecture, the conjectured behavior remains largely untested in our domain. Applying this line of reasoning to our setting may be too far of a reach as the proposed mechanisms may not be well suited to explain behavior from small-scale outcomes that occurred months or years earlier. Other studies have shown that false memories can be planted in subjects' minds so that they believe them to be true and they subsequently act according to this planted memory in a reinforcing manner (Clifasefi *et al.*, 2013). Though these studies are even further from the current one, an alternative conjecture could claim that memories (in addition to outcomes) do impact choices, which leaves this as a largely 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 the reinforcement received alters the propensity to play the chosen strategy and those that are similar to, or adjacent to, this strategy.¹¹ Conceptually, let a strategy j be played that yields a payoff of x . As before, this payoff will be used to update $q^j(t+1)$, but will also be used to update the propensity of adjacent strategies in $t+1$. The extent to which the chosen and adjacent strategies are updated will determine the

¹¹ Take as an illustration the different strategies available in the ultimatum game. If the proposer chooses to give 30% of the pie to the responder, the reinforcement received from this choice will update the propensity to play 30%, but will also update strategies close to it, such as 25% and 35%, for example.

resulting propensities. To account for this, Roth and Erev introduce $0 \leq \varepsilon \leq 1$ and a modified updating process of $q^j(t+1) = q^j(t) + (1 - \varepsilon)x$. If $\varepsilon = 0$, the updating process is as before. However, if $\varepsilon > 0$, then εx will be used to update the propensities of strategies adjacent (or similar) to j . For instance, let strategies $j+1$ and $j-1$ be the two adjacent strategies which are updated and let these strategies receive an equal share of εx , or each receives $\frac{\varepsilon}{2}x$. This will produce updated propensities $q^{j+1}(t+1) = q^{j+1}(t) + \frac{\varepsilon}{2}x$ and analogously for $q^{j-1}(t+1)$. This implies that strategies $j+1$ and $j-1$ need not be played in order to receive reinforcement. Importantly, if adjacent strategies are to be updated, such strategies must exist. If the chosen 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 the least risk averse are already using the riskiest strategies. Positive feedback that might typically be used to update adjacent strategies will not do so in this case given that the chosen strategy is already at the boundary, i.e., 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 given they are the furthest from this boundary. Positive feedback may lead to the largest change in behavior for this group. This intuition is reflected in our final conjecture.

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.

The third conjecture does not imply that the least risk-averse subjects will not respond at all to the feedback from the first lottery; it simply implies their response will not be as strong as the response from the relatively more risk-averse subjects. That is, we still expect to find support for the other conjectures regardless of the initial levels of risk aversion exhibited.

3 Experimental design of the main experiment

The two risk preference games were conducted in two subsequent academic years, the first in December 2011 and the second in September/October 2012 in the city of Meran in the province of South Tyrol, Italy. The large majority of parents (87%) of all primary school children in Meran (that has a population of 38,000 inhabitants) gave permission for their child to participate in the experiment, which implies that our sample is comprehensive and has almost no self-selection. The experimental sessions were conducted during regular school hours and participation was voluntary for children, but all except one child consented to participate.¹² In total, 1,122 children participated in the two experimental sessions. In the first year of our study, they attended grades one to four and were six to ten years old; in the second year of our study they were in grades two to five and ten months older.¹³ The children's decisions were always incentivized with tokens that could be exchanged for small presents (e.g., stickers, wristbands, sweets, pencils).¹⁴

Risk preferences in the first game – which we called the SMILEY-game and to which we will refer to as Game 1 in the following – were elicited with three binary decision tasks. In each task the subject could either get 2 tokens for sure or play a lottery with a 50% probability of winning 3, 4, or 5 tokens and a 50% probability of losing all tokens (see Figure B1 in Appendix B for the experimental instructions and decision sheets). 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 afterwards the subject had to decide for each of the three decision tasks whether to play the lottery. One choice was randomly selected and implemented for payment. In case a subject chose the safe option in the randomly selected decision task, the lottery was not played. In line with our research question and conceptual outline, our analytical sample consists of those 743 subjects who chose to play the lottery in the randomly selected

¹² The experiment 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 (see section 4.3). A review of outcomes in the other games is provided in Sutter *et al.* (2019).

¹³ In Italy, children start primary school at the age of six and it lasts for five years. In the first year of our study, we ran the experiment also with the fifth graders who left primary school after that year, meaning that for them we do not have data from September/October 2012. These fifth graders are not included in the total number of 1,122 participants across both waves of data collection.

¹⁴ In order to avoid zero profits, all children received one extra token, independent of the outcomes in the risk preference game.

decision task, i.e., those who experienced a win or a loss in Game 1. We take the number of risky choices (ranging from one to three) as an indicator of risk tolerance.

The second risk preference game – denoted Game 2 henceforth – was a simple investment task (Charness and Gneezy, 2010). In this task, 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.¹⁵ Non-invested tokens were safe earnings for the subject. We take the number of invested tokens (ranging from zero to five) as an indicator of risk tolerance. After the decision in the investment game, we elicited the memory of the outcome of Game 1 by asking the subjects the following question (that was 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 some 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”*

The answer to this question serves as our variable for a subject’s memory. Together with estimating the impact of the actual outcome in Game 1 on behavior in Game 2, we can check whether – and if so, how – memories of Game 1 outcomes are faulty. 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 intuitive 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

¹⁵ Even though this lottery structure would suggest indifference by a risk-neutral individual, variances in choices are observed and correlated with risk taking in the SMILEY-game. For our study, we simply need a relative comparison and this structure satisfies our goal by eliciting the level of risk investments an individual is making.

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.

The experiments took place in a separate room with several individual workplaces where trained experimenters explained the tasks to the subjects one-on-one and each subject had to repeat the rules of the game in order to check for understanding.¹⁶ In total, 34 out of 743 subjects had some difficulties in understanding the first, the second, or both risk elicitation tasks. For completeness, we include those 34 subjects in our analysis, but note that our results remain unchanged when excluding them (see section 4.3).

4 Results

Our primary interest is in empirically establishing whether actual (randomized) outcomes of a past small-scale event influence future choices. Additionally, we can examine whether memories of past outcomes have an effect. We will use the conceptual framework provided in Section 2 as a guide to proceed in our analysis.

4.1 Descriptive statistics

In total, 1,122 subjects played both games.¹⁷ Out of these, 743 subjects experienced a win or a loss in Game 1 because they had decided to play the lottery in the randomly selected task. For these 743 subjects we have a randomized outcome for one of the lotteries they played, and hence these subjects are our basis for examining whether the outcome of Game 1 (and the memory thereof) affects risk taking in Game 2. In the following, we therefore do not study the choices of the other 379 subjects who chose the safe option in the randomly picked task and thus could not experience a win or a loss in the lottery. Naturally, since the probability of experiencing a win or a loss increases in the number of risky choices, these 379 subjects are more risk averse than the 743 subjects with a random lottery outcome (see Table A.2 in Appendix A), but they do not differ on other observables. Importantly, the identification of the

¹⁶ The individual workplaces were equipped with partition walls and placed with sufficient distance to each other so that children could take their decisions in private.

¹⁷ Note that there was almost no attrition between the two games played. Only 26 subjects who participated in Game 1 did not participate in Game 2 and their background characteristics are not different from the background characteristics of the 1,122 subjects participating in both games (see Table A.1 in Appendix A).

causal effects of winning or losing the lottery in Game 1 on risky investments in Game 2 is not compromised by this selection since winning or losing was randomly assigned *conditional* on playing the lottery. Rather, the fact that less risk-averse subjects constitute our analytical sample is intuitive from an external-validity perspective, because it is exactly those subjects who would select into risky situations (and experience good or bad outcomes) in everyday life.

Of the 743 subjects, 45.76% were female, 183 were 7/8-year-olds, 192 were 8/9-year-olds, 183 were 9/10-year-olds and 185 were 10/11-year-olds. Table 1 summarizes background characteristics, separately for subjects who won and those who lost the lottery in Game 1. From this table, we see that the outcome – winning or losing the lottery – was random (383 wins vs. 360 losses; $p\text{-value} = 0.419$ in a binomial test) and that there is practically no difference in the demographic characteristics across both groups. The only (random) exception is the number of siblings. Even in this instance, the difference is small. On average, the 743 subjects chose to play the risky lottery in 2.27 out of the three tasks in Game 1. Next, we look at their memories ten months later.

Table 1: Descriptive statistics and balancing tests

Background characteristics	Winners	<i>N</i>	Losers	<i>N</i>	<i>p-value</i>
Age in years	8.42 [1.17]	383	8.47 [1.24]	360	0.616
Female (=1)	47.78 %	383	43.61 %	360	0.254
German school (=1)	46.48 %	383	43.33 %	360	0.390
No. of siblings	1.27 [0.95]	371	1.47 [1.04]	350	0.008
No. of risky choices in Game 1 ^a	2.27 [0.70]	383	2.26 [0.67]	360	0.605
IQ ^b	20.41 [4.36]	379	20.42 [4.58]	358	0.668
Income father ^c	1,811 [419]	341	1,840 [446]	304	0.584
Income mother ^c	1,664 [383]	290	1,653 [413]	246	0.483
Father self-employed (=1) ^c	19.01 %	342	14.75 %	305	0.151
Mother self-employed (=1) ^c	6.53 %	291	7.66 %	248	0.609

Notes: Column 2 (4): sample means for winners (losers); standard deviations in brackets (for non-dummy variables); Column 3 (5): number of observations winners (losers); Column 6: p -values from two-sample χ^2 -Tests (for dummy variables) and Mann-Whitney U Tests (for non-dummy variables).

The numbers of observations sometimes differ between rows because of missing values.

^a Number of risky choices in the three binary decision tasks of risk Game 1.

^b The IQ was measured with a modified version of Raven's Colored Progressive Matrices

^c For a detailed description of these variables see the notes in Appendix C. We did not get information about parents' professions for all children.

4.2 Outcomes in Game 1 and memories thereof ten months later

In the second experimental session ten months after Game 1, subjects were asked about their memory of Game 1. Of the subjects who experienced a win or a loss in Game 1 ($N=743$), 54% could not remember the outcome, 1% ($N=7$) stated (incorrectly) that they had not played the lottery in the randomly selected choice task, 15% believed they had lost and 30% believed they had won the lottery.

We will now test conjecture 1 that winning leads to a higher probability of correct recall. Figure 1 depicts the share of subjects who remembered the outcome of the lottery correctly, broken up by the actual outcome of the lottery. The figure shows that 42% of the subjects who won, but only about 26% of those who lost, remembered the outcome correctly ($p\text{-value} < 0.01$, $\chi^2\text{-test}$). Thus, winning the lottery results in a more accurate memory, as expected.

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


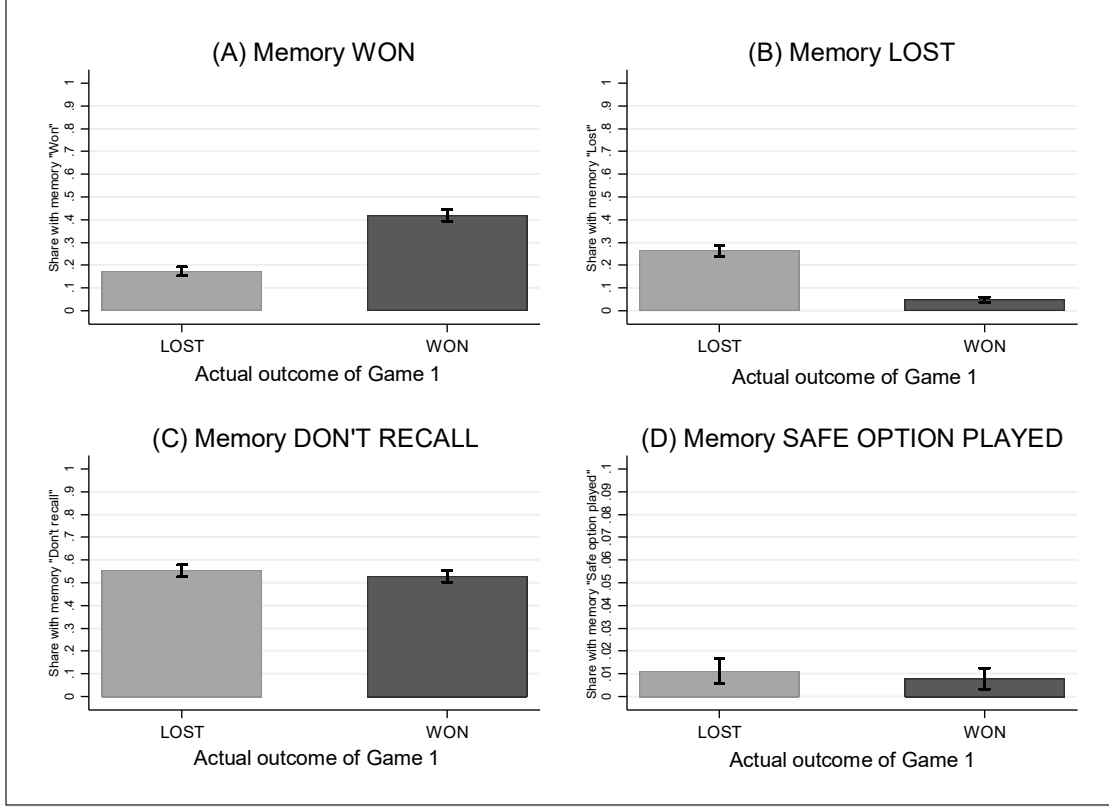


Figure 2: Memory of the outcome of Game 1, conditional on the actual outcome of Game 1 on the horizontal axis. Error bars indicate mean \pm standard error.



In Figures A1 and A2 in Appendix A, we show that the general difference in memory for good and bad outcomes is present in each age group and for all stake sizes of lotteries played in Game 1. The figures also show that older children are better able to recall the outcome if they had won ($p\text{-value} < 0.05$, Cuzick's Wilcoxon-type test for trend) and that the stake size increases the accuracy of memory for those who won ($p\text{-value} < 0.05$, Cuzick's Wilcoxon-type test for trend).¹⁸

These results provide support for our first conjecture, but we also wish to test the notion that, even when they misremember, those who lost are more likely to change their memory to a good one. In Figure 2, we show on the horizontal axis whether someone actually lost or won

¹⁸ We also ran a probit regression analysis in which we estimate the probability of remembering the outcome of Game 1 correctly and control in various models for the outcome, the number of risky choices and the stake size in Game 1 as well as an interaction term of the outcome and the stake size. Moreover, we add age, and interaction of age and the outcome of Game 1, gender, number of siblings, German school, relative IQ as well as average number of risky choices and fraction of winners in Game 1 in class to account for possible peer effects. Table A.3 in Appendix A reports the results of this analysis and confirms that correct recall is indeed significantly more likely if Game 1 was won and this effect can be seen for all age cohorts as well as stake sizes except for the lowest stake size. Moreover, increasing stake sizes as well as increasing age have a positive impact on accuracy of memory for those who won.

the lottery, and on the vertical axis the fraction of subjects with a particular memory. The figure shows that out of those who actually lost Game 1, 17.22% thought they had won (panel A, left bar). On the other hand, out of those who won Game 1, only 4.7% thought they had lost (panel B, right bar) ($p\text{-value} < 0.01$, $\chi^2\text{-Test}$). In addition, panels C and D show that there is no systematic relationship between not recalling the outcome or believing having chosen the safe option, and actually winning or losing the lottery ($p\text{-value} > 0.1$, $\chi^2\text{-Tests}$). Taken together, this leads to our first result.

Result 1: Those who won the first lottery are able to form a more accurate memory of the outcome than those who lost. The share of those who lost and believe they won is larger than the share of those who won and believe they lost.


4.3 Determinants of choices in Game 2

Having established that memories may differ from the actual experience, we now examine if either affects choices. In Figure 3, we display the average level of risk investment in Game 2 on the vertical axis. On the horizontal axis, we distinguish subjects by their level of risk tolerance in Game 1, with the most risk-averse subjects on the left hand side (“1 risky choice in Game 1”) and the most risk-tolerant subjects on the right hand side (“3 risky choices in Game 1”). Moreover, we split up each group contingent on whether a subject experienced a win or a loss in Game 1.

In line with conjecture 2, those who won the lottery in Game 1 invested, on average, 2.56 tokens in Game 2, but those who lost invested only 2.42 tokens. In line with our conjecture 3, we see in Figure 3 that this effect is driven by the two groups that are most risk averse ($p\text{-value} < 0.1$ and $p\text{-value} < 0.05$ for those who chose one or two risky choices in Game 1 respectively; $p\text{-value} < 0.01$ if we pool both subsets).¹⁹ There is no significant impact of the outcome in Game 1 on average risk investments in Game 2 for the most risk-tolerant group on the right-hand side of the figure. Overall, the pattern in Figure 3 is in line with our conjecture 3.

¹⁹ Unless noted otherwise, the p -values are the result of two-sided Mann-Whitney U-tests.


Figure 3: Average number of risk investment in Game 2, conditional on the outcome of Game 1 and conditional on the number of risky choices made in Game 1. Error bars indicate mean \pm standard error.



Turning to memory, Figure 4 highlights the level of risk investments made in Game 2 as a function of memory about Game 1 outcomes. We classify the memory as “Memory Won”, “Memory Lost” and “Memory Other”, where “Memory Other” is a composite of “Memory Don’t Recall” and “Memory Safe Option Played”.²⁰ Again, we display results for each of the three subgroups based on the level of risk tolerance in Game 1. Figure 4 shows that memory does not have as much influence on choices as the actual outcome. In fact, we never find a statistically significant difference between “Memory Won” and “Memory Lost” on the level of risk investments in Game 2 ($p\text{-value} > 0.1$ in all three comparisons). In sum, our non-parametric analyses support conjectures 1, 2, and 3.

²⁰ We pool the two groups because there are no distinct effects between these two groups on risk investments and in addition there are only 7 subjects who think they had not played the lottery but instead had chosen the safe option (when in fact they had played the lottery).

Figure 4: Average number of risk investment in Game 2, conditional on the beliefs of the outcome of Game 1 and conditional on the number of risky choices made in Game 1. Error bars indicate mean \pm standard error.



Next, we turn to our regression analysis.²¹ Our dependent variable is the level of risk investments made in Game 2 (ranging from 0 to 5). The primary independent variables are the following: *Won Game 1* (=1 if they won in Game 1, 0 otherwise), *# of risky choices Game 1* (= the number of risky choices made in Game 1, ranging from 1 to 3), an interaction of these two variables to account for differences observed in Figure 3, *Memory WON* (=1 if they believed they won, 0 otherwise), and *Memory OTHER* (=1 if they believed they had chosen the safe option or if they couldn't recall the outcome, 0 otherwise). This means that the omitted memory-category is "Memory LOST". We also control for *Stake size Game 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*²², *No. of siblings* (number of siblings), *German school* (=1 if a subject attended 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

²¹ We employ OLS regressions throughout. Note that the results remain qualitatively unchanged when we apply ordered probit models instead (see Table A.4 in Appendix A).

²² Although our research question does not consider gender differences, the regression results in Table 2 reveal a significant gender effect in risk investments in Game 2. In line with prior literature (Croson and Gneezy, 2009), female subjects are willing to invest significantly fewer tokens into the second lottery.

fraction of winners within a subject's class to account for potential peer effects. In column (1) of Table 2, 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.

Table 2: OLS regression with the number of risk investment in Game 2 as the dependent variable

VARIABLES	(1)	(2)	(3)	(4)
Won Game 1 (=1)	0.795*** (0.270)		0.781*** (0.275)	0.890*** (0.282)
# risky choices Game 1	0.314*** (0.077)	0.161** (0.065)	0.315*** (0.078)	0.324*** (0.078)
Won \times # risky choices Game 1	-0.289** (0.118)		-0.287** (0.118)	-0.337*** (0.121)
Memory WON (=1)		0.121 (0.126)	0.025 (0.134)	-0.026 (0.137)
Memory OTHER (=1)		-0.049 (0.117)	-0.109 (0.117)	-0.118 (0.122)
Stake size Game 1				-0.018 (0.054)
Age				-0.016 (0.038)
Female (=1)				-0.215** (0.091)
No. of siblings				-0.122 (0.085)
German school (=1)				-0.030 (0.048)
Relative IQ				0.035 (0.255)
Average # risky choices Game 1 in class				0.068 (0.168)
Fraction of winners of Game 1 in class				0.034 (0.285)
Constant	1.706*** (0.182)	2.469*** (0.099)	1.762*** (0.195)	1.959*** (0.629)
Observations	743	743	743	718
R-squared	0.020	0.004	0.023	0.036
<i>Wald tests (p-values)</i>				
H_0 : No effect of winning Game 1 for				
... those who had 1 risky choice in Game 1	0.002		0.004	0.002
... those who had 2 risky choices in Game 1	0.010		0.021	0.027
... those who had 3 risky choices in Game 1	0.565		0.524	0.361

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.

There are several noteworthy patterns in Table 2. First, if subjects experienced a win in Game 1, they made significantly larger risk investments in Game 2 (see positive and statistically significant coefficients on *Won Game 1* in columns (1), (3), and (4)). Second, this increase is attenuated for those who are less risk averse, i.e., more risk tolerant, as seen by the negative and statistically significant sign on the interaction, *Won \times #risky choices Game 1*, in the same three columns. The Wald tests in the last three rows of the table confirm that the actual experience of winning or losing is only significant for subjects who had played 1 or 2 risky choices, but not for the most risk-tolerant subjects who had picked 3 (out of 3) risky choices in Game 1. For example, subjects who chose the risky lottery only once in Game 1, and who won in this lottery, are estimated to increase the investment in the second lottery by 0.55 tokens, compared to subjects who lost in that lottery (controlling for additional background variables as in column (4)).²³ Given an average investment of 2.49 tokens in the second lottery for the entire sample, an estimated difference of 0.55 tokens, depending upon having won or lost the first lottery, represents an economically large effect of more than 20% of the investments in the second lottery. Hence, winning the lottery matters, in particular for the most risk-averse subjects who increase their risk-taking most strongly. From the regressions, we also see another important result, namely that the effects of memory are never found to be significant – see the coefficient on *Memory WON* in columns (2), (3), and (4). We can summarize the evidence from Figures 1 and 2 and from Table 2 in our second result.

Result 2: Winning a lottery ten months prior leads to significantly more risk taking in a subsequent lottery. This effect is driven by the relatively more risk-averse subjects. The memory of winning a small-stakes lottery ten months prior does not affect subsequent risk taking.

4.4 Replication of the effect of the lottery outcome on later risk taking in a second experiment

The fact that winning or losing a small-stakes lottery affects risk taking ten months later may be surprising to some. Acknowledging the importance of replicating experimental findings to mitigate concerns of false-positive results (e.g., Maniadis *et al.*, 2014), we next show that the result replicates in another experimental dataset. In this second experiment, Austrian school

²³ Our regression analysis is insensitive to several robustness checks reported in Appendix A. First, in Table A.5 we additionally control for subjects' patience, cooperation behavior and altruism. Second, in Table A.6 we drop those 243 subjects with inconsistent choice patterns in Game 1 (i.e., choosing the save option in tasks where the potential lottery win is higher and choosing to play the lottery when it is lower). Third, we drop 34 subjects with comprehension problems in Table A.7.

students completed two identical choice-list tasks which were separated by more than 3 years (in fact, 37 months) from one another. Both tasks included 20 binary risk choices (instead of three binary choices in the first lottery, and a risky investment task in the second lottery of our main experiment). In each choice, subjects had to choose between a lottery (which always yielded either zero or 10 Euro with equal probability) and a safe amount of money (which increased in steps of 50 cents from 0.50 Euro to 10 Euro). For payment, one of the 20 choices was randomly picked, and if the lottery had been chosen in the respective choice, a random draw (from an urn) determined whether a subject won or lost the lottery. Again, subjects received feedback about the lottery outcome directly after the first risk-taking task. The results from this first lottery task are reported in Sutter *et al.* (2013). The second task was run more than 3 years later, and has not been reported elsewhere previously. We utilize the choices of 136 school students who had chosen the lottery in the choice that was selected for payment in the first task and who made a decision in the second task more than three years later. In the first task, the children were in fifth, seventh, and ninth grade, while in the second task the same children were in eighth, tenth, and twelfth grade.

Out of the 136 subjects, 65 won the lottery in the first task, and 71 lost it. We can calculate for them the certainty equivalent which indicates the safe amount of money that makes them indifferent between earning this safe amount and playing the lottery.²⁴ Larger certainty equivalents indicate more risk taking. A simple pairwise comparison indicates that those subjects who won in the first task take more risks in the second task (two-sided t-test, $p\text{-value} = 0.025$). We can confirm this effect in a simple linear regression that is shown in Table 3. We take the certainty equivalent of the lottery in task 2 (3 years after task 1) as the dependent variable. In addition to further controls (see the notes to Table 3), we examine the effects of having won the lottery in the first task and control for the certainty equivalent in the first task. We see that having won the lottery 3 years prior increases the certainty equivalent by about 0.6 Euro to 0.7 Euro, which represents 12% to 14% of the lottery's expected value of 5 Euro. While we did not ask for memories in this second experiment, we can confirm the main result from the first experiment, i.e., that the random determination of a lottery outcome has a causal effect on the extent of risk taking much later, in this case more than 3 years later. Note that the confirmation in the second experiment is based on slightly older children (11 to 15 years in the first task, and 14 to 18 years in the second task), originates from a different country (Austria instead of Italy), comprises a longer time span between both tasks (37 months versus 10

²⁴ We do this by taking the midpoint between the two sure payoffs where the subject switched from the lottery to the sure payoff.

months), and has used a different risk elicitation method. Despite those differences in methods and subject pool, we can confirm our previously established Result 2: Subjects who have won the first lottery (more than three years earlier in this experiment) are taking more risk in the second lottery.

Table 3: OLS regression with the certainty equivalent in the second lottery task as the dependent variable

VARIABLES	(1)	(2)	(3)
First lottery won (=1)	0.705** (0.311)	0.602** (0.300)	0.608** (0.302)
Certainty equivalent in first lottery task (3 years earlier)		0.166* (0.085)	0.162* (0.085)
Further controls	No	No	Yes
Constant	4.722*** (0.214)	3.844*** (0.498)	3.679*** (1.210)
Observations	136	136	136
R-squared	0.037	0.070	0.071

Notes: Robust standard errors in parentheses. ***Significance at 1% level. **Significance at 5% level. *Significance at 10% level. The dependent variable is the certainty equivalent measured in the second risk-taking task (37 months after the first task). A higher certainty equivalent represents a higher willingness to take risks. The dummy variable “First lottery won” indicates whether the outcome of the first lottery was a win (=1) or a loss (=0). “Certainty equivalent in first lottery task” measures the willingness to take risks in the first lottery task (i.e., the certainty equivalent). We also control for further variables: For some participants, we had varied the stake size of the lottery (but normalize everything here to a lottery with a 50:50 chance of earning 10€ or nothing). For most of our participants (111 out of 136) the prize was 10€. For the rest, we introduced a prize variation with age, increasing the prize from 6€ for the youngest age cohort to 8€ for the middle and 10€ for the oldest age cohort. The stake size does not have any significant influence. Further we control whether the probabilities of winning were known to be 50:50 (N=72) or were so only in expectation (as in Ellsberg’s, 1962, classical experiment) (N=64). We find no differences in certainty equivalents between both variants.

5 Conclusion

Humans have a remarkable ability to learn and thus adapt their behavior to frequently changing environments. Our environment thus shapes economic preferences and, more generally speaking, cognitive and non-cognitive skills. Naturally, the family and social environment have an enormous influence on preferences and economic choices (Almås *et al.*, 2016; Alan *et al.*, 2017; Alan and Ertac, 2018; Brenoe and Epper, 2018; Castillo *et al.*, 2020; Falk *et al.*, 2021), as do large macroeconomic shocks, like economic depressions, the last global financial crisis, or the current COVID-19 pandemic (Malmendier and Nagel, 2011; Dohmen *et al.*, 2016; Terrier *et al.*, 2021).

In this paper, we have been the first to study whether seemingly negligible outcomes from the past can also 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 reasons (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 data from our main experiment with 743 subjects, we have investigated if risky choices depend upon the randomized outcome of another lottery played ten months earlier. In line with the directional predictions of reinforcement learning, we have found that subjects who won the first lottery take significantly more risks in the second one. Since the outcome of the first lottery 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 remember – 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. 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 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. To the best of our knowledge, we are also the first to document selective memories of children, i.e., from a

time in life where human behavior is strongly shaped by one's family and environment and the associated experiences (Heckman, 2006, 2008; Cappelen *et al.*, 2020; Kosse *et al.*, 2020).

Most importantly, we have shown that the outcomes of small-scale events have a causal and relatively long-lasting effect on risk taking later on. The results from our main experiment are replicable in a second experiment in which more than 3 years passed between the first lottery task – and its resolution – and risk taking in a second lottery task. Again, the random determination of winning or losing the first task had a powerful effect on subsequent risk taking. Since risk preferences are so important for lifetime outcomes (Castillo *et al.*, 2018; Charness *et al.*, 2020; Schneider and Sutter, 2020; Samek *et al.*, 2021), our findings imply that a grain of randomness in small-scale, seemingly negligible, events can have important long-lasting consequences.

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.
- Alan, S. and S. Ertac (2018). Fostering Patience in the Classroom: Results from Randomized Educational Intervention. *Journal of Political Economy* **126(5)**, 1865-1911.
- 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.
- Almås, I., A. W. Cappelen, K. G. Salvanes, E. Ø. Sørensen and B. Tungodden (2016). Willingness to Compete: Family Matters. *Management Science* **62(8)**, 2149-2162.
- Arrow, K. J. (1965). *Aspects of the Theory of Risk-Bearing* Helsinki: Yrjö Jahnssonin Säätiö.
- 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'Économie 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.
- Brenoe, A. A. and T. Epper (2018). The Intergenerational Transmission of Time Preferences Persists across Four Decades. *Working Paper. University of Zurich*.
- 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 (2020). Detecting Drivers of Behavior at an Early Age: Evidence from a Longitudinal Field Experiment. NBER Working Paper 28288.
- 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., T. Garcia, T. Offerman, and M. C. Villeval (2020). Do Measures of Risk Attitude in the Laboratory Predict Behavior under Risk in and outside of the Laboratory? *Journal of Risk and Uncertainty*, **60**(2), 99-123.
- 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.
- Clifasefi, S. L., D. M. Bernstein, A. Mantonakis and E. F. Loftus (2013). "Queasy does it": False Alcohol Beliefs and Memories May Lead to Diminished Alcohol Preferences. *Acta Psychologica* **143**(1), 14-19.
- Cooper, D. J. and J. H. Kagel (2003). Lessons Learned: Generalizing Learning across Games. *American Economic Review* **93**(2), 202-207.
- Cooper, D. J. and J. H. Kagel (2005). Are Two Heads Better Than One? Team Versus Individual Play in Signaling Games. *American Economic Review* **95**(3), 477-509.
- Corkin, S. (2002). What's New with the Amnesic Patient H.M.?. *Nature Reviews Neuroscience* **3**(2), 153-160.

- Croson, R. and U. Gneezy (2009). Gender Differences in Preferences. *Journal of Economic Literature* **47**(2), 448-474.
- 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.
- Ellsberg, D. (1961). Risk, Ambiguity and the Savage Axioms. *Quarterly Journal of Economics*, **75**(4): 643-669.
- Enke, B., F. Schwerter and F. Zimmermann (2020). Associative Memory and Belief Formation. Working Paper. Harvard University.
- Epstein, S. (1994). Integration of the Cognitive and the Psychodynamic Unconscious. *American Psychologist* **49**, 709-724.
- Erev, I., E. Ert and A. E. Roth (2010). A Choice Prediction Competition for Market Entry Games: An Introduction. *Games* **1**, 117-136.
- 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 (2019). Investor Memory. Maastricht University, Working Paper.
- 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 (2019). Persistent Overconfidence and Biased Memory: Evidence from Managers. Working Paper. University of Pittsburgh.
- 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.
- Köszegi, B. (2006). Ego Utility, Overconfidence, and Task Choice. *Journal of the European Economic Association* **4(4)**, 673-707.
- 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.
- Levitt, S. D. (2021). Heads or Tails: The Impact of a Coin Toss on Major Life Decisions and Subsequent Happiness. *Review of Economic Studies* **88(1)**, 378-405.
- List, J. A., R. Petrie, and A. Samek (2021). How Experiments with Children Can Inform Economics. University of Chicago, Friedman Becker Institute Working Paper 2021-62.
- Loewenstein, G. (1987). Anticipation and the Valuation of Delayed Consumption. *Economic Journal* **97(387)**, 666-684.
- Loomes, G. and R. Sugden (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *Economic Journal* **92(368)**, 805-824.
- Malmendier, U. and S. Nagel (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?. *Quarterly Journal of Economics* **126(1)**, 373-416.,
- Maniadis, Z., F. Tufano and J. A. List (2014). One Swallow Doesn't Make a Summer: New Evidence on Anchoring Effects. *American Economic Review* **104(1)**, 277-290.
- Necker, S. and M. Ziegelmeyer (2016). Household Risk Taking after the Financial Crisis. *Quarterly Review of Economics and Finance* **59(C)**, 141-160.

- 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.
- Samek, A., A. Gray, A. Datar and N. Nicosia (2021). Adolescent Time and Risk Preferences: Measurement, Determinants and Field Consequences. *Journal of Economic Behavior and Organization* **184**, 460-488.
- Saucet, C. and M. C. Villeval (2019). Motivated Memory in Dictator Games. *Games and Economic Behavior* **117**, 250-275.
- Schneider, S. and M. Sutter (2020). Higher Order Risk Preferences: New Experimental Measures, Determinants and Related Field Behavior. IZA Working Paper 13646.
- Schwardmann, P. and J. van der Wee (2019). Deception and Self-Deception. *Nature Human Behavior* **3**, 1055-1061.
- Sutter, M., M. G. Kocher, D. Glätzle-Rützler 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.
- Terrier, C., Chen, D., Sutter, M. (2021). COVID-19 within Families Amplifies the Prosociality Gap between Adolescents of High and Low Socioeconomic Status. Working Paper. Max Planck Institute for Research on Collective Goods.
- Thorndike, E. L. (1898). Animal Intelligence: An Experimental Study of the Associative Processes in Animals. *Psychological Review Monograph Supplements* **2**(4).
- Weber, R. A. (2003). Learning' with No Feedback in a Competitive Guessing Game. *Games and Economic Behavior* **44**(1), 134-144.
- 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

Appendix A

Additional figures and tables

Figure A1: Share of subjects remembering the outcome correctly conditional on the outcome of Game 1 and conditional on age groups. Error bars indicate mean \pm standard error.





Figure A2: Share of subjects remembering the outcome correctly conditional on the outcome of Game 1 and conditional on the stake size of the lottery played in Game 1. Error bars indicate mean \pm standard error.



Notes:

In the sample that we analyze, 215 played the “Low”, 214 the “Medium” and 314 the “High” stake size lottery in Game 1. Note that there are relatively more subjects playing the “High” stake size lottery as relatively more risk-averse subjects play only the lottery with the highest stakes (if at all) which increases the number of observations for the “High” stakes lottery.

Table A.1: Background characteristics of the 26 subjects who participate only in Game 1 compared to the 1,122 subjects who participated in both games

Background characteristics	Both games	<i>N</i>	Only game 1	<i>N</i>	<i>p</i> -value ¹
Age in years	7.60 [1.18]	1,122	7.34 [1.49]	26	0.216
Female (=1)	45.37%	1,122	50%	26	0.639
German school (=1)	45.99%	1,122	42.31%	26	0.710
No. of siblings	1.305 [1.038]	1,121	1.231 [1.070]	26	0.697
No. of risky choices ^a	1.907 [0.855]	1,122	1.577 [1.027]	26	0.138
Income father ^c	1,815 [431]	965	1,630 [188]	4	0.404
Income mother ^c	1,653 [392]	813	1,600 [528]	3	0.680
Father self-employed (=1) ^c	16.13%	967	0%	4	0.381
Mother self-employed (=1) ^c	6.61%	817	33.33%	3	0.065

Notes:

Mean with standard deviation in brackets for non-dummy variables.

¹ *p*-values from χ^2 -Tests for dummy variables and Mann-Whitney U-tests for non-dummy variables

Table A.2: Descriptive statistics and balancing tests of subjects who played the lottery (N=743) and subjects who did not play the lottery (N=379)

Background characteristics	Played lottery	<i>N</i>	No lottery	<i>N</i>	<i>p</i> -value
Age in years	8.44 [1.20]	743	8.50 [1.18]	379	0.377
Female (=1)	45.76 %	743	44.59 %	379	0.710
German school (=1)	44.95 %	743	47.76 %	379	0.373
No. of siblings	1.37 [1.00]	721	1.33 [1.11]	369	0.213
No. of risky choices Game 1 ^a	2.27 [0.69]	743	1.20 [0.70]	379	0.000
No. of risk investment Game 2 ^b	2.49 [1.17]	743	2.31 [1.26]	379	0.012
IQ ^c	20.42 [4.47]	737	20.10 [4.93]	372	0.621
Income father ^d	1,825 [432]	645	1,795 [430]	320	0.407
Income mother ^d	1,659 [397]	536	1,641 [383]	277	0.697
Father self-employed (=1) ^d	17.00 %	647	14.37 %	320	0.296
Mother self-employed (=1) ^d	7.05 %	539	5.76 %	278	0.480

Notes:

Column 2 (4): sample means for winners (losers); standard deviations in brackets (for non-dummy variables); Column 3 (5): number of observations winners (losers); Column 6: *p*-values from two-sample χ^2 -Tests (for dummy variables) and Mann-Whitney *U* Tests (for non-dummy variables).

The numbers of observations sometimes differ between rows because of missing values.

^a Number of risky choices in the three binary decision tasks of risk Game 1.

^b Number of risk investment in the lottery of risk Game 2.

^c The IQ was measured with a modified version of Raven's Colored Progressive Matrices

^d For a detailed description of these variables see the notes in Appendix C. We did not get information about parents' professions for all children.

Table A.3: Probit regression with the probability of remembering the outcome of Game 1 correctly as the dependent variable

VARIABLES	(1)	(2)	(3)
Won Game 1 (=1)	0.154*** (0.038)	0.151*** (0.038)	0.125*** (0.045)
# risky choices Game 1		-0.033 (0.026)	-0.043 (0.029)
Stake size Game 1		0.072*** (0.021)	0.075*** (0.022)
Age			0.028* (0.015)
Female (=1)			-0.057 (0.037)
No. of siblings			0.020 (0.019)
German school (=1)			-0.102*** (0.035)
Relative IQ			0.132 (0.089)
Average # risky choices Game 1 in class			0.003 (0.075)
Fraction of winners of Game 1 in class			0.389*** (0.135)
Observations	743	743	718
Pseudo R-squared	0.021	0.036	0.063

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.

Table A.4: Ordered probit regression with the number of risk investment in Game 2 as dependent variable

VARIABLES	(1)	(2)	(3)	(4)
Won Game 1 (=1)	0.721*** (0.241)		0.713*** (0.246)	0.822*** (0.256)
# risky choices Game 1	0.277*** (0.070)	0.134** (0.059)	0.279*** (0.070)	0.300*** (0.071)
Won × # risky choices Game 1	-0.270** (0.105)		-0.269** (0.105)	-0.319*** (0.109)
Memory WON (=1)		0.094 (0.111)	0.015 (0.119)	-0.032 (0.123)
Memory OTHER (=1)		-0.059 (0.103)	-0.109 (0.103)	-0.119 (0.109)
Stake size Game 1				-0.015 (0.049)
Age				-0.005 (0.035)
Female (=1)				-0.193** (0.082)
No. of siblings				-0.033 (0.044)
German school (=1)				-0.112 (0.077)
Relative IQ				0.040 (0.232)
Average # risky choices Game 1 in class				0.043 (0.155)
Fraction of winners of Game 1 in class				0.043 (0.260)
Constant cut1	-1.407*** (0.178)	-1.787*** (0.158)	-1.468*** (0.186)	-1.595*** (0.572)
Constant cut2	-0.127 (0.164)	-0.508*** (0.153)	-0.183 (0.174)	-0.305 (0.572)
Constant cut3	0.762*** (0.170)	0.377** (0.156)	0.707*** (0.181)	0.604 (0.570)
Constant cut4	1.561*** (0.179)	1.173*** (0.164)	1.507*** (0.191)	1.426** (0.571)
Constant cut5	2.264*** (0.180)	1.874*** (0.178)	2.212*** (0.193)	2.144*** (0.588)
Observations	743	743	743	718
Pseudo R-squared	0.006	0.004	0.007	0.011
<i>Wald tests (p-values)</i>				
H_0 : No effect of winning Game 1 for				
... those who had 1 risky choice in Game 1	0.002		0.003	0.001
... those who had 2 risky choices in Game 1	0.016		0.030	0.039
... those who had 3 risky choices in Game 1	0.435		0.404	0.265

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.

Table A.5: OLS regression as in Table 2 model (4) with additional experimental measures as controls (number of risk investment in Game 2 as dependent variable)

VARIABLES	(4)
Won Game 1 (=1)	1.013*** (0.294)
# risky choices Game 1	0.365*** (0.078)
Won \times # risky choices Game 1	-0.411*** (0.125)
Memory WON (=1)	-0.043 (0.143)
Memory OTHER (=1)	-0.130 (0.124)
Stake size Game 1	-0.020 (0.058)
Age	-0.090** (0.040)
Female (=1)	-0.292*** (0.100)
No. of siblings	-0.045 (0.053)
German school (=1)	-0.238*** (0.082)
Relative IQ	0.028 (0.255)
Average # risky choices Game 1 in class	0.066 (0.165)
Fraction of winners of Game 1 in class	0.213 (0.288)
Cooperation (tokens sent to partner from own class, 0-5)	0.048 (0.044)
Patience (tokens invested in time investment task, 0-5)	0.100*** (0.029)
Donation (tokens donated to charity, 0-6)	0.068* (0.040)
Constant	1.770*** (0.516)
Observations	645
R-squared	0.079
<i>Wald tests (p-values)</i>	
H_0 : No effect of winning Game 1 for	
... those who had 1 risky choice in Game 1	0.001
... those who had 2 risky choices in Game 1	0.063
... those who had 3 risky choices in Game 1	0.115

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.

Table A.6: OLS regression with the number of risk investment in Game 2 as dependent variable for subjects with consistent choice patterns in Game 1 (N=500)

VARIABLES	(1)	(2)	(3)	(4)
Won Game 1 (=1)	1.177*** (0.390)		1.076*** (0.394)	1.189*** (0.417)
# risky choices Game 1	0.381*** (0.110)	0.130 (0.081)	0.379*** (0.110)	0.407*** (0.126)
Won × # risky choices Game 1	-0.417*** (0.155)		-0.398** (0.154)	-0.438*** (0.163)
Memory WON (=1)		0.273* (0.158)	0.203 (0.162)	0.123 (0.163)
Memory OTHER (=1)		-0.006 (0.158)	-0.046 (0.154)	-0.085 (0.152)
Stake size Game 1				0.053 (0.081)
Age				-0.003 (0.045)
Female (=1)				-0.169 (0.108)
No. of siblings				-0.025 (0.060)
German school (=1)				-0.074 (0.105)
Relative IQ				0.018 (0.255)
Average # risky choices Game 1 in class				0.158 (0.235)
Fraction of winners of Game 1 in class				-0.133 (0.358)
Constant	1.520*** (0.284)	2.143*** (0.244)	1.512*** (0.313)	1.252 (0.864)
Observations	500	500	500	486
R-squared	0.020	0.016	0.029	0.038
<i>Wald tests (p-values)</i>				
H_0 : No effect of winning Game 1 for				
... those who had 1 risky choice in Game 1	0.002		0.007	0.005
... those who had 2 risky choices in Game 1	0.006		0.026	0.023
... those who had 3 risky choices in Game 1	0.576		0.375	0.379

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.

Table A.7: OLS regression with the number of risk investment in Game 2 as dependent variable for subjects who understood Game 1 (N=709)

VARIABLES	(1)	(2)	(3)	(4)
Won Game 1 (=1)	0.753*** (0.273)		0.753*** (0.281)	0.877*** (0.290)
# risky choices Game 1	0.325*** (0.0794)	0.195*** (0.0638)	0.327*** (0.0796)	0.344*** (0.080)
Won × # risky choices Game 1	-0.242** (0.117)		-0.243** (0.117)	-0.310** (0.122)
Memory WON (=1)		0.132 (0.123)	-0.00459 (0.131)	-0.040 (0.134)
Memory OTHER (=1)		-0.0241 (0.117)	-0.108 (0.115)	-0.109 (0.121)
Stake size Game 1				-0.002 (0.054)
Age				-0.021 (0.036)
Female (=1)				-0.218** (0.093)
No. of siblings				-0.005 (0.050)
German school (=1)				-0.121 (0.083)
Relative IQ				0.104 (0.253)
Average # risky choices Game 1 in class				0.074 (0.148)
Fraction of winners of Game 1 in class				0.090 (0.270)
Constant	1.662*** (0.186)	2.033*** (0.169)	1.718*** (0.198)	1.770*** (0.516)
Observations	709	709	709	686
R-squared	0.028	0.018	0.030	0.043
<i>Wald tests (p-values)</i>				
H_0 : No effect of winning Game 1 for				
... those who had 1 risky choice in Game 1	0.003		0.004	0.002
... those who had 2 risky choices in Game 1	0.002		0.004	0.011
... those who had 3 risky choices in Game 1	0.819		0.843	0.696

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.

Appendix B

Experimental instructions for main experiment (translated from German/Italian) – for the instructions for the second experiment (in section 4.4), see Sutter *et al.* (2013)

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

Risk preferences Game 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*).

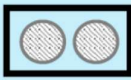

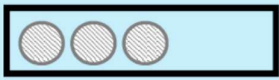

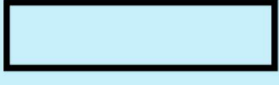
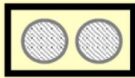

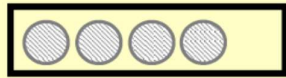

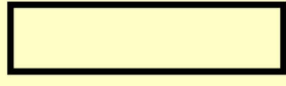
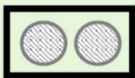

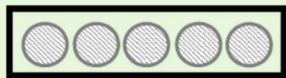

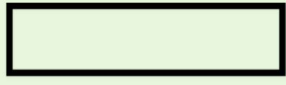
The green part is very similar to the blue and yellow part. Here you see the decision sheet for the green part. Again, there are 2 tokens on the left-hand side, but on the right-hand side there are 5 tokens now. What do you think will happen if you tick THIS box? (*please point at the*

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 B1: Decision Sheets “Risk Preferences Game 1”

 <input data-bbox="305 546 354 598" type="checkbox"/>	 =   =  <input data-bbox="568 546 617 598" type="checkbox"/>
 <input data-bbox="300 997 349 1050" type="checkbox"/>	 =   =  <input data-bbox="563 997 612 1050" type="checkbox"/>
 <input data-bbox="300 1438 349 1491" type="checkbox"/>	 =   =  <input data-bbox="563 1438 612 1491" type="checkbox"/>

Risk preferences Game 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 B2: Decision Sheet “Risk Preferences Game 2”



Appendix C

Notes on the estimated income of parents

In order to get a measure for income, we asked the children to state their parent’s profession as precisely as possible. The children’s answers were categorized with the use of the Public Employment Service Austria (AMS). They provide information on the average gross starting salary per month of almost 1,800 different types of professions. If a child could only give information on the company the parent works at, we used the most common profession within the same company. We used the Austrian Public Employment Service (AMS) classification because the information provided there, on different types of professions is much more detailed than the information provided by the census bureau in South Tyrol (ASTAT). However, the average gross starting salary provided by both the AMS and the ASTAT have a highly significant positive correlation. Note that we did not get information about parents’ professions for all children participating in our experiment.