Discussion Papers of the
Max Planck Institute for
Research on Collective Goods
2023/6

Income Risk, Precautionary Saving, and Loss Aversion – An Empirical Test

Marcela Ibanez
Sebastian O. Schneider

MAX PLANCK
SOCIETY

Discussion Papers of the
Max Planck Institute
for Research on Collective Goods 2023/6

Income Risk, Precautionary Saving, and Loss Aversion – An Empirical Test

Marcela Ibanez / Sebastian O. Schneider

June 2023

Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, D-53113 Bonn
https://www.coll.mpg.de

Income Risk, Precautionary Saving, and Loss Aversion – An Empirical Test*

Marcela Ibanez
marcela.ibanez@zentr.uni-goettingen.de
University of Goettingen, Germany

Sebastian O. Schneider
sschneider@coll.mpg.de
Max Planck Institute for Research on Collective Goods, Bonn, Germany

May 20, 2023

Abstract

This paper empirically examines the behavioral precautionary saving hypothesis that uncertainty about future income triggers an increase in saving because of loss aversion. Guided by the theoretical model of Kőszegi and Rabin (2009), we first extend their theoretical analysis to also consider the internal margin, i.e., the strength, of loss aversion, and then empirically study the relation between income risk, experimentally elicited loss aversion, and precautionary savings. We do so using a sample of 640 individuals from the low-income population of Bogotá, characterized by limited financial education and subject to substantial income risk. In line with the theoretical predictions, we find that an increase in income risk is associated with higher savings for loss-averse individuals, and that this increase in savings grows with the degree of loss aversion. An accompanying laboratory experiment confirms that an exogenous increase in income risk causally leads to this observed pattern. Thus, consistent with the theoretical predictions derived from the model of Kőszegi and Rabin (2009), but in contrast to common assumptions, our findings establish that loss aversion is not necessarily an obstacle to saving, and thus identify new approaches of increasing saving among individuals with low financial education.

JEL-classification: D11, D14, D15, D81, D90, G40, J65, O16

Keywords: Reference-dependent utility, expectations, consumption plans, precautionary savings, loss aversion, risk preferences, income risk, low income, Bogotá, experiment.

*We thank Zvonimir Bašić, Claudia Cerrone, Sanjit Dhami, Christoph Engel, Stephan Klasen, Katja Michaelowa, and Matthias Sutter, as well as participants at the Behavioral Finance Working Group Conference in London, the Symposium on Economic Experiments in Developing Countries in Kenya, and the European ESA meeting in Bergen for helpful comments and suggestions. We thank Margarita Cabra, Lukas Hermann, Aida Takhmazova, Sonia Triviño, and Jamshidbek Odiljonov for excellent research assistance. Sebastian Schneider gratefully acknowledges the financial support provided by the German Research Foundation (DFG) via RTG 1723 and the Foundation of German Businesses (sdw).

1. Introduction

Loss aversion has typically been associated with poor financial status and poor financial decision-making – both theoretically and empirically (e.g., Benartzi and Thaler, 1995; Odean, 1998; Thaler and Benartzi, 2004; Cardenas and Carpenter, 2013; Barberis, 2013). However, in reference-dependent models of inter-temporal consumption, loss aversion takes a diametrically opposed role: individuals who are loss-averse with respect to future consumption save more today when facing income risk, in order to decrease the utility loss associated with possible consumption levels below their reference point (Aizenman, 1998; Siegmann, 2002; Kőszegi and Rabin, 2009; Park, 2016; Pagel, 2017).[1] Yet, there is no empirical test of this loss aversion-based precautionary motive for savings on the individual level.

In this paper, we extend the analysis of the model by Kőszegi and Rabin (2009), derive and investigate three specific hypotheses: i) loss-averse individuals save more when exposed to greater income risk (precautionary saving motive); ii) holding (positive) income risk constant, a higher degree of loss aversion implies more saving compared to lower degrees of loss aversion; and iii) the more loss-averse individuals are, the higher the increase in saving associated with an increase in income risk. Using real-world data on savings and income risk and experimentally elicited, incentivized measures of loss aversion, we examine the relationship between income uncertainty, loss aversion, and savings empirically. In an accompanying laboratory experiment on saving, we alter income uncertainty exogenously to address causality. This is the first paper that directly addresses the proposed theoretical relation empirically, and the first to investigate the relationship between household saving and loss aversion using a direct measure of loss aversion at the individual level.

The topic is important not only from a decision-theoretic perspective, but also from a policy point of view. Saving rates are notoriously low in developing and developed nations alike (e.g., World Bank, 2014), resulting in a widespread inability to cope with income shocks. For example, in the U.S., nearly 40% of the population cannot easily come up with $400 to cover unexpected expenses, and 25% of adults skipped medical treatments due to an inability to pay (Federal Reserve Board, 2020). Income shocks, and consequently the unpreparedness to cope with them,

[1] Following Kőszegi and Rabin (2009), we use the terms "income uncertainty" and "income risk" as synonyms in this paper.

have serious negative impacts such as poor mental health and increased suicide rates (e.g., Clark et al., 2016; Christian et al., 2019). This calls for a an improvement in the savings situation – especially for those facing limitations in access to credit markets and weak social-protection systems. For them, savings are one of the few alternatives – if not the only one – to prepare for, and mitigate the effect of, income shocks (Dercon, 2010; Tovar and Urrutia, 2017). A prerequisite to improving this situation is understanding household saving behavior and its drivers. In this respect, the combined positive effect of income uncertainty and loss aversion has been completely neglected so far in empirical work.

We study the derived hypotheses focusing on the relation between income risk, loss aversion, and savings among the low-income population in Bogotá, Colombia. Although the financial means to save are certainly lower in this population group than in the middle-income class, we observe considerable heterogeneity in savings in our sample, implying that many individuals do engage in saving even though they are poor.[2] What is more important, however, is that this population group is exposed to substantial income uncertainty, similar to comparable population groups in other developing countries (Loayza et al., 2009; Stampini et al., 2016). The majority of the employed population in our sample depend on informal employment (52%) or are underemployed (30%), i.e., working less than full time. Yet, informal risk-sharing networks that are commonly observed in rural areas (e.g., D'Exelle and Verschoor, 2015) only play a minor role in poor urban settings. Hence, due to its exposure to various income shocks without (formal and informal) social-security systems in place, this population group is ideally suited to investigate the precautionary savings motive emerging from a reference-dependent model of consumption and saving. Moreover, since income shocks may pose a threat of falling into a poverty trap, they might be particularly consequential. Thus, the prospect of reducing future consumption can induce a severe loss in utility, thus generating a strong incentive to save.

Our investigation of a loss aversion-based precautionary saving motive is guided by the theoretical framework of Kőszegi and Rabin (2009). Others have discussed the positive effect of loss aversion on precautionary saving using reference-dependent models before (e.g., Aizenman, 1998; Siegmann, 2002). However, the Kőszegi and Rabin model arguably is the most general one

[2] Several studies from development settings report similar findings; see, e.g., the review in Karlan et al. (2014).

that is embedded in a rich theory of expectation-based preferences with numerous applications, which makes empirical investigations particularly useful.

We construct a measure of savings as the total value of liquid assets. To measure loss aversion and risk preferences, we use the experimental method by Abdellaoui et al. (2007). To measure income risk, we use subjective (ex-ante) measures of the probability of becoming unemployed to obtain a prediction of the regional unemployment risk (Lusardi, 1998; Lugilde et al., 2018; Benito, 2006; Guariglia and Kim, 2004). We use a perception-based measure, as individuals are expected to have better information on their employment prospect than an external observer (Deaton, 1991), which might apply particularly to the informal workforce. In addition, Weil (1993) shows that labor-income shocks are a crucial determinant of the precautionary saving motive. Lastly, from the individual perspective, the local average unemployment risk can arguably be considered exogenous.

We find, in line with a precautionary savings motive, that individuals hold higher savings when they are exposed to greater income risk, which is particularly true for the loss-averse sub-sample (Hypothesis 1). Similarly, and again consistent with the reference-dependent model of intertemporal consumption by Kőszegi and Rabin (2009), we find that, controlling for income risk, more loss-averse individuals hold higher savings than less loss-averse individuals (Hypothesis 2). Moreover, the larger the degree of loss aversion, the larger the value saved that is associated with an increase in income risk (Hypothesis 3). These findings are robust to different measures of loss aversion, the inclusion of controls on a large set of socioeconomic characteristics, and different econometric specifications.

We further extend this analysis using a two-stage laboratory experiment to causally identify the impact of increases in income risk. In the first stage, we measure loss aversion, following again Abdellaoui et al. (2007). In the second stage, we implement the saving experiment by Xu et al. (2022) to vary exogenously the degree of income risk. The empirical results support the predictions derived from the reference-dependent model by Kőszegi and Rabin (2009) of a positive effect of income risk on saving, particularly among loss-averse individuals (Hypothesis 1). As with the field data, this effect increases with the degree of loss aversion (Hypothesis 3).

We make several contributions to the economic literature. First, we contribute to the literature investigating household saving behavior in general. Our results provide support for the predictions emerging from reference-dependent precautionary saving models – in particular from the one by Kőszegi and Rabin (2009). We empirically establish that the role of loss aversion on saving is more complex than commonly proposed, and in particular, that loss aversion can even foster saving. For example, Thaler and Benartzi's (2004) 'Save More Tomorrow Program' builds on the argument that loss-averse individuals perceive saving a portion of their current income as a loss. However, Thaler and Benartzi (2004) do not measure loss aversion, nor do they account for uncertainty of future income. This is the first paper using individual level data to investigate empirically the predictions of reference-dependent precautionary saving models in general regarding the role of loss aversion and income risk on savings (e.g., Aizenman, 1998; Siegmann, 2002), and those derived from the model by Kőszegi and Rabin (2009) in particular.

Our second contribution is to the literature on precautionary savings (e.g., Leland, 1968; Sandmo, 1970; Dreze and Modigliani, 1975; Carroll et al., 2000, 2021). Bowman et al. (1999) were among the first to consider loss aversion in a consumption-saving model with uncertain income, and show that when there is sufficient income uncertainty, a loss-averse person resists lowering consumption in response to bad news about future income. This resistance is greater than the resistance to increasing consumption in response to good news, generating an asymmetry in consumption behavior. Support for this prediction was provided on the macro level by Fisher and Montalto (2011).[3] However, when income uncertainty is resolved in the second period as is typically the setting in precautionary saving frameworks, in the reference-dependent models, loss aversion induces a precautionary savings motive (Aizenman, 1998; Siegmann, 2002; Kőszegi and Rabin, 2009; Park, 2016; Pagel, 2017). The precautionary savings motive resulting from these reference-dependent models is arguably more intuitive than the traditional theory (e.g., Leland, 1968; Kimball, 1990), which assumes that individuals conduct a rather complex theoretical optimization of expected utility. For example in Kőszegi and Rabin (2009), individuals save to mitigate (over-weighted) utility losses due to consumption levels below their expectations that might occur if income falls.[4] For this reason, these models might (more) adequately capture the

[3] See also the model by Eeckhoudt et al. (2016).

[4] Other models use different reference points than the expectation, and incorporate loss aversion differently, but the intuition stays the same.

saving behavior of populations that fail to perform the most sophisticated financial planning.

While the classical theory has been empirically studied in the Western world (e.g., Guiso et al., 1992; Dynan, 1993; Lise, 2013; Bayer et al., 2019; Christelis et al., 2020),[5] we contribute to this literature by studying precautionary savings among low-income populations in developing countries (Jalan and Ravallion, 2001; Giles and Yoo, 2007; Paxton and Zhuo, 2011; Michler and Balagtas, 2017).

The role of loss aversion on precautionary savings has empirically been considered by Pagel (2017). The paper presents an expectations-based reference-dependent life-cycle consumption model that depends on the loss aversion-based precautionary savings motive considered here, and provides empirical support for it. We contribute to this literature by providing micro evidence based on experimentally-elicited preference measures.[6]

Lastly, we contribute to the broader economic literature that empirically relates loss aversion with real-world consequences, including reference-dependent expectations and consumption decisions (e.g., Ericson and Fuster, 2011, or Karle et al., 2015). Benartzi and Thaler (1995) show that loss aversion can explain the equity premium puzzle and affects participation in equity markets. Moreover, there is empirical evidence linking (expectation-based) loss aversion and labor markets (e.g., Camerer et al., 1997; Fehr and Goette, 2007; Farber, 2008; Abeler et al., 2011; Imas et al., 2016), sports (Pope and Schweitzer, 2011; Bartling et al., 2015; Allen et al., 2016; Markle et al., 2018), domestic violence (Card and Dahl, 2011), innovation (Rosokha and Younge, 2020), housing markets (Andersen et al., 2019), and student performance (Karle et al., 2022), among others. We contribute to this line of research by focusing on the relation between loss aversion and savings. To our knowledge, we are the first to investigate this relationship empirically with a *direct measure* of loss aversion at the individual level. Previous work *argued* that loss aversion is (negatively) related to saving (e.g., Thaler and Benartzi, 2004) or *attributed observed behavior* to loss aversion (e.g., Fisher and Montalto, 2011), but lacked an appropriate

[5]See also the reviews by Carroll (2001), Browning and Crossley (2001), Lugilde et al. (2019), and Baiardi et al. (2020).on the empirical literature on precautionary savings and dynamic optimization models that build on this idea.

[6]Even in the (older) empirical literature investigating the classical precautionary savings hypothesis (e.g., Leland, 1968; Sandmo, 1970), a direct measurement of preference parameters has been seldomly used. Notable exceptions are Noussair et al. (2014) and Xu et al. (2022), who elicit higher-order risk preferences, but do not consider loss aversion.

measure of loss aversion at the individual level to *measure* the relation. The parameter-free elicitation procedure due to Abdellaoui et al. (2007) that we implement in our study provides a state-of-the-art measure of loss aversion on a highly relevant population group.

Our paper may help calibrate models building on the examined precautionary savings motive. Moreover, it might also inform policy design when considering behavioral approaches to increase the savings rate, in particular among those with limited financial education and among the poor (for a comprehensive overview of the research on saving among the poor, see Karlan et al., 2014). Based on our findings, we suggest that interventions might focus on the loss associated with failing to save – given that saving is a viable option at all. This could be done, for example, in a similar vein to Karlan et al. (2016), who compare the effectiveness of reminders to save that are framed as a loss ("your dreams won't come true") to reminders that are framed as a gain ("your dreams will come true"). They find no significant effects of the frame on a household's saving rate. However, they do not consider income uncertainty. As our research indicates, focusing on unacceptably low levels of consumption or well-being in the future, resulting from income uncertainty, could help in designing effective interventions to increase savings.

This paper is structured as follows. The next section presents the model of inter-temporal consumption by Kőszegi and Rabin (2009) and extends its analysis, from which the hypotheses of the study are derived. Section 3 explains the empirical strategy used to test the predictions of this model, and Section 4 explains how the different measures were obtained. Results are presented in Section 5 for the Colombian field data. Section 6 presents the design and the results of the laboratory experiment. The approaches and findings of the paper are discussed in Section 7, and Section 8 concludes.

2. Theoretical Framework

The conceptual framework that we use in the analysis is based on the reference-dependent utility model of inter-temporal consumption by Kőszegi and Rabin (2009). First, we introduce this model and derive the precautionary motive for saving as presented in their paper. In a second step, we extend the analysis to derive hypotheses relating the strength of the precautionary savings motive with the degree of loss aversion.

The model considers a two-period consumption-saving decision problem where individuals face uncertainty regarding their future wealth. Here, we present the model for the case in which wealth W is a binary random variable and uncertainty is resolved in the second period. In Appendix A.1, we present the two-period model for a more general case where W is non-binary random wealth.

We assume that with equal probabilities, wealth takes two possible values: W_0+s and W_0-s, where W_0 is deterministic income and $s>0$ a scalar, reflecting income risk.[7] An individual has to divide wealth W between consumption c_t in two periods, $t=1,2$, maximizing the sum of instantaneous utility in the first period and the expected instantaneous utility in t=2,

$$U = u_1(c_1) + \mathbb{E}[u_2(c_2)],$$

subject to the budget constraint $c_1 + c_2 = W$.

In the first period, there is no uncertainty on income and instantaneous utility is given by

$$u_1(c_1) = m(c_1),$$

where m is the utility of consumption that is assumed to be three times differentiable, increasing and strictly concave.[8]

The expected instantaneous utility in the second period, $\mathbb{E}[u_2(c_2)]$, depends on the expected utility of consumption in that period, m, and on the so-called 'gain-loss utility'. Before the first period starts, it is assumed that agents choose their favorite credible consumption plan, which specifies possibly stochastic consumption levels for each period. This plan is called the Personal Preferred Equilibrium (PPE).[9] When uncertainty is resolved in the second period and consumption decisions are implemented, plans are updated and lead to new beliefs. Changes in beliefs induce a gain or a loss in utility through 'gain-loss utility' depending on whether new beliefs imply a higher or lower consumption level than previously believed. Following Kahneman

[7] Results generalize to non-binary random income in the more general model; see the corresponding Proposition 8 in Kőszegi and Rabin (2009), as well as Proposition 1 and its Corollary in this study, for more general results.

[8] We abstract from overconsumption and assume that a deviation in period 1 from the ex-ante optimal plan cannot increase the assessment of the overall utility in period 1; see Proposition 5 in Kőszegi and Rabin (2009).

[9] Details on this concept are given in Appendix A.2 or in Kőszegi and Rabin (2009).

and Tversky (1979), it is assumed that individuals weight utility losses different than utility gains by using a factor $\lambda > 0$ that captures the degree of loss aversion. For an individual who is loss averse, we have $\lambda > 1$, whereas for a gain-seeking individual we have $\lambda < 1$.

If income is high (i.e., if $W = W_0 + s$), which occurs with probability 1/2, there is a gain in utility from changes in beliefs, as the individual had planned a lower consumption level (c_2^-) with probability 1/2. This change is weighted by $\eta > 0$, which captures the weight attached to gain-loss utility. Conversely, if income is low, there is a loss in utility since the agent had planned a higher consumption level (c_2^+), again with probability 1/2; this change is weighted by $\eta > 0$ and $\lambda > 0$ to account for loss-averse ($\lambda > 1$) or gain-seeking ($\lambda < 1$) behavior.[10]

Summarizing, the expected instantaneous utility in the second period is given by

$$\begin{aligned}\mathbb{E}[u_2(c_2)] = &\frac{1}{2}\left(m(c_2^+) + \frac{1}{2}\eta\left(m(c_2^+) - m(c_2^-)\right)\right) \\ &+ \frac{1}{2}\left(m(c_2^-) - \frac{1}{2}\lambda\eta\left(m(c_2^+) - m(c_2^-)\right)\right),\end{aligned}$$

where m is the utility of consumption as defined above, $c_2^+ = W_0 - c_1 + s$ and $c_2^- = W_0 - c_1 - s$.

As shown by Kőszegi and Rabin (2009), for an interior solution, the optimal consumption path satisfies

$$m'(c_1) = \frac{1}{2}m'(c_2^+) + \frac{1}{2}m'(c_2^-) + \frac{1}{4}\eta(\lambda - 1)[m'(c_2^-) - m'(c_2^+)]. \tag{1}$$

To see whether increases in risk, s, increase $m'(c_1)$, i.e., decrease c_1 (since m is strictly concave), we apply a Taylor approximation of the right-hand side of (1) around $s = 0$ to obtain[11]

$$m'(c_1) \approx m'(c_2) + \frac{1}{2}m'''(c_2)s^2 + \frac{1}{2}\eta(\lambda - 1)(-m''(c_2))s. \tag{2}$$

From this derivation, we see that, for a loss-averse individual (i.e., when $\lambda > 1$), uncertainty causes an increase in savings as consumption decreases in period 1 when $m''' > 0$, but also

[10]The assumption of putting a higher weight on utility below the reference point, hence assuming loss aversion (i.e., $\lambda > 1$), is common in reference-dependent models. Kőszegi and Rabin (2009) call it the "clearly correct assumption", although empirical studies also document 'gain-seeking' behavior (e.g., Schmidt and Traub, 2002). Therefore, we only assume $\lambda > 0$ and allow for gain-seeking behavior. See Appendix A.1 for further details.

[11]See Equation (11) in Kőszegi and Rabin (2009).

when the last term dominates the second term in (2). The first condition corresponds to the classical theory of precautionary saving, as initiated by Leland (1968), where a positive third derivative of consumption utility causes the individual to save. Following the assumption of Kőszegi and Rabin (2009) that m is a global utility function, small risks in the model might still be substantial in "practical terms", and thus the last term dominates the second term in (2).[12]

Generalizing to the case where second-period income has more than just two realizations and where people might overconsume in the first period leads to the first hypothesis from this model.[13]

Hypothesis 1. *For loss-averse agents, higher levels of income uncertainty are associated with higher levels of savings.*

We now extend the analysis of Kőszegi and Rabin (2009) to consider how the degree of loss aversion affects the precautionary savings motive. From both (1) and (2), we see that savings increase in the degree of loss aversion.[14] This finding can be generalized to non-binary income risk and to overconsuming individuals, i.e., those increasing their consumption relative to the ex-ante optimal level in the first period: Suppose wealth is now equal to $W_0 + sy$, where y is a non-deterministic mean-zero lottery that is resolved in period 2. Overconsumption is linked to a parameter $\gamma \geq 0$: If it remains below a certain threshold, individuals overconsume in this more general framework (Appendix A.1).[15]

Proposition 1. *For any increasing, strictly concave, three times differentiable consumption utility function m, any $\eta > 0, \lambda > 0, \gamma \geq 0$, and s small and positive, the personal preferred equilibrium consumption rule satisfies $dc_1/d\lambda < 0$.*

[12] This assumption is needed, since technically, this is true only for small s; see their comment in Footnote 25 in Kőszegi and Rabin (2009).

[13] See Proposition 8 in Kőszegi and Rabin (2009).

[14] As in the case of uncertainty, it cannot technically be said for sure whether the second-order condition for a utility maximum is satisfied for any amount of risk (as (2) is a Taylor approximation around $s = 0$). However, for small amounts of risk, the condition holds. Yet, following the argument of Kőszegi and Rabin (2009) that small risks in the model can be substantial in practical terms, this restriction comes, just as in their analysis, without major consequences.

[15] Kőszegi and Rabin (2009) show that an individual increases consumption in the first period relative to the ex-ante optimal level if $\gamma < 1/\lambda$; see Appendix A.1 for further details.

The proof of Proposition 1 is in Appendix A.3. From this proposition, we derive the following hypothesis:

Hypothesis 2. *When facing income uncertainty, a higher degree of loss aversion is associated with higher savings. This also includes coefficients of loss aversion* $\lambda \leq 1$.

Irrespective of Hypotheses 1 and 2 being true (i.e., the degree of loss aversion or uncertainty for loss-averse individuals being positively related to savings), from (1) and (2) we see that the effect of loss aversion on savings increases in uncertainty and that the effect of uncertainty on savings increases in loss aversion. As expected, this result generalizes to non-binary income lotteries and holds independently of individuals overconsuming in the first period:

Corollary 1. *For any increasing, strictly concave, three times differentiable consumption utility function* m *and any* $\eta > 0$, $\lambda > 0$, $\gamma \geq 0$, *the personal preferred equilibrium consumption rule satisfies* $d^2c_1/(dsd\lambda)|_{s=0} < 0$.

From Corollary 1 and following the interpretation of small risks in the model by Kőszegi and Rabin (2009), we can derive the third hypothesis:

Hypothesis 3. *The increase in savings associated with an increase in the degree of loss aversion is an increasing function of income uncertainty. Equivalently, the (positive) relation between income uncertainty and savings is an increasing function of the degree of loss aversion. As in Hypothesis 2, this also includes coefficients of loss aversion* $\lambda \leq 1$.

3. Empirical Strategy

To test the hypotheses derived from Kőszegi and Rabin's (2009) model in Section 2 and to investigate the relationship between income risk, loss aversion on the individual level and savings, we use individual measures of savings, income risk, and loss aversion in our analysis. To test Hypothesis 1 on the positive association between income risk and savings for loss-averse individuals, we run the following regression – with and without restricting our sample to loss-averse individuals:[16]

$$\text{Savings}_i = \beta_1 s_i + \zeta X_i + \beta_0 + \varepsilon_i, \qquad \text{(Model 1)}$$

[16] Although agents are generally assumed to be loss averse (e.g., Kőszegi and Rabin, 2009), for a rigorous test of Hypothesis 1, the sample has to be restricted to the loss-averse sub-sample.

where Savings_i denotes individual accumulated liquid savings of individual i, s_i is the individual's income uncertainty, X_i is a vector of socioeconomic characteristics for individual i, and ε_i is the error term; β_1 and ζ are regression coefficients estimating the relation between savings and income uncertainty and socioeconomic characteristics, respectively, and β_0 is the intercept of this model. The data would support Hypothesis 1 if $\beta_1 > 0$ (and in a more rigorous sense, if this is the case when restricting the sample to loss-averse individuals).

To test the second hypothesis, postulating that the degree of loss aversion is associated with higher savings when facing income risk, we run the following regression:

$$\text{Savings}_i = \beta_1 s_i + \beta_2 \lambda_i + \zeta X_i + \beta_0 + \varepsilon_i, \qquad \text{(Model 2)}$$

where λ_i is the degree of loss aversion of individual i with corresponding regression coefficient β_2. Although the population group under study is highly exposed to income uncertainty, we control for the degree of income uncertainty s_i for a rigorous test. A positive β_2 would support Hypothesis 2.

Finally, we test Hypothesis 3, claiming that the relation between income uncertainty and savings is an increasing function of the degree of loss aversion, by estimating the following equation:

$$\text{Savings}_i = \beta_3 (s_i \times \lambda_i) + \beta_1 s_i + \beta_2 \lambda_i + \zeta X_i + \beta_0 + \varepsilon_i, \qquad \text{(Model 3)}$$

where β_3 is the regression coefficient of the interaction term of individual loss aversion λ_i and individual income uncertainty s_i. Hypothesis 3 is supported if $\beta_3 > 0$. Note that we center income risk measures around mean values and loss aversion around 1, and, to maintain consistency, do so in all models. Hence, β_1 is the main effect of uncertainty for a loss-neutral individual, while β_2 is the main effect of loss aversion estimated at a mean level of income uncertainty. The theoretical model does not provide definitive predictions on savings for loss-neutral individuals facing income risk (in contrast to gain-seeking and loss-averse individuals).

The next section presents the definitions of loss aversion, savings, and income risk, and the data used for these measures.

4. Field Data

The primary field data used in the study were collected between October and November 2013 as part of a project investigating the financial vulnerability of low-income households in Bogotá. The study comprised an extensive survey of the financial situation of the households and incentivized economic experiments on risk and time preferences.

For our primary data collection, we conducted a two-step sampling process. First, low-income neighborhoods were identified by assessing the proportion of people belonging to the two lowest socioeconomic strata. Neighborhoods with a larger proportion of low-income population, and which were assessed as safe for the team to visit, were eligible for the study. Participants for the study were then selected from a list of households in the area in 2010. The criterion for selecting participants was that they should be beneficiaries of the social health insurance, SISBEN. This condition would guarantee that the participants were from low socioeconomic strata.

In total, 640 participants completed the survey and the experiment. The survey lasted around 90 minutes. The experiment was completed at a different location a few days later and took about 20 minutes. Participants received a participation fee of 10,000 Colombian Pesos (COP; about USD 5.35 at that time). In addition, participants were paid with a 20% probability according to their choices. This way, participants earned on average about COP 15,000 or about USD 8 (min: COP 10,000, max: COP 40,000) for the two-hour task, which is sizable given the official daily minimum wage of COP 19,650 in 2013 (which is of course not necessarily paid in the informal sector).

4.1. Definition of Variables

Savings We measure savings as the total value of an individual's monetary assets. This includes total savings in checking accounts, certificates of deposit, mutual funds, savings in cash or in other currencies, the value deposited in savings plans (i.e., money to buy a house or to pay for the education of their children), and the net value of loans given. We use the sum of those categories, since in cases of emergency it is possible to withdraw money from all of these savings devices.

Income Risk In the analysis, income risk is measured by the (self-reported and aggregated) risk of becoming unemployed based on survey data.[17] In order to handle the endogeneity and imprecision associated with self-reported unemployment risk, we use the average probability at the local planning unit, *UPZ*, or, put differently, the predicted unemployment risk at the regional level where the UPZ is the single predictor. Typically, a *localidad* consists of several *UPZs*; for example, the *localidad Suba* consists of 12 *UPZs*.[18]

We use this measures since unemployment is one of the main sources of income risk with which our population group is confronted, as Weil (1993) shows that labor-income shocks are a crucial determinant of the precautionary saving motive, and since health risk, another major source of income risk, is rather difficult to quantify, which would only result in an arguably more noisy and less objective measure. In addition, unemployment is quite high in Colombia. DANE estimated the unemployment rate at 8.64 percent in Bogotá for 2013. Using a perception-based measure for unemployment accounts for individuals having better information on their employment prospect than an external observer (Deaton, 1991), which might apply particularly to the informal workforce. The advantage of the (aggregated) subjective measure we use is that it can be considered to be exogenous for a single individual who cannot affect the aggregated unemployment risk. Moreover, using aggregated values increases precision. The assumption that we use is that individuals observe when neighbors lose employment. This makes their subjective individual risk of income loss salient. Objectively, in big cities like Bogotá, where housing is segregated and to a certain degree informal in many areas, it can be assumed that individuals living in the same (poor) neighborhood work on similar job markets and are constrained in relocating for a job, and thus face similar unemployment risks (cf. the spatial mismatch hypothesis going back to Kain (1968), who documented that housing segregation of nonwhite workers in Detroit and Chicago affected their employment outcomes – among them unemployment risk –, and see for example Andersson et al. (2018) for recent empirical support of the now broader question "whether a worker with locally inferior access to jobs is likely to have worse labor market outcomes"). As detailed below, we find support for this conjecture also in our data.

[17] All individuals who took part in the experiment and who were working at the time of the interview answered the question "What is the probability that you will lose your job next year".

[18] There are 20 localities in Bogotá aggregating more than 110 UPZs; *UPZs* with fewer than 25 observations were grouped with their neighboring *UPZ(s)*.

Loss Aversion For the experimental elicitation of loss aversion, we used the non-parametric method for joint elicitation of utility in the loss and the gain domain introduced by Abdellaoui et al. (2007). This method is based on simple lottery choices where individuals compare two lotteries over a series of decision tasks that vary payoffs and probabilities of good and bad states of the world.[19] The elicitation format is thus equivalent to standard procedures used to elicit utility-based risk coefficients (e.g., Holt and Laury, 2005). The method applied corrects for the misperception of probability. Using these choices, it elicits utility points iteratively over a large range of values. The elicited utility points can then be connected to yield a utility function over the gain and loss domain. The advantage of this method is that it is very flexible as it does not require any parametric assumptions over a utility function or probability weighting. Competing methods to determine the shape of the utility function mostly focus on the elicitation of preferences at just one or a few (arbitrarily) selected points in the interval of interest (e.g., Binswanger, 1980; Holt and Laury, 2005). This is also true for the loss-aversion task used in Gächter et al. (2021), which is based on simple binary-outcome lottery choices, such as the method we use, yet leaving probabilities constant. Consequently, it cannot correct for misperceptions of probability. Moreover, it does not allow the derivation of multiple (utility-based) measures of loss aversion building on the various competing definitions used in the literature (e.g., Kahneman and Tversky, 1979; Köbberling and Wakker, 2005).

To test whether participants understood the decision tasks employed here, we included comprehension questions (see Appendix E). Participants could only continue the experiment when they could solve the task correctly. In the ex-post survey for interviewers, they reported that they had the impression that participants understood the task well (average response was 7.9 on a 0 to 10 scale). Moreover, analyzing the data ex-post on the aggregate level, we can reject random decision-making.[20]

[19]Following the terminology of Andersen et al. (2006), this corresponds to an interactive multiple price list (iMPL), where the individual makes choices from refined options according to previous decisions.

[20]More specifically, using binomial tests for the whole set of decision tasks as well as for different subsets of the tasks (e.g., decisions involving non-negative outcomes, decisions involving other probabilities than 0.5), we reject the probability of choosing the left (the right) lottery with a probability of 50%. We can also check single decision items, where we find no evidence for random decision-making neither in decision tasks involving non-negative outcomes nor in those involving non-positive outcomes. For mixed gambles, we cannot reject that the left lottery is chosen in half of the cases (i.e., with a probability of 50%), which is consistent with our findings with respect to loss aversion (based on a different classification strategy, see below): Aggregating

The definitions of loss aversion that we apply and the resulting coefficients of loss aversion build on the elicited utility functions. Following Abdellaoui et al. (2007), we use five different definitions of loss aversion, since so far there is no agreement on a definition of loss aversion, and empirically, measures differ considerably (e.g., Bouchouicha et al., 2019; Brown et al., 2021). To be precise, we rely on the definitions by Kahneman and Tversky (1979), Wakker and Tversky (1993), Bowman et al. (1999), Neilson (2002) and Köbberling and Wakker (2005). We also follow Abdellaoui et al. (2007) in the operationalization of resulting measures of loss aversion (summarized in Table 5 in Appendix D or Table 5 in Abdellaoui et al., 2007), which we explain in detail in Appendix B.1. To incorporate the different definitions of loss aversion, and to ensure that our results are independent of the exact definition of loss aversion, we compute different meta-measures of loss aversion. The first meta-measure is the geometric mean of the coefficients resulting from the definitions by Neilson (2002) and Köbberling and Wakker (2005), available for all individuals.[21] The second meta-measure additionally includes the coefficient based on the definition by Kahneman and Tversky (1979), which we can compute for a subset of 579 participants of the experiment, as this involves derivatives; see Section 5 for details. Finally, we compute a measure relying on all five definitions.

Other Risk Preferences The experimental data were also used to estimate risk preferences corresponding to the curvature of consumption utility m and probability weighting that we use as control variables. We use similar procedures as Abdellaoui et al. (2007). In particular, after appropriate rescaling, we estimate the curvature of utility by estimating the parameters of the following power utility function for the gain and loss domains (see Appendix B.2 for details):

$$m(x) = \begin{cases} -(-x)^a & \text{for } a > 0,\ -1 \leq x < 0 \\ 0.5 \cdot (x)^b & \text{for } b > 0,\ 0 \leq x \leq 1. \end{cases}$$

The estimated parameters, referred to as coefficients of risk aversion in expected utility (EU) theory, are used to describe and classify utility curvature.[22]

individual data, the sample can neither be classified as loss averse nor gain seeking.

[21] We apply the geometric instead of the arithmetic mean, since all coefficients are ratios, thus centered around 1, where it is desirable that coefficients of .5 and 2 have a mean of 1 instead of 1.25. Additionally, the geometric mean is the adequate choice when ranges of single components differ, which is the case for the loss aversion coefficients resulting from the different definitions.

[22] For $x > 0$, the utility function is strictly concave if $0 < b < 1$, linear if $b = 1$, and strictly convex if $b > 1$. For $x < 0$, we have that the function is strictly concave if $a > 1$, linear if $a = 1$, and strictly convex if $0 < a < 1$.

The experimental method proposed by Abdellaoui et al. (2007) also allows a non-parametric estimation of subjective probabilities. We use this method to estimate the objective probability corresponding to a subjective probability of 50 percent.

Time Preferences To elicit the near future impatience or interest rate, which we also use as control variable, we followed the experimental design by Andersen et al. (2008). Participants were asked whether they would prefer to receive an amount x in 30 days or an amount $x(1+r/12)$ with $r > 0$ in 60 days. This question was asked for different and increasing values of r and participants usually switch from choosing x in 30 days to $x(1+r/12)$ in 60 days for a sufficiently high r. This switching point allows the calculation of a lower and an upper bound of the interest rate. Since choices deal with the concept of receiving money, interpreting the results as impatience is likely more accurate.

In addition, we let participants perform the same task with a more distant time-framing (180 vs. 210 days). This allows us to consider consistency in inter-temporal choice.

Other Control Variables In the empirical analysis, we also control for other covariates that have been found to affect the likelihood of saving or the amount of savings directly or indirectly: Age, gender and the relation to the household head (to capture female-headed households), number of children, family size, whether or not parents are still alive (to capture inheritance), exercising habits and the BMI (to capture health status), education, financial literacy, income situation, homeownership, and financial planning behavior (e.g., Conley and Ryvicker, 2004; Finke et al., 2006; Devaney et al., 2007; Sanders and Porterfield, 2010; Fisher and Montalto, 2010, 2011; Van Rooij et al., 2012; Gorbachev and Luengo-Prado, 2019).

5. Field Results

5.1. Descriptive Statistics

Summary statistics are reported in Table 3 in Appendix D. Participants in the study were between the ages of 24 to 87, with a mean age of 49 years. Our sample consists of roughly

70 percent women half of which are single, separated/divorced or widowed.[23] On average, the highest educational attainment was passing the sixth year of school. Financial literacy was also relatively low: the average individual was able to answer roughly only half of the 18 questions concerning, for example, simple math or interest-rate topics correctly. These figures reinforce the need to investigate "behavioral" drivers of saving behavior such as loss aversion. Mainly, they illustrate that the sample at hand is arguably more likely to be guided by heuristics rather than by a sophisticated optimization process when it comes to saving.

The mean monthly income in an average household was 319,000 COP, which at that time was roughly 170 USD. The poverty line at the date of the interview was approximately 155 USD. Half of the sample was assigned to the lowest socioeconomic strata according to the SISBEN classification. This explains that around 85 percent of the sample does not engage in saving money and that the overall mean of savings is 256,000 COP – approximately 130 USD –, thus less than the average per-capita household income per month. The mean savings of those who were actually saving was around 1,761,000 COP, which corresponds to roughly 900 USD. Those reporting non-zero savings save exclusively in cash (27 percent), in a savings account (20 percent), or exclusively for housing (34 percent). The majority of the sample (74 percent) reported carrying out their financial planning on a day-by-day basis, and more than half of the sample never, or hardly ever, exercises. The sample's mean BMI is 25.7, which is close to the Colombian average BMI of 25.9 (World Health Organization, 2014).

Summary statistics on income risk are reported in Table 10 in Appendix D. The average self-reported unemployment risk ranges from 15.2 percent in the *UPZ* Corabastos (*localidad* Kennedy) to 36.5 in Comuneros/Alfonso López (*localidad* Usme). Comparing unaggregated, individual self-reported unemployment risk between these two UPZs shows significant differences, both in the mean and in the distribution, using a two-sided t-test ($p<0.03$) and an exact Kolmogorov-Smirnov test ($p<0.1$), respectively. This suggests that labor markets differ for inhabitants of different neighborhoods, in line with the spatial mismatch hypothesis. The mean (and median) perceived risk of unemployment is about 25%.

[23] According to Colombia's statistical office, a share of 50% of women being single, divorced or widowed is representative for women between 15 and 49 in Bogota (Departamento Administrativo Nacional de Estadística (DANE), 2022).

The summary statistics for the different measures of loss aversion are presented in Table 5 in Appendix D.[24] Depending on the measure used, we find that on average participants display loss aversion. In our lab-in-the-field experiment in Colombia, the lowest mean and median values for the loss aversion coefficients are based on the definitions by Neilson (2002) and Bowman et al. (1999), while the highest coefficients are based on the definition by Köbberling and Wakker (2005). This pattern has also been documented in a recent meta-analysis of empirical estimates of loss aversion (Brown et al., 2021) and by Abdellaoui et al. (2007) (whose protocol we follow; however, they find higher mean and median values for all definitions).[25] Similar to them we find that some individuals display gain-seeking behavior.

In general, the measures of loss aversion display a large and significant degree of positive correlation with each other (summarized in Table 6 in Appendix D). The only measure that has a low degree of correlation with all other measures is the Bowman measure that is positively and significantly correlated only with the Neilson measure.

5.2. Empirical Results

To test our hypotheses, we run three different models, as explained with detail in Section 3. As the outcome variable – savings – is skewed, non-negative and contains a large number of zeros resulting from not engaging in saving, we apply a Negative Binomial Hurdle model. This model is a two-part model, where the probability of engaging in saving and the amount of savings is estimated separately and by different models. Using a logit-model, the probability of 'passing the hurdle' and actually engaging in saving is estimated. The (strictly positive) amount

[24] For all individuals, we can compute loss aversion coefficients based on the definitions by Neilson (2002) and Köbberling and Wakker (2005). Other definitions are more difficult to operationalize, in particular the ones relying on derivatives. Because some choice tasks involved stochastic dominant options for some individuals, which was a result from the iterative characteristic of the protocol, the number of available utility points differs. We exclude choices resulting from such choice tasks from the analysis, following e.g., Bleichrodt and Pinto (2000), who elicit probability-weighting functions non-parametrically with a comparable protocol. As a result, this hinders the operationalization of the loss aversion coefficients in some cases.

[25] Abdellaoui et al. (2007) report that loss aversion coefficients based on the definitions by Bowman et al. (1999) and Neilson (2002) have the lowest mean (0.74 and 1.07) and median values (0.74 and 0.43) in their study, where the latter is below 1 for both definitions. The highest value for the mean and median they obtain for loss aversion, as defined by Köbberling and Wakker (2005), with a mean of 8.27 and a standard deviation of 15. Other studies focusing on monetary or health outcomes have found individual mean values between 0.11 and 19.861 (Brown et al., 2021) relying on different definitions of loss aversion.

of savings once the hurdle is passed is estimated using a Truncated Negative Binomial model. In Appendix C.1, we explain the model and its choice in more detail, and give the formula for computation of marginal effects on the outcome. In Appendix C.2, we discuss possible alternative models.

Table 1 presents the estimated coefficients. We present the results separately by the likelihood to save (upper panel) and the amount of savings, given that an individual is actually saving (i.e., conditional saving, lower panel). All models control for main socioeconomic variables affecting savings, as well as for region and occupation (if indicated), and risk and time preference and time consistency measures (see the list in Tables 3 and 5 in Appendix D).[26] Further, to account for the small number of clusters (UPZs) in our study, we use the wild cluster bootstrap procedure (Cameron et al., 2008) adapted to ML estimation (score bootstrap, Kline and Santos, 2012) for inference.

Estimation results of a direct test of Hypothesis 1 are presented below 'Model 1' in Table 1. The first two columns present the results for the entire sample, while the columns labeled 'Loss-Averse Subsample' present the results for the sub-population of loss-averse individuals in the sample. The likelihood to save is practically unrelated to income risk both for the entire sample and for the loss-averse sub-sample. Yet, consistently with the predictions, we find that income risk is positively and significantly related with the amount of savings, given that an individual is actually saving. This result holds both for the entire sample, although only significant at the 10% level, and in particular for the sub-population of loss-averse individuals ($p < 0.01$). Thus, in total we find support for Hypothesis 1:[27]

Result 1. *An increase in income risk is associated with an increase in savings (due to an increase in conditional savings) for loss-averse individuals (as specified in Hypothesis 1).*

[26] When restricting the sample to the loss-averse sub-population, however, we can only control for a subset of these variables due to the low number of observations.

[27] Strictly speaking, for the loss-averse sub-sample and the extensive margin (likelihood of saving), the coefficient of income risk is not exactly zero, and negative (-0.001; average marginal effect between 0 and -0.0001). As for the amount of saving, the coefficient is 0.130 (average marginal effect 2.825), we see from the formula to compute overall marginal effects (Equation (4) in the Appendix) that the estimated coefficient from the extensive margin is indeed negligible for the change in (overall) predicted savings.

Estimation results of a direct test of Hypothesis 2 are reported in 'Model 2' of Table 1. With respect to the likelihood of saving, coefficients of both income risk and loss aversion are positive, but not statistically different from zero. Regarding the intensive margin, we observe that an increase in loss aversion is associated with an increase in the amount of conditional savings. This result holds when controlling for income risk, and it is robust to using alternative measures of loss aversion.[28] We summarize our findings with respect to Hypothesis 2:

Result 2. *An increasing degree of loss aversion is associated with an increase in savings (due to an increase in conditional savings).*

Model 3 allows us to test Hypothesis 3, but we may also use these insights to gain a deeper understanding of Results 1 and 2 – in addition to the results from the direct tests (Model 1 and Model 2). We find that the likelihood to save is, by and large, unrelated with income risk (and only the interaction term is significantly positive, and only at the 10% level). Yet, supporting the theoretical model and in line with the results from Model 1, we find that conditional savings are positively and significantly associated with income risk. As the measure of loss aversion is centered around one, the coefficient of income risk shows the relation between income risk and the likelihood to save or savings for a loss-neutral agent. As the interaction terms of loss aversion and income risk are positive and significant (albeit for the likelihood to save only at a significance level of 10%), they indicate that the precautionary saving motive, i.e., savings associated with income risk, is even larger for (more) loss-averse individuals. Hence, we find additional support for Hypothesis 1 from an estimation that uses the full sample and controls for all variables listed in Tables 3 and 5 in Appendix D.

The two just mentioned positive coefficients of the interaction terms between loss aversion and income risk ($p < 0.1$ for the likelihood to save, and $p < 0.01$ for conditional savings) indicate that the positive relation of savings with income risk is larger, the more loss-averse individuals are. Hence, this result supports Hypothesis 3.

Result 3. *The (positive) association between income risk and savings is larger for more loss-averse individuals.*

Finally, the coefficient of loss aversion indicates a negative relation with conditional savings. As the measure of income risk is centered around its mean value, this finding, and the results

[28]Estimation results for different measures of loss aversion are presented in Table 12 in Appendix D.

from Model 2, suggest that Result 2 is driven by individuals facing a somewhat elevated income risk. Indeed, re-estimating Model 3 with an indicator variable for facing a high income risk, defined here as an average perceived local unemployment risk above the 75% percent quantile (Model 3b, column labeled 'Risk: High IR'), reveals an unambiguously positive and significant relation of an increase in loss aversion for those facing a high income risk – at the intensive and extensive margin (see the linear combinations in Panel C of Table 1).[29]

In terms of the magnitude of the relation, the estimations from Model 1 imply that, when the average perceived unemployment risk rises by one standard deviation from 24.7 to 30.7, for a loss-averse sample, conditional savings are predicted (at mean covariate values) to be more than COP 600,000 (about USD 315) higher, while the average total savings are predicted to increase by about COP 90,000, or roughly USD 45 (according to Equation (4) in Appendix C.1, i.e., when weighting for the probability of actually saving).[30] This corresponds to 33 percent higher total savings.

5.3. Robustness Tests

In the literature, various definitions of loss aversion have been proposed, but none can be considered standard (Abdellaoui et al., 2007). To address this issue, we have constructed different meta-measures of loss aversion combining two or more definitions. Our results are robust to using any of the different meta-measures of loss aversion for both Model 2, and Model 3. We report results of these robustness tests in Table 12 and 13 in Appendix D.

[29] In Tables 7 and 8 in Appendix D, we provide summary statistics for the sub-sample of individuals facing a high income risk. The sub-sample looks fairly comparable to the whole sample – with the exception of savings. It is maybe noteworthy that the higher income risk seems to come along with a slight shift in planning horizon: In the sample with high income risk, subjects' (financial) planning horizon is less often the next day, and more often the next month; this seems like a shift of 5% of the (sub-)sample, underlying the salience of income risk in that sub-sample.

[30] To predict the total increases in conditional savings one might instead of just using mean values of covariates for the prediction predict individually and average the predictions. This leads to larger values.

Table 1: Results from Estimating Model 1, Model 2, and Model 3 Using a Negative Binomial Hurdle Model – Colombian Data

	Model (1)				Model (2)		Model (3)		Model (3b) Binary Classification	
			Loss-Averse Subsample						Risk: High IR ($\geq$ 75% Quantile)	
Panel A: Likelihood of Saving										
Income Risk (IR)	0.026	(1.01)	-0.001	(-0.02)	0.027	(1.07)	0.025	(1.02)		
High IR (=1)									0.278	(0.95)
Loss Aversion (LA)					0.042	(1.61)	0.032	(1.01)	0.011	(0.36)
LA $\times$ IR							0.012*	(1.70)		
LA $\times$ High IR (=1)									0.068*	(2.05)
Panel B: Amount of Savings										
Income Risk (IR)	0.139*	(1.93)	0.130***	(6.17)	0.145**	(2.06)	0.150***	(2.67)		
High IR (=1)									2.254***	(2.90)
Loss Aversion (LA)					0.062**	(2.13)	-0.177***	(-3.57)	-0.117**	(-2.76)
LA $\times$ IR							0.022***	(5.29)		
LA $\times$ High IR (=1)									0.216***	(4.34)
Panel C: Linear Combinations										
Likelihood of Saving										
Loss Aversion									0.079***	(3.108)
Amount of Savings										
Loss Aversion									0.099***	(4.535)
AIC	1080		189		1072		1225		1231	
Controls	25		6		25		25		25	
Region	Yes		No		Yes		Yes		Yes	
Occupation	Yes		No		Yes		Yes		Yes	
Observations	640		97		640		640		640	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Wild cluster (score) bootstrapped t-/chi2-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various saving devices, see Section 4. In this Negative Binomial Hurdle model, the participation equation estimates the likelihood to engage in saving, while the second equation estimates conditional savings – the amount of savings, given that a person is saving any money. The coefficient of loss aversion is centered at one and measured by a continuous and experimentally elicited meta-measure, based on the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005), see Appendix B.1. Income risk is centered at the mean; see Section 4 for details. We control for variables listed in Tables 3 and 5 in Appendix D. In particular, we include individual level control variables for risk aversion, impatience, inconsistent time preferences, and furthermore include individual control variables for region, and for the working sectors according to the ISIC classification of economic activities, if indicated. We account for the cluster structure (at the UPZ level) in our data by using wild cluster (score) bootstrapping (Cameron et al., 2008; Kline and Santos, 2012), considering the “small” number of clusters.

6. Addressing Causality: A Lab Experiment on Precautionary Saving and Loss Aversion

Despite the plausibility of the assumption that perceived income risk, regionally aggregated, is exogenous on the individual level, i.e., that it cannot be affected by a single individual, our results up to now have to be interpreted as correlational. Being able to control for a broad range of individual characteristics increases confidence that indeed loss aversion and income uncertainty drive the reported results. Yet, we cannot clearly establish causality using cross-sectional data.

To overcome this limitation, we conducted a lab experiment in which we endogenously alter income risk in a within-subject design to investigate whether an increase in income risk indeed *causally* triggers an increase in saving because of loss-aversion. The incentivized experiment consists of two parts. In the first part, individual measures of loss aversion are elicited, following again Abdellaoui et al. (2007). For the second part, the savings experiment, we follow Xu et al. (2022), in particular the tasks focusing on individuals' reaction to income risk. While Xu et al. (2022) test for the effects of higher-order risk preferences on savings in these tasks as predicted by the classical precautionary saving theory (Leland, 1968), we consider the impact of loss aversion here, as modelled in Kőszegi and Rabin (2009) or Aizenman (1998).

Choices were incentivized: Part 1 and part 2 were payed out each with a probability of 50%. If part 1 was randomly selected for payout, one of the 108 decisions was randomly chosen to be payoff relevant (all with equal probability). If the participant's choice involved a lottery, the computer randomly resolved the uncertainty according to the probabilities stated, which was clearly explained to participants (see the instructions in Appendix F). If part 2 was randomly selected for payout, again with equal probability, one of the 14 saving decisions was selected to be payout relevant by the computer. As for the first part, any remaining uncertainty was resolved by the computer according to the specified probabilities, which was explained to participants.

6.1. Participants and Data Collection

We recruited 200 participants from the joint participant pool of labs connected to our institutions using hroot (Bock et al., 2014). Besides the general requirement of the lab to be fluent in the local language, we did not impose additional restrictions. The experiment was implemented in otree (Chen et al., 2016). All instructions were presented on screen, and participants could proceed in their own speed (see Appendix F for the instructions). Research assistants or one of the authors were available at any time of the experiment to resolve any remaining questions privately. Part one took on average 22 minutes, and task two took about 7 minutes (excluding about 14 minutes of general and task specific instructions including trial rounds and examples, and about 10 minutes for displaying results and filling in a brief questionnaire). Participants were paid in cash before leaving the lab.

During the sessions, the lab team noticed that 7 participants had major understanding problems due to a lack of proficiency in the local language, and that 2 subjects were obviously making decisions completely at random. Either in the open-ended question in the post-experimental questionnaire or when leaving the lab, 8 subjects indicated that they felt that their choices should not be analyzed due to a lack of concentration over the approximately one-hour session (5 participants), or because they felt that saving was meaningless for reasons outside the experiment (e.g., because of negative or very low interest rate on their bank's savings account and geopolitical uncertainties; 3 participants). In the analysis, we exclude these 17 observations.

6.2. The Savings Task

In the savings tasks, participants confronted 14 two-period saving decisions. In period one participants received a fixed income \$$x$, while in the second period, participants could receive either a high \$$a$ or a low \$$b$ payment with equal probability. The expected value in period two was kept constant across decision tasks, but the spread was altered. In particular, x could take one of two values $x \in \{48, 51\}$ and the a_b combinations could be from the following set: $\{70\_20, 66\_24, 62\_28, 58\_32, 54\_36, 50\_40, 45\_45\}$.[31] In each task, consisting of a combination of first-period income \$$x$ and second-period income combinations \$$a\_b$, participants can save or borrow up to \$20 in \$2 steps free of interest. A slider implementation (see Figure 1) allows to save (borrow) by moving the bar upwards (downwards). In the case of saving \$$s$, the

[31] To investigate consumption smoothing and its relation to risk aversion, Xu et al. (2022) investigate eight different parameters for x ranging from 33 to 54 in steps of EUR 3. As we focus on participants' reaction to income risk, we limit this set of fixed period-one income figures to the two values 48 and 51, to have two decisions for each of the seven possible income spreads per individual. We opt for values above period-two expected income to encourage saving.

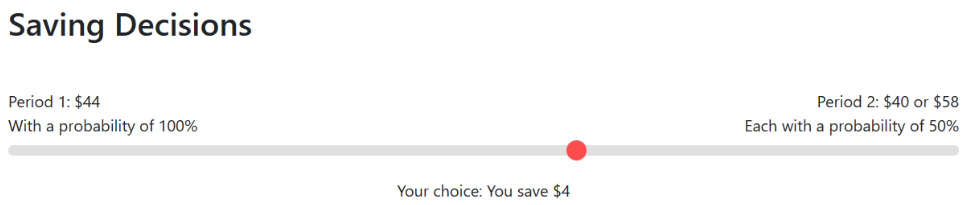

Figure 1: Saving Decision

Note: Using a slider, participants could increase (decrease) their saving in each saving task, thereby decreasing (increasing) the fixed period-one payment and increasing (decreasing) the two possible period-two payments. The red dot indicating the current choice was revealed only once a decision was selected by clicking on the slider. In this example, for the first-period income, $\$x = 48$ before saving $\$s = 4$, and the two second-period payments were $\$a = 54$ and $\$b = 36$ (i.e., a spread of \$18) before adding $\$s = 4$.

period-one payment $\$x$ is reduced by $\$s$ and the two possible period-two payments $\$a$ and $\$b$ are increased by $\$s$. The tasks were presented in random order on individual basis, except for the first task, which was always fixed with first-period income $\$x = 48$ and second-period income $\$a = b = 45$ to study the *reaction* to income risk. If the savings experiment was randomly selected for payment, only the payout from one decision situation and one period were implemented and paid. If the decision of the second period was selected for payment, a random draw determined whether participants received $\$a$ or $\$b$. The conversion rate in the saving experiment was $\$1 =$ EUR 0.8.

Recall from Section 2 that we expect saving to increase with income risk (Hypothesis 1), that loss-averse individuals will save more (Hypothesis 2), and that the precautionary motive, i.e., the increase in saving due to income risk, increases with the degree of loss aversion (Hypothesis 3).

6.3. Results

Following the empirical strategy in Section 3, and analogously to the analysis of the field data in Section 5, we estimate three models using a meta measure of loss aversion (Meta Measure 1), and present the results in Table 2. As expected, neither are there excessive zero-saving decisions nor is there a highly skewed outcome variable to be observed in the lab, and we thus, following Xu et al. (2022), model the amount of savings using a random effects panel model (instead of a hurdle model as in Section 5 with the Colombian field data).[32]

[32]When facing income risk, participants save (or borrow) in 93.35% of saving choices and only one participant ($<1\%$) does never save (or borrow). Skewness is below 1, whereas for the field data, it is about 10.

Consistent with a precautionary saving motive, we find that saving increases when participants are confronted with higher income risk. This result holds for the entire sample and for the sub-sample of loss-averse individuals. A $10 (EUR 8) increase in the spread of period-two income causes a $1.23 (EUR 0.98) increase in saving; in the loss-averse sub-sample, the corresponding increase amounts to $1.62 (EUR 1,30). Hence, the lab experimental findings causally confirm Result 1.

Model 2 investigates the role of an increase in the degree of loss aversion on saving – keeping constant the degree of income risk. For the average amount of income spread in the savings task ($34), the coefficient of the degree of loss aversion is negative, but insignificant. Using lab data, we therefore cannot find support for Hypothesis 2 as derived from the model by Kőszegi and Rabin (2009) regarding the role of an increase in the degree of loss aversion on saving.

Finally, we test Hypothesis 3 asserting that more loss-averse individuals increase their saving more as a reaction to income risk increases. The results from estimating Model 3 confirm this prediction and we find that the interaction coefficient of income risk and loss aversion is significantly positive at the 1% significance level. Hence, also with respect to Hypothesis 3, the lab experimental findings causally confirm our findings from the field, in particular Result 3: Income uncertainty leads to higher saving, and the higher the degree of loss aversion, the higher is the increase in saving triggered by income uncertainty.

Regarding Hypothesis 2, field data implied that the overall impact of loss aversion on the average precautionary saving motive is positive mainly due to those facing high or very high levels of income risk. Recalling the (insignificant) negative coefficient of loss aversion in Model 2 with lab data suggests a similar picture here, considering the arguably lower income risk induced in laboratory savings tasks. We re-investigate this relationship by re-estimating Model 3 using a binary categorization of income risk tasks instead of the continuous spread measure used in Model 1-Model 3. Decision tasks involving an income risk of at leat the 75% quantile of income spread ($42) are coded 1, decision tasks below the 75% quantile are coded 0.

Indeed, the estimated coefficients in Model 3b show that when facing high or very high levels of income risk, the degree of loss aversion is associated with a significantly larger degree of saving compared to decisions with lower income risk (interaction term), implying a positive association

Table 2: Results from Estimating Model 1, Model 2, and Model 3 Using a Random Effects Panel Regression – Laboratory Experiment

	Model (1)				Model (2)		Model (3)		Model (3b) Binary Classification	
			Loss-Averse Subsample						Risk: High IR (≥ 75% Quantile)	
Panel A: Saving										
Income Risk (IR)	0.123***	(9.658)	0.162***	(5.835)	0.123***	(9.656)	0.132***	(10.476)		
High IR (= 1)									3.525***	(9.409)
Loss Aversion (LA)					-0.221	(-0.749)	-0.100	(-0.336)	-0.464	(-1.533)
LA × IR							0.028***	(2.775)		
LA × High IR (= 1)									0.851***	(3.093)
Panel B: Lin. Combination										
Loss Aversion									0.386	(0.284)
Observations	2562		546		2562		2562		2562	

* p < 0.1, ** p < 0.05, *** p < 0.01.

Note: The dependent variable is the amount saved in a given saving decision task, and for every individual, we have 14 saving decisions, see Section 6.2. Following Xu et al. (2022), we employ a random effects panel regression to account for the multiple (dependent) observations from every individual. The coefficient of loss aversion is centered at one and measured by a continuous and experimentally elicited meta-measure, based on the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005), see Appendix B.1. Income risk is centered at the mean.

with precautionary saving. In absolute terms, this coefficient is about twice as large as the (insignificantly negative) coefficient of loss aversion without income risk or for low levels of income risk (main effect). Put differently, the net effect of loss aversion at high levels of income risk is positive, but insignificantly so (as the linear combination in Panel B is insignificantly different from zero due to the imprecision of the negative coefficient on loss aversion for low levels of income risk). Thus, although more loss-averse individuals save (insignificantly) less at average levels of income risk ($34), as soon as income risk increases, they compensate their "undersaving" compared to the less loss-averse participants, and, for high levels of income risk, end up accumulating a higher share of their income for saving than their less loss-averse counterparts (although, insignificantly so within the limited income risk triggered in the lab).

7. Discussion

In this paper we study the relation between income uncertainty, loss aversion and savings. The field and lab data show that when facing higher income risk, loss-averse individuals hold, on average, higher savings compared to when facing lower income risk (Hypothesis 1). While previous research finds empirical support for a precautionary motive for savings, this is the first paper that uses micro data and incentivized measures of loss aversion to empirically establish the theoretical (positive) link between income risk, loss aversion, and saving. We show – using lab and field data – that the increase in savings that is associated with an increase in income risk is higher, the larger the degree of loss aversion (Hypothesis 3). Especially this latter contribution is relevant as in line with the classical theory of precautionary saving (e.g., Leland, 1968), it would be possible that we observe higher savings independently of loss aversion – for example because individuals are, on average, prudent (i.e., because they have a convex marginal consumption utility, $m''' > 0$).[33] Our results clearly link the degree of loss aversion to the precautionary saving motive.

The empirical findings support the essence of the reference-dependent inter-temporal consumption model by Kőszegi and Rabin (2009), i.e., that (the degree of) loss aversion is positively related to the precautionary saving motive. Further, the results suggest that the degree of income risk matters to assess whether a more loss-averse individual saves more: The field data indicate that for below-average levels of income risk more loss-averse individuals save less – in line with assumptions in, e.g., Thaler and Benartzi (2004). For high levels of income risk, due to their stronger reaction to income risk, more loss-averse individuals save more. Given that the income risk in the laboratory is somewhat limited, we see the null result with respect to Hypothesis 2 in the laboratory to be in line with these findings.

While the empirical findings in the lab show that individuals increase saving in the face of income uncertainty due to a precautionary savings motive, our results from the field can only suggest this. We must recognize that we are unable to cleanly identify the specific factor that

[33] See the empirical patterns on the prevalence of prudence (e.g., Noussair et al., 2014) and the discussion following Equation 2 in Section 2.

motivates increases in savings. Future field work should extend the analysis to disentangle the effect of precautionary behavior from other factors affecting savings (e.g., financial inclusion).[34]

It is noteworthy that the support we find in our field data for the predictions of the model by Kőszegi and Rabin (2009) come from the intensive margin mainly – that is: The model describes behavior of those that engage in saving rather than explaining the shift from not saving to saving. This of course might have many reasons, such as a budget constraint. In light of the fact that the majority of the Colombian sample performs their financial planning on a day-by-day basis, it may as well be the case that many simply lack the financial sophistication or fail to plan (Bone et al., 2009; Brown et al., 2009; Duffy and Li, 2019). Alternatively, Pagel (2017) shows that time inconsistencies could explain a hump-shaped life cycle consumption pattern. Early in life, consumption is low due to a precautionary savings motive. However, as uncertainty resolves over time, this motive is dominated by time-inconsistent over-consumption, eventually leading to declining consumption towards the end of life. Future work should assess how dynamic inconsistent plans interact with loss aversion and income uncertainty to explain savings.

Although our results are consistent with the loss-aversion based precautionary saving motive (Aizenman, 1998; Siegmann, 2002; Kőszegi and Rabin, 2009; Park, 2016; Pagel, 2017), there could be concerns regarding the particular characteristics of our sample in Colombia. First, we used convenient sampling and recruited participants who were at home during working hours (9-16). This could lead to measurement error if the respondents are not the ones participating in financial decisions in the household. We think that this is unlikely. Most of the respondents report to have been working for all 12 months in the last year (77% of female respondents and 85% of male respondents). In addition, most of the respondents reported to be head of the household (56% of female and 82% of male respondents) indicating that they have a large degree of control and knowledge of financial decisions. For the single, divorced or widowed women, it is clear that they take their financial decisions alone (recall that all participants are older than 24). For the married women - perceiving themselves as head of household or not –, it is also to be expected that they are involved in (financial) household decision-making: According to the Gender Data

[34] This could be achieved, for example, with a more precise measure of precautionary savings, as used by Deidda (2013), or with wealth to income measures to quantify the share of precautionary savings (Guiso et al., 1992; Carroll and Samwick, 1998; Lusardi, 1998).

Portal of the World Bank (2022), 80.2% of married women say that they alone or jointly have the final say in major household purchase decisions (for healthcare decisions, it is even 92.6%).

Moreover, we find that our sample is relatively less loss-averse than the one in Abdellaoui et al. (2007) – a sample of business school students in France. Other studies also find considerably lower shares of loss aversion at the individual level (e.g., Schmidt and Traub, 2002; Bleichrodt and Pinto, 2002). Had, however, our implementation of the method resulted in an equal underestimation of the degree of loss aversion for every individual, and we should thus assume higher coefficients for every individual, this would leave our results qualitatively unaffected.[35] Further, according to the model by Park (2016), when loss aversion is low (i.e., the individual is loss-tolerant), the decision maker may deviate into consuming more if they face a small level of uncertainty relative to the intensity of their loss aversion. Our empirical results suggest that the precautionary saving motive is still relevant for our population group.

Another concern one might have regarding our loss aversion measures, particularly with respect to those from the field, is that individuals with similar preferences might behave differently in the preference elicitation tasks – depending on the income volatility they face. It is thus important to highlight that the relation between loss aversion measures and precautionary saving in the field is correlational. Our measures from the laboratory experiment are not susceptible to this concern, as the measure of income risk that is used in the analysis is cleanly controlled in a within-subject design and therefore independent of the elicited parameter of loss aversion.

With respect to saving, it is sometimes argued that the poor are too poor to save. In our sample, the mean household per-capita income lies near the poverty level, and employment is mostly informal. Our findings, however, build on variation *within* the sample, and establish that despite possible budget constraints, income uncertainty and loss aversion induce savings. We believe that these findings are important exactly *because* our sample might face budget constraints, as it illustrates how even those who have very little to save could be triggered to increase their saving, for example by illustrating the possible consequences of uncovered income shocks, and highlighting the associated "loss" in consumption.[36]

[35] These results can be obtained from the authors.

[36] This approach could, however, also induce stress; caution is needed when implementing such an approach with vulnerable population groups.

It would be of great interest to validate this study's findings with other data from other samples which have other socio-demographic characteristics. A first step in that direction is the lab experiment that we implemented with university students. The extension is relevant as it allows to generalize the finding to populations from the Western world. Further, the controlled environment allows us to build a panel data set, where we can obtain ex-ante measures of risk preferences and loss aversion at the time that we exogenously manipulate income uncertainty. While it is already compelling to find support for the field experimental findings, further work should elaborate on the generalizability of our findings.

8. Conclusion

In this paper, we have tested whether the theoretical predictions of reference-dependent intertemporal models of consumption and saving with respect to loss aversion and precautionary savings can be empirically supported (Aizenman, 1998; Siegmann, 2002; Berkelaar et al., 2004; Kőszegi and Rabin, 2009; Park, 2016; Pagel, 2017). More specifically, guided by the model of Kőszegi and Rabin (2009), we have tested whether loss-averse individuals who face higher income risk hold higher savings, which would be consistent with a precautionary savings motive based on loss aversion (Hypothesis 1). Our results from the lab and from the field support this hypothesis. An increase in income risk by one standard deviation (i.e., an increase in the probability of job loss by about 6 percentage points) is associated with an increase of about 33 percent of total savings for the loss-averse sub-population in our Colombian data.

Secondly, we have tested whether individuals who exhibit a higher degree of loss aversion hold more savings than individuals with a lower degree of loss aversion, given that they face income risk (Hypothesis 2). The empirical analysis provides mixed evidence regarding this hypothesis. The field data show a positive association between loss aversion and savings on average, where the effect is mainly driven by individuals facing a somewhat elevated level of income risk. Perhaps less surprisingly, the lab experiment, inducing income uncertainty at levels of about half a daily wage for students, does not support a positive effect of loss aversion on savings, and if anything indicates no effect, or even a negative effect, on average.

Lastly, we have investigated whether the increase in savings associated with an increase in income risk is larger the higher the degree of loss aversion is (Hypothesis 3). Also for this last hypothesis, we find very strong support both with the field and experimental data.

Our findings can be used to calibrate models, but also to inform policy-makers. Our study's results indicate that savings are counter-cyclical, with (loss-averse) individuals saving more in anticipation of negative income shocks. However, our results suggest that the capacity to save for low-income populations might be limited. A large share of the population does not hold any savings, and those who do accumulate relatively low values. In addition, the response to income risk allows savings equivalent to one third of the average per-capita monthly income. These findings suggest a relevant role for social-protection programs to smooth household consumption. For instance, unemployment insurance could be an essential policy instrument to stabilize household consumption and protect families against asset depletion that could trigger poverty traps (Gerard and Naritomi, 2021; Landais and Spinnewijn, 2021; Kroft and Notowidigdo, 2016).

Although we admit that raising the saving rate remains a challenge in developing countries, several studies have demonstrated that this is possible (e.g., Ashraf et al., 2006; Dupas and Robinson, 2013; Karlan et al., 2014), and the heterogeneity we observe in our data also supports this claim. Moreover, we like to stress that the relevance of our results is in no way limited to developing nations. Social-security systems are far from perfect in numerous countries in the Western world, too, as our examples in the introduction illustrate. Thus, in both settings our insights regarding a saving model that accounts for behavioral aspects should be equally helpful.

References

Abdellaoui, M., H. Bleichrodt, and C. Paraschiv (2007). Loss aversion under prospect theory: A parameter-free measurement. *Management Science 53*(10), 1659–1674.

Abeler, J., A. Falk, L. Goette, and D. Huffman (2011). Reference points and effort provision. *American Economic Review 101*(2), 470–492.

Aizenman, J. (1998). Buffer stocks and precautionary savings with loss aversion. *Journal of International Money and Finance 17*(6), 931–947.

Allen, E. J., P. M. Dechow, D. G. Pope, and G. Wu (2016). Reference-dependent preferences: Evidence from marathon runners. *Management Science 63*(6), 1657–1672.

Andersen, S., C. Badarinza, L. Liu, J. Marx, and T. Ramadorai (2019). Reference dependence in the housingmarket. Discussion Paper 14147, Centre for Economic Policy Research.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2006). Elicitation using multiple price list formats. *Experimental Economics 9*(4), 383–405.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008). Eliciting risk and time preferences. *Econometrica 76*(3), 583–618.

Andersson, F., J. C. Haltiwanger, M. J. Kutzbach, H. O. Pollakowski, and D. H. Weinberg (2018). Job displacement and the duration of joblessness: The role of spatial mismatch. *The Review of Economics and Statistics 100*(2), 203–218.

Ashraf, N., D. Karlan, and W. Yin (2006). Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *The Quarterly Journal of Economics 121*(2), 635–672.

Baiardi, D., M. Magnani, and M. Menegatti (2020). The theory of precautionary saving: An overview of recent developments. *Review of Economics of the Household 18*, 513–542.

Barberis, N. C. (2013). Thirty years of prospect theory in Economics: A review and assessment. *Journal of Economic Perspectives 27*(1), 173–196.

Bartling, B., L. Brandes, and D. Schunk (2015). Expectations as reference points: Field evidence from professional soccer. *Management Science 61*(11), 2646–2661.

Bayer, C., R. Lütticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Benartzi, S. and R. H. Thaler (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics 110*(1), 73–92.

Benito, A. (2006). Does job insecurity affect household consumption? *Oxford Economic Papers 58*(1), 157–181.

Berkelaar, A. B., R. Kouwenberg, and T. Post (2004). Optimal portfolio choice under loss aversion. *The Review of Economics and Statistics 86*(4), 973–987.

Binswanger, H. P. (1980). Attitudes toward risk: Experimental measurement in rural India. *American Journal of Agricultural Economics 62*(3), 395–407.

Bleichrodt, H. and J. L. Pinto (2000). A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science 46*(11), 1485–1496.

Bleichrodt, H. and J. L. Pinto (2002). Loss aversion and scale compatibility in two-attribute trade-offs. *Journal of Mathematical Psychology 46*(3), 315–337.

Bock, O., I. Baetge, and A. Nicklisch (2014). hroot: Hamburg Registration and Organization Online Tool. *European Economic Review 71*, 117–120.

Bone, J., J. D. Hey, and J. Suckling (2009). Do people plan? *Experimental Economics 12*, 12–25.

Bouchouicha, R., L. Deer, A. G. Eid, P. McGee, D. Schoch, H. Stojic, J. Ygosse-Battisti, and F. M. Vieider (2019). Gender effects for loss aversion: Yes, no, maybe? *Journal of Risk and Uncertainty 59*(2), 171–184.

Bowman, D., D. Minehart, and M. Rabin (1999). Loss aversion in a consumption-savings model. *Journal of Economic Behavior & Organization 38*(2), 155–178.

Brown, A. L., Z. E. Chua, and C. F. Camerer (2009). Learning and visceral temptation in dynamic saving experiments. *The Quarterly Journal of Economics 124*(1), 197–231.

Brown, A. L., T. Imai, F. Vieider, and C. Camerer (2021). Meta-analysis of empirical estimates of loss-aversion. Working Paper 8848, CESifo.

Browning, M. and T. F. Crossley (2001). The life-cycle model of consumption and saving. *Journal of Economic Perspectives 15*(3), 3–22.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997). Labor supply of New York City cabdrivers: One day at a time. *The Quarterly Journal of Economics 112*(2), 407–441.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics 90*(3), 414–427.

Card, D. and G. B. Dahl (2011). Family violence and football: The effect of unexpected emotional cues on violent behavior. *The Quarterly Journal of Economics 126*(1), 103–143.

Cardenas, J. C. and J. Carpenter (2013). Risk attitudes and economic well-being in Latin America. *Journal of Development Economics 103*, 52–61.

Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives 15*(3), 23–45.

Carroll, C. D., M. B. Holm, and M. S. Kimball (2021). Liquidity constraints and precautionary saving. *Journal of Economic Theory 195*, 105276.

Carroll, C. D., J. Overland, and D. N. Weil (2000). Saving and growth with habit formation. *American Economic Review 90*(3), 341–355.

Carroll, C. D. and A. A. Samwick (1998). How important is precautionary saving? *The Review of Economics and Statistics 80*(3), 410–419.

Chen, D. L., M. Schonger, and C. Wickens (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij (2020). Consumption uncertainty and precautionary saving. *The Review of Economics and Statistics 102*(1), 148–161.

Christian, C., L. Hensel, and C. Roth (2019). Income shocks and suicides: Causal evidence from Indonesia. *The Review of Economics and Statistics 101*(5), 905–920.

Clark, A. E., C. D'Ambrosio, and S. Ghislandi (2016). Adaptation to poverty in long-run panel data. *The Review of Economics and Statistics 98*(3), 591–600.

Conley, D. and M. Ryvicker (2004). The price of female headship: Gender, inheritance, and wealth accumulation in the United States. *Journal of Income Distribution 13*(3/4), 41–56.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica 59*, 1221–1248.

Deidda, M. (2013). Precautionary saving, financial risk, and portfolio choice. *Review of Income and Wealth 59*(1), 133–156.

Departamento Administrativo Nacional de Estadística (DANE) (2022). Mujeres por estado conyugal entre 15 y 49 años. https://geoportal.dane.gov.co/servicios/atlas-estadistico/src/Tomo_I_Demografico/mapa/4.2_Mujeres15a49a%C3%B1osSegunEstadoConyugal_2005.pdf. Last accessed: 2023-05-17.

Dercon, S. (2010). Risk, poverty, and human development: What do we know, what do we need to know? In R. Fuentes-Nieva and P. A. Seck (Eds.), *Risk, Shocks, and Human Development: On the Brink*, pp. 15–39. London, UK: Palgrave Macmillan.

Devaney, S. A., S. T. Anong, and S. E. Whirl (2007). Household savings motives. *Journal of Consumer Affairs 41*(1), 174–186.

D'Exelle, B. and A. Verschoor (2015). Investment behaviour, risk sharing and social distance. *The Economic Journal 125*(584), 777–802.

Dreze, J. H. and F. Modigliani (1975). Consumption decisions under uncertainty. In W. T. Ziemba and R. G. Vickson (Eds.), *Stochastic optimization models in Finance*, pp. 459–486. Cambridge, MA: Academic Press.

Duffy, J. and Y. Li (2019). Lifecycle consumption under different income profiles: Evidence and theory. *Journal of Economic Dynamics and Control 104*, 74–94.

Dupas, P. and J. Robinson (2013). Why don't the poor save more? Evidence from health savings experiments. *American Economic Review 103*(4), 1138–1171.

Dynan, K. E. (1993). How prudent are consumers? *Journal of Political Economy 101*(6), 1104–1113.

Eeckhoudt, L., A. M. Fiori, and E. R. Gianin (2016). Loss-averse preferences and portfolio choices: An extension. *European Journal of Operational Research 249*(1), 224–230.

Ericson, K. M. M. and A. Fuster (2011). Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments. *The Quarterly Journal of Economics 126*(4), 1879–1907.

Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of New York City taxi drivers. *American Economic Review 98*(3), 1069–1082.

Federal Reserve Board (2020). Report on the economic well-being of U.S. households in 2019. https://www.federalreserve.gov/publications/2020-economic-well-being-of-us-households-in-2019-preface.htm. Last accessed: 2023-05-17.

Fehr, E. and L. Goette (2007). Do workers work more if wages are high? Evidence from a randomized field experiment. *American Economic Review 97*(1), 298–317.

Finke, M. S., S. J. Huston, and D. L. Sharpe (2006). Balance sheets of early boomers: Are they different from pre-boomers? *Journal of Family and Economic Issues 27*(3), 542–561.

Fisher, P. J. and C. P. Montalto (2010). Effect of saving motives and horizon on saving behaviors. *Journal of Economic Psychology 31*(1), 92–105.

Fisher, P. J. and C. P. Montalto (2011). Loss aversion and saving behavior: Evidence from the 2007 U.S. Survey of Consumer Finances. *Journal of Family and Economic Issues 32*(1), 4–14.

Gerard, F. and J. Naritomi (2021). Job displacement insurance and (the lack of) consumption-smoothing. *American Economic Review 111*(3), 899–942.

Giles, J. and K. Yoo (2007). Precautionary behavior, migrant networks, and household consumption decisions: An empirical analysis using household panel data from rural China. *The Review of Economics and Statistics 89*(3), 534–551.

Gorbachev, O. and M. J. Luengo-Prado (2019). The credit card debt puzzle: The role of preferences, credit access risk, and financial literacy. *The Review of Economics and Statistics 101*(2), 294–309.

Grogger, J. T. and R. T. Carson (1991). Models for truncated counts. *Journal of Applied Econometrics 6*(3), 225–238.

Guariglia, A. and B.-Y. Kim (2004). Earnings uncertainty, precautionary saving, and moonlighting in Russia. *Journal of Population Economics 17*, 289–310.

Guiso, L., T. Jappelli, and D. Terlizzese (1992). Earnings uncertainty and precautionary saving. *Journal of Monetary Economics 30*(2), 307–337.

Gächter, S., E. J. Johnson, and A. Herrmann (2021). Individual-level loss aversion in riskless and risky choices. *Theory and Decision 92*(3-4), 599–624.

Holt, C. A. and S. K. Laury (2005). Risk aversion and incentive effects: New data without order effects. *American Economic Review 95*(3), 902–904.

Imas, A., S. Sadoff, and A. Samek (2016). Do people anticipate loss aversion? *Management Science 63*(5), 1271–1284.

Jalan, J. and M. Ravallion (2001). Behavioral responses to risk in rural China. *Journal of Development Economics 66*(1), 23–49.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292.

Kain, J. F. (1968). Housing segregation, negro employment, and metropolitan decentralization. *The Quarterly Journal of Economics 82*(2), 175–197.

Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2016). Getting to the top of mind: How reminders increase saving. *Management Science 62*(12), 3393–3411.

Karlan, D., A. L. Ratan, and J. Zinman (2014). Savings by and for the poor: A research review and agenda. *Review of Income and Wealth 60*(1), 36–78.

Karle, H., D. Engelmann, and M. Peitz (2022). Student performance and loss aversion. *The Scandinavian Journal of Economics 124*(2), 420–456.

Karle, H., G. Kirchsteiger, and M. Peitz (2015). Loss aversion and consumption choice: Theory and experimental evidence. *American Economic Journal: Microeconomics 7*(2), 101–120.

Kőszegi, B. and M. Rabin (2009). Reference-dependent consumption plans. *American Economic Review 99*(3), 909–936.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica 58*(1), 53–73.

Kline, P. and A. Santos (2012). A score based approach to wild bootstrap inference. *Journal of Econometric Methods 1*(1), 23–41.

Köbberling, V. and P. P. Wakker (2005). An index of loss aversion. *Journal of Economic Theory 122*(1), 119–131.

Kroft, K. and M. J. Notowidigdo (2016). Should unemployment insurance vary with the unemployment rate? theory and evidence. *The Review of Economic Studies 83*(3), 1092–1124.

Landais, C. and J. Spinnewijn (2021). The value of unemployment insurance. *The Review of Economic Studies 88*(6), 3041–3085.

Leland, H. E. (1968). Saving and uncertainty: The precautionary demand for saving. *The Quarterly Journal of Economics 82*(3), 465–473.

Lise, J. (2013). On-the-job search and precautionary savings. *The Review of Economic Studies 80*(3), 1086–1113.

Loayza, N., L. Servén, and N. Sugawara (2009). Informality in Latin America and the Caribbean. Working Paper 4888, World Bank Policy Research.

Lugilde, A., R. Bande, and D. Riveiro (2018). Precautionary saving in Spain during the great recession: Evidence from a panel of uncertainty indicators. *Review of Economics of the Household 16*, 1151–1179.

Lugilde, A., R. Bande, and D. Riveiro (2019). Precautionary saving: A review of the empirical literature. *Journal of Economic Surveys 33*(2), 481–515.

Lusardi, A. (1998). On the importance of the precautionary saving motive. *American Economic Review 88*(2), 449–453.

Markle, A., G. Wu, R. White, and A. Sackett (2018). Goals as reference points in marathon running: A novel test of reference dependence. *Journal of Risk and Uncertainty 56*(1), 19–50.

Michler, J. D. and J. V. Balagtas (2017). The importance of the savings device in precautionary savings: Empirical evidence from rural Bangladesh. *Agricultural Economics 48*(2), 129–141.

Neilson, W. S. (2002). Comparative risk sensitivity with reference-dependent preferences. *Journal of Risk and Uncertainty 24*(2), 131–142.

Noussair, C. N., S. T. Trautmann, and G. v. d. Kuilen (2014). Higher order risk attitudes, demographics, and financial decisions. *The Review of Economic Studies 81*(1), 325–355.

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance 53*(5), 1775–1798.

Pagel, M. (2017). Expectations-based reference-dependent life-cycle consumption. *The Review of Economic Studies 84*(2), 885–934.

Park, H. (2016). Loss aversion and consumption plans with stochastic reference points. *The BE Journal of Theoretical Economics 16*(1), 303–336.

Paxton, J. and F. Zhuo (2011). Economic shocks and savings behavior by the rural poor. *Economics Bulletin 31*(4), 3286–3293.

Pope, D. G. and M. E. Schweitzer (2011). Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes. *American Economic Review 101*(1), 129–157.

Rosokha, Y. and K. Younge (2020). Motivating innovation: The effect of loss aversion on the willingness to persist. *The Review of Economics and Statistics 102*(3), 569–582.

Sanders, C. K. and S. L. Porterfield (2010). The ownership society and women: Exploring female householders' ability to accumulate assets. *Journal of Family and Economic Issues 31*(1), 90–106.

Sandmo, A. (1970). The effect of uncertainty on saving decisions. *The Review of Economic Studies 37*(3), 353–360.

Schmidt, U. and S. Traub (2002). An experimental test of loss aversion. *Journal of Risk and Uncertainty 25*(3), 233–249.

Siegmann, A. (2002). Optimal saving rules for loss-averse agents under uncertainty. *Economics Letters 77*(1), 27–34.

Stampini, M., M. Robles, M. Sáenz, P. Ibarrarán, and N. Medellín (2016). Poverty, vulnerability, and the middle class in Latin America. *Latin American Economic Review 25*(4), 1–44.

Thaler, R. H. and S. Benartzi (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy 112*(S1), S164–S187.

Tovar, J. and M. Urrutia (2017). The impact of social safety net programs on household savings in Colombia. *The Developing Economies 55*(1), 23–37.

Van Rooij, M. C., A. Lusardi, and R. J. Alessie (2012). Financial literacy, retirement planning and household wealth. *The Economic Journal 122*(560), 449–478.

Wakker, P. (2008). Explaining the characteristics of the power (CRRA) utility family. *Health Economics 17*(12), 1329–1344.

Wakker, P. and A. Tversky (1993). An axiomatization of cumulative prospect theory. *Journal of Risk and Uncertainty 7*(2), 147–175.

Weil, P. (1993). Precautionary savings and the permanent income hypothesis. *The Review of Economic Studies 60*(2), 367–383.

Wooldridge, J. M. (2013). *Introductory econometrics: A modern approach* (5 ed.). Mason, OH: South-Western, Cengage Learning.

World Bank (2014). *Global financial development report 2014: Financial inclusion*, Volume 2. Washington, DC: World Bank.

World Bank (2022). Gender Data Portal. https://genderdata.worldbank.org/countries/colombia/. Last accessed: 2023-05-17.

World Health Organization (2014). *Global status report on noncommunicable diseases 2014*. Geneva, Switzerland: WHO Press.

Xu, Y., M. Boksem, C. N. Noussair, S. T. Trautmann, G. van de Kuilen, and A. Sanfey (2022). Higher order risk preferences and economic decisions. Working paper, Utrecht School of Economics, Utrecht University.

A. Theoretical Framework: Details (For Online Publication)

A.1. General Version of the Two-period Model by Kőszegi and Rabin (2009)

As in Section 2, we assume that an individual has to distribute wealth, W, for consumption across two periods such that $W = c_1 + c_2$, where c_t denotes consumption in period t for $t = 1, 2$. As in the main text, we consider the case in which wealth is stochastic and uncertainty is resolved in the second period.

Consumption in the first period (and thus saving) is determined by maximizing the expectation of the sum of instantaneous utilities u_t in both periods, where no discounting is assumed, i.e.,

$$U = u_1(c_1) + \mathbb{E}[u_2(c_2)]. \tag{3}$$

As in the simplified version of the model introduced in the main text, individuals are assumed to choose their favorite credible consumption plan before the first period starts (i.e., in period $t = 0$). Credible means that they anticipate whether or not they would be able to stick to the plan, and only consider those plans where they do not see an incentive to deviate from later on.[37] Favourite means that there are possibly several such credible plans, and the decision-maker chooses his or her preferred one according to the maximization principle. This plan is called preferred personal equilibrium (PPE) by Kőszegi and Rabin (2009) and at the time of planning in period $t = 0$, it leads to possibly stochastic 'rational beliefs' $F_{0,1}$ and $F_{0,2}$ about consumption in Period 1 and Period 2. Mathematically, these beliefs are simply probability distributions assigning a probability to any possible consumption level. Plans about consumption in period t that are made in the same period (i.e., $F_{t,t}$) assign a probability of 1 to the actual consumption level c_t. When uncertainty is resolved and consumption decisions are implemented, plans are updated and lead to new beliefs.

Instantaneous utility in periods $t = 1, 2$ is given by

$$u_t = m(c_t) + \sum_{\tau=t}^{2} \gamma_{t,\tau} N(F_{t,\tau}|F_{t-1,\tau}),$$

where $m(\cdot)$ is consumption utility that is three times differentiable, increasing and strictly concave, and corresponds to a "classical utility function". The 'gain-loss utility', $N(F_{t,\tau}|F_{t-1,\tau})$, reflects utility gains or losses due to changes in current 'beliefs' $F_{t,\tau}$ compared to former 'beliefs' $F_{t-1,\tau}$ about contemporaneous ($\tau = t$) and future ($\tau > t$) consumption. Depending on the

[37] Details about how these plans are formed are given in Appendix A.2 or in Kőszegi and Rabin (2009).

distance of a period $\tau \geq t$ in the future, the impact of changes in beliefs about consumption in that period via the 'gain-loss utility' differs, which is reflected by weights $\gamma_{t,\tau} \geq 0$ with $\gamma_{t,t} = 1$. For simplicity, we use the notation $\gamma_{1,2} = \gamma$. The weight $\gamma_{1,2} = \gamma$ is decisive for an individual to adhere to her plan, i.e., to resist overconsuming in the first period relative to the previously set consumption level, as explained below.

'Gain-loss utility' N compares every percentile of the distributions of consumption according to 'beliefs' $F_{t,\tau}$ and $F_{t-1,\tau}$, using a "universal gain-loss utility function" μ. More specifically, for a possibly discrete distribution F_d, $c_{F_d}(p/100)$ is a percentile for $0 \leq p \leq 100$ with $p \in \mathbb{N}$ if $F_d(c_{F_d}(p/100)) \geq p/100$ and $F_d(c) < p/100$ for all $c < c_{F_d}(p/100)$. Then, gain-loss utility from the change in beliefs from $F_{t-1,\tau}$ to $F_{t,\tau}$ is defined as

$$N(F_{t,\tau}|F_{t-1,\tau}) = \sum_{p=1}^{100} \mu(c_{F_{t,\tau}}(p/100), c_{F_{t-1,\tau}}(p/100)),$$

where

$$\mu(\hat{c}, \tilde{c}) = \begin{cases} \eta(m(\hat{c}) - m(\tilde{c})) & \text{if } \hat{c} \geq \tilde{c} \\ -\lambda\eta(m(\tilde{c}) - m(\hat{c})) & \text{if } \hat{c} < \tilde{c}. \end{cases}$$

for two consumption levels $\hat{c}$ and $\tilde{c}$, m as defined above and parameters $\eta > 0$ and $\lambda > 0$.[38]

The parameter $\eta > 0$ simply scales the difference in consumption utility, and $\lambda > 0$ may account for loss-averse ($\lambda > 1$) or gain-seeking ($\lambda < 1$) behavior.

The parameter $\gamma \geq 0$ 'discounts' anticipated future gains or losses in 'gain-loss' utility that affect utility already in period 1. For $\gamma > 1/\lambda$, the anticipated future loss is weighted high enough to prevent the consumer from deviating from the optimal ex-ante plan, i.e., they resist overconsuming; see Proposition 5 in Kőszegi and Rabin (2009). When $\lambda > 1$, following Kőszegi and Rabin (2009), we can assume $\gamma < 1$. As we allow for gain-seeking behavior, i.e., $\lambda < 1$, we leave γ unrestricted, to allow for $\gamma > 1/\lambda$. Then, the proof of Proposition 8 in Kőszegi and Rabin (2009) holds for $\lambda < 1$, although they do not consider this case.

If the agent resists deviating from the plan, instantaneous utility in Period 1 is given by

$$u_1 = m(c_1) + N(F_{1,1}|F_{0,1}) + \gamma N(F_{1,2}|F_{0,2}) = m(c_1),$$

[38] This choice of the "gain-loss utility function" fulfills certain desirable characteristics of a reference-dependent utility function for $\lambda > 1$; see Kőszegi and Rabin (2009), p. 914. In particular, it fulfills "the explicit or implicit assumptions" about the 'value function' by Kahneman and Tversky (1979), as formulated by Bowman et al. (1999).

as beliefs do not change in the first period (i.e., $F_{0,t} = F_{1,t}$ for $t = 1, 2$), since in addition to adherence to the plan, no uncertainty is resolved. In Period 2, utility is given by

$$u_2 = m(c_2) + N(F_{2,2}|F_{1,2}).$$

With that, the optimization problem can be solved by equalizing the marginal utility of saving and consumption in the first period.

If the agent cannot resist deviating from the ex-ante optimal plan, their PPE specifies a higher consumption level in Period 1 compared to the optimal one; see Proposition 5 in Kőszegi and Rabin (2009).

A.2. Rational Beliefs

In this Appendix, we explain the intuition behind 'rational beliefs'. We refer to Kőszegi and Rabin (2009) for a precise definition.

'Beliefs' are the result of a plan: They "must be rationally based on credible plans for state-contingent behavior".[39] One concept of what a credible plan could be was termed 'preferred personal equilibrium (PPE)' by Kőszegi and Rabin (2009) and was used in their text, although they note that other theories of forming beliefs could also be combined with their model. Roughly speaking, a plan is a PPE if it is the preferred "plan among those that are credible". A plan is credible if it maximizes the mathematical expectation of the reference-dependent utility in every period given the beliefs which the plan induced *and* if continuation plans are consistent. That is: If an individual plans for very low consumption in Period 1 in order to save for Period 2, but would not make the same choice if solving the maximization problem in Period 1 – e.g., because they are present-biased or cannot live with such a low level of consumption –, this would not be a credible plan, and it is not a PPE. Using backwards induction, they would anticipate their behavior in Period 1 and consume more in Period 1 from the beginning until their entire plan is consistent with solutions evolving from a similar maximization process in Period 1. This PPE reflects the idea that individuals anticipate the implications of their plans and only make plans they know they would adhere to them.

[39] The most simple example of a state-contingent plan could be: "If things go well, I will spend x\$ for consumption in Period 1. If things do not work out well, I will only spent y\$ in this period" (where $x > y > 0$).

A.3. Proofs

Proof of Proposition 1. This proof follows the rationale of the proof of Proposition 8 in Kőszegi and Rabin (2009).

We prove that the derivative of the marginal utility of increasing savings with respect to λ is positive. Equivalent to the argument in the proof of Kőszegi and Rabin's Proposition 8, this implies that $dc_1/d\lambda < 0$ for both $\gamma > 1/\lambda$ and $\gamma \leq 1/\lambda$, since in the first case, the ex-ante optimal plan involves a lower c_1 and the person adheres to this plan. In the latter case, a higher marginal utility in Period 2 makes a lower c_1 become consistent. Furthermore, since, for $\gamma \leq 1/\lambda$, the chosen c_1 will be higher than for $\gamma > 1/\lambda$, see Kőszegi and Rabin (2009), a lower c_1 will become consistent, as the agent adheres to the ex-ante optimal plan for a lower γ.

The derivation of marginal utility of increasing savings is due to Kőszegi and Rabin (2009): Let F be the cumulative distribution function of the (mean-zero) random variable y. The expected utility in Period 2 is

$$\begin{aligned}
&\int m(c_2 + sy)\, dF(y) + \iint \mu(m(c_2 + sy) - m(c_2 + sy'))\, dF(y')\, dF(y) \\
&= \int m(c_2 + sy)\, dF(y) \\
&\quad - \frac{1}{2}\eta(\lambda - 1) \iint m(c_2 + s\max\{y, y'\}) - m(c_2 + s\min\{y, y'\})\, dF(y')\, dF(y).
\end{aligned}$$

Hence, the derivative of the expected utility in Period 2 with respect to c_2, i.e., the marginal utility from increasing savings is

$$\begin{aligned}
&\int m'(c_2 + sy)\, dF(y) \\
&\quad + \frac{1}{2}\eta(\lambda - 1) \iint m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\})\, dF(y')\, dF(y).
\end{aligned}$$

Now, unlike in the proof or Proposition 8 in Kőszegi and Rabin (2009), we take the derivative of the expression above with respect to λ:

$$\frac{1}{2}\eta \iint m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\})\, dF(y')\, dF(y).$$

This derivative is positive for any strictly concave m, any $s > 0$, $\eta > 0$, and any non-degenerate random variable y. Thus, the marginal utility from increasing savings is an increasing function of λ. □

Proof of Corollary 1. As in the proof of Proposition 1, the derivative of the marginal utility from increasing savings with respect to λ is given by

$$\frac{1}{2}\eta \iint m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\})\, dF(y')\, dF(y).$$

The derivative of this expression with respect to s evaluated at $s = 0$ is

$$\frac{1}{2}\eta(-m''(c_2)) \iint |y' - y|\, dF(y')\, dF(y),$$

which is positive for any strictly concave consumption utility function m, $\eta > 0$ and any non-degenerate random variable y. □

B. Data: Details (For Online Publication)

B.1. Details on the Measures of Loss Aversion

In this section we describe how we operationalized the different measures of loss aversion with our data, following Abdellaoui et al. (2007).

Kahneman-Tversky (KT) Kahneman and Tversky (1979) define an individual as loss-averse, if for all amounts of money x the utility μ of receiving this amount is lower than the disutility of losing that same amount, i.e., if $\forall\, x > 0 : -\mu(-x) > \mu(x)$. A natural coefficient of loss aversion emerging from this definition is $-\mu(-x)/\mu(x)$ for every elicited amount $x > 0$. If $\mu(-x)$ for any of these eight elicited amounts of money $x > 0$ was not elicited, it was linearly interpolated. As the coefficient of loss aversion, we took the median of the computed coefficients.

Neilson (N) Neilson (2002) proposes computing the ratio of 'relative steepness', which is the utility value $\mu(x)$ divided by the corresponding x-value. This figure incorporates information about steep parts of the utility function at any point of the interval of interest – even in flat regions. If the relative steepness of the utility function over the loss domain is bigger than the one on the gain domain at any point, the individual is classified as loss averse, i.e., $\mu(-x)/x \geq \mu(y)/y$, $\forall\, x, y > 0$. For this definition, we computed the coefficient of loss aversion as the ratio of the infinum of $\mu(-x)/(-x)$ over the supremum of $\mu(y)/y$.

The remaining definitions rely on the steepness of the utility function as expressed by the derivative of the latter on both domains.

Wakker-Tversky (WT) Wakker and Tversky (1993) suggest applying the concept of Kahneman and Tversky (1979) to the derivative of utility, i.e., to compare the value of the derivative of the utility function for gains and losses 'point-wise' at certain absolute values: $\mu'(-x) > \mu'(x), \forall\, x > 0$. At every elicited utility point $x > 0$ on the gain domain, the derivative $\mu'(x)$ was operationalized as the mean of the two connecting slopes to the left-hand side and to the right-hand side. $\mu'(-x)$ was operationalized as the slope of the linearly interpolated utility function at the point $-x$. Similar to the case for the definition by Kahneman and Tversky (1979), a natural coefficient emerging from the definition $\mu'(-x) > \mu'(x), \forall\, x > 0$, is $\mu'(-x)/\mu'(x)$ for $x > 0$. In this case, we also took the median of the coefficients thus computed.

Bowman (B) Bowman et al. (1999) propose performing this comparison 'domain-wise', that is, $\mu'(-x) > \mu'(y), \forall\, x, y > 0$. As in the case for the definition by Neilson (2002), the definition $\mu'(-x) > \mu'(y),\ \forall\, x, y > 0$ can be transformed into a coefficient of loss aversion by computing $\inf \mu'(-x)/\sup \mu'(y)$ for $x, y > 0$, where the derivatives where operationalized as just described.

Köbberling-Wakker (KW) Finally, Köbberling and Wakker (2005) define an individual as loss-averse if the slope of the utility function on the left-hand side of the reference point is steeper than the slope of the utility function on the right-hand side of the reference point: $\mu'(0_-) > \mu'(0_+)$. The natural coefficient of loss aversion resulting from this definition, $\mu'(0_-)/\mu'(0_+)$, was computed as the ratio of slopes connecting 0 with the elicited utility points that are closest to 0 on both domains.

B.2. Parametric Estimation of a Power Utility Function

General Form for Positive Arguments Usually, the power family is defined for $x > 0$ by

$$m(x) = \begin{cases} x^b & \text{for } b > 0 \\ \ln(x) & \text{for } b = 0 \\ -x^b & \text{for } b < 0. \end{cases}$$

Considering Non-Positive Arguments Since $\ln(x)$ is not defined for $x < 0$, the case $b = 0$ must be excluded, if negative arguments are of interest. Furthermore, $b < 0$ has to be excluded as well, if the point $x = 0$ is to be considered.[40] Thus, when allowing for gains and losses, the *power family* reduces to

$$m(x) = \begin{cases} -(-x)^a & \text{for } a > 0,\ x < 0 \\ x^b & \text{for } b > 0,\ x \geq 0. \end{cases}$$

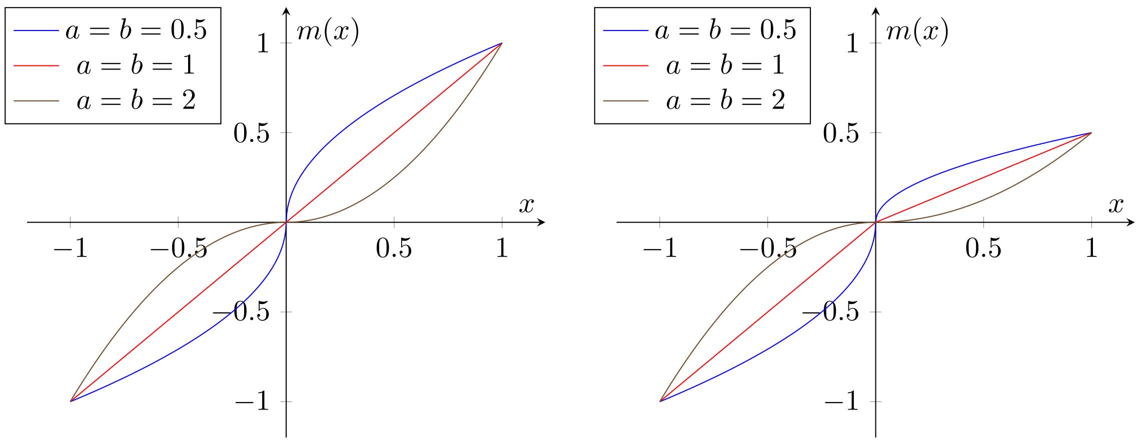

(a) Curvature of the power family for different values of a and b.

(b) Estimated power utility functions plotted for different values of a and b.

Figure 2: Illustration of the Power Family Utility Function with Different Values of a and b

Figure 2(a) illustrates the curvature of the power family for different values of a and b.

Rescaling Arguments Arguments x of the utility function must be rescaled in order to lie within the interval $[-1, 1]$ for all the subjects in the study in order to be able to compare estimated parameters.

Due to the method used, the minimal x-value observed is $L_1 = -5{,}000{,}000$. Thus, for losses, we need a transformation $x \mapsto -\frac{x}{L_1}$, where $x \in [L_1, 0]$.

For Gains, $G_{0.5}$ is the maximum x-value for any individual, we therefore transform $x \mapsto \frac{x}{G_{0.5}}$, where $x \in [0, G_{0.5}]$.

Rescaling Outputs By the method chosen, we need to have $m(L_1) = -1$, $m(0) = 0$ and $m(G_{.5}) = .5$. We check this: For the negative domain, we have

$$m(L_1) = -\left(\frac{L_1}{L_1}\right)^a = -(1)^a = -1,$$

[40] Wakker (2008, p.1336) gives a less technical explanation: “With both positive and negative x present, a negative power a or b generates an infinite distance between gains and losses. Such a phenomenon is not empirically plausible, so that negative a and b should then not be expected to occur.”

independent of $a > 0$, so there is no need to rescale outputs. However, for the positive domain,

$$m(G_{0.5}) = \left(\frac{G_{0.5}}{G_{0.5}}\right)^b = 1^b = 1,$$

independent of $b > 0$. Therefore, and also to have estimates comparable for the negative and the positive domain, we rescale $m(x)$ for $x \geq 0$ and set:

$$m(x) = 0.5 \cdot \left(\frac{x}{G_{0.5}}\right)^b \qquad \text{for} \qquad x \geq 0.$$

Note that we could also leave the estimation formula untouched and multiply our outcomes by the factor 2, making them lie within the interval $[0, 1]$ instead of $[0, .5]$.

Estimation Equation The final estimation equation is thus

$$m(x) = \begin{cases} -\left(\frac{x}{L_1}\right)^a & \text{for } a > 0,\ x < 0 \\ 0.5 \cdot \left(\frac{x}{G_{0.5}}\right)^b & \text{for } b > 0,\ x \geq 0. \end{cases}$$

This equation is illustrated in Figure 2(b).

Curvature In order to classify a utility function as convex or concave based on the estimated values of the parameters a or b, we can deduct the curvature of the utility function from Figure 2 for the given values of a and b. Analytically, for classifying an individual's utility function, we calculate the second derivative of the estimated utility function.

$$m''(x) = \begin{cases} -\left(\frac{x}{L_1}\right)^a \cdot \frac{1}{x^2} \cdot a(a-1) & \text{for } a > 0,\ x < 0 \\ 0.5 \cdot \left(\frac{x}{G_{0.5}}\right)^b \cdot \frac{1}{x^2} \cdot b(b-1) & \text{for } b > 0,\ x > 0, \end{cases}$$

where $x = 0$ has to be excluded from the domain.

We immediately see that for $x > 0$,

$$m''(x) \begin{cases} < 0 \text{ thus } m \text{ strictly concave} & \text{if } 0 < b < 1 \\ = 0 \text{ thus } m \text{ linear} & \text{if } b = 1 \\ > 0 \text{ thus } m \text{ strictly convex} & \text{if } b > 1, \end{cases}$$

and for $x < 0$ we have

$$m''(x) \begin{cases} < 0 \text{ thus } m \text{ strictly concave} & \text{if } a > 1 \\ = 0 \text{ thus } m \text{ linear} & \text{if } a = 1 \\ > 0 \text{ thus } m \text{ strictly convex} & \text{if } 0 < a < 1. \end{cases}$$

C. Results: Details (For Online Publication)

C.1. Econometric Model

The outcome variable used in our analysis – savings (in 100,000 COP) – does not include negative values and is therefore a limited dependent variable according to the definition in Wooldridge (2013, Chapter 17). Furthermore, the empirical frequency of zeros in the distribution of the amount of savings in our sample exceeds the frequency of zeros according to any commonly used theoretical distribution in such cases (e.g., the Poisson distribution or the Negative Binomial distribution). This is to be expected, since not everybody actually engages in saving. Thus, the outcome variable is a so-called Corner Solution Response.[41]

The distribution of the value of saving in our sample is skewed, and values are reported repeatedly and are usually divisible by 100,000 COP. Therefore, we should assume a discrete rather than a continuous dependent variable. Given these characteristics of the outcome variable, we apply a Negative Binomial Hurdle model to study the relationship between income risk, loss aversion, and savings. The Poisson Hurdle model is nested in the Negative Binomial Hurdle model we fit and differences between the log-likelihoods of both models mostly exceed 100 by far. This indicates that a likelihood ratio test (conservatively assuming the test statistic to follow a chi-square distribution with one degree of freedom) would reject the hypothesis of no overdispersion.

This model is a so-called two-part model, where the probability of engaging in savings and the amount of savings is estimated separately by different models. For the Hurdle models applied here, the likelihood of both equations can be calculated separately. Using a logit-model, the probability 'that the hurdle is passed' and that a person engages in savings is estimated. The second model estimates the amount of savings once the hurdle is passed, using a Truncated Negative Binomial model. In Appendix C.2, we discuss alternative models and their suitability in this context.

[41] The options to deny the response or to indicate that they did not know about the amount of savings were allowed and treated separately. Four respondents denied answering and five respondents did not know the amount of savings they held at the time of the interview. Together, this corresponds to about 1% of the respondents whose savings amount we could not observe. These cases were excluded from the analysis.

Following Grogger and Carson (1991), we compute marginal effects of loss aversion and income risk on the predicted amount of unconditional savings using the estimates resulting from fitting Model 3 with a Negative Binomial Hurdle model. Denoting savings for individual i with Y_i, the overall marginal effect of X_{ih}, i.e., of covariate h for individual i, on his or her predicted savings can be computed as

$$\frac{\partial \mathbb{E}(Y_i|X_i)}{\partial X_{ih}} = \frac{\partial}{\partial X_{ih}}[\mathbb{E}(Y_i|X_i, Y_i > 0)][1 - F(0)] + \mathbb{E}(Y_i|X_i, Y_i > 0)\frac{\partial}{\partial X_{ih}}[1 - F(0)], \quad (4)$$

where $1 - F(0)$ is the share of the population for which we observe $Y_i > 0$. This means the overall effect can be decomposed into two effects: The effect on those who are saving, weighted by the probability of saving, plus the effect on the proportion that 'passes the hurdle' and is saving, weighted by the mean amount of savings in the saving population. We compute marginal effects using mean values of covariates, unless otherwise indicated.

C.2. Discussion: Model Choice

In this part, we briefly discuss alternatives to the model chosen and assess their appropriateness in the setting of this paper.

Usually, OLS regression is a suitable starting point for modelling empirical relationships. However, a large share of the non-savers with zero COP of savings could mask relationships observed for the fraction of participants that actually saves. It seems appropriate to take the large share of the non-savers observed in our data into account when selecting a suitable model.

A Tobit model is frequently used in similar situations. Here, it is not suitable. A central assumption of the Tobit model is that the process determining participation is the same as the process determining the amount of saving. The signs of the coefficients of the independent variables in Table 1 differ in the two equations where many are significantly different from zero, showing that this assumption is violated. Second, normality and homoscedasticity of the dependent variable model are prerequisites for using a Tobit model. In contrast to OLS, where departures from these assumptions still lead to unbiased and consistent estimates, it is less clear how sensitive the Tobit model is to departures from these assumptions. The empirical distribution of the outcome variable we observe in our data is discrete. This observed empirical distribution is a rather bad approximation of any continuous probability distribution, so the assumption of normality is not likely to hold.

More flexible models for corner solution responses that can model the participation process and the savings process separately are – in addition to the Hurdle model applied in this study – so-called inflated models. For example, the Zero-Inflated Poisson model or the Zero-Inflated Negative Binomial model for the case of a discrete dependent variable.

Zero-inflated models rely on the assumption that a zero COP value of savings can be the result of two cases: In the first case, an individual would decide to save and then chooses a saving amount of zero. In the second case, an individual would decide not to save at all. We believe that the first case is rather unrealistic, since we did not ask for changes in savings in a given limited time, but rather look at the stock of savings. We therefore conclude that these models are not appropriate in our setting.

It is noteworthy that the excess zeros in the distribution of the outcome variable are not a problem of data observability, where models for censored data or sample correction models (e.g., the Heckman model) would be adequate. Only for around 1 percent of the participants are data actually missing, and these cases were excluded.

When only focusing on the positive amount of savings, no special care is needed to account for excessive zeros in the distribution of the outcome variable. In such cases, a traditional OLS model could be applied, or a log OLS model, if we expect the relationship to be proportional to the response.

Given the discrete character of the outcome variable, and its heavily non-symmetric empirical distribution, a model that accounts for this characteristic should be applied, such as the Zero-Truncated Poisson or the Zero-Truncated Negative Binomial model. The latter is the second part of the two-part model we apply, the Negative Binomial Hurdle model. Thus, if not accounting for excess zeroes, we would model conditional savings in the same way that we do in this study, while accounting for a large proportion of non-savers.

D. Further Results and Robustness (For Online Publication)

Table 3: Summary Statistics – Colombian Data

	Mean	s.d.	Min	Max	Obs.
Individual Information					
Age	49.0	13.4	24	87	640
Male (=1)	0.28	0.45	0	1	640
Relationship to head of HH					
Head of household (=1)	0.64	0.48	0	1	640
Partner (=1)	0.23	0.42	0	1	640
Son/Daughter or their partner (=1)	0.07	0.25	0	1	640
Other (=1)	0.06	0.24	0	1	640
Household Characteristics					
Number of adult household members	2.8	1.4	1	12	640
Number of adolescents	1.2	1.3	0	7	640
Father still alive (=1)	0.31	0.46	0	1	640
Mother still alive (=1)	0.51	0.50	0	1	640
Exercising					
Every day (=1)	0.17	0.37	0	1	640
At least once a week (=1)	0.18	0.38	0	1	640
At least once a month (=1)	0.09	0.28	0	1	640
Never or hardly ever (=1)	0.57	0.50	0	1	640
Other Health Indicators					
BMI	25.7	4.3	12.9	43.0	640
Education					
Highest year passed	5.8	3.3	0	11	640
Financial literacy score (max. 18)	9.3	3.4	0	16	640
Financial Situation of the Household					
SISBEN Level 2 (=1)	0.50	0.50	0	1	640
Size of safety net (# Persons)	2.5	3.5	0	60	640
Monthly HH income per capita[a]	3.19	2.26	0.01	18.00	640
Market price of house[a]	180.10	408.86	0.00	3000.00	640
Debt[a]	17.24	65.68	0.00	588.04	640
Savings[a]	2.56	13.91	0.00	200.00	640
Engaging in saving (=1)	0.15	0.35	0	1	640
Conditional savings	17.61	32.82	0.20	200.00	93
Planning Horizon					
Day to day (=1)	0.74	0.44	0	1	640
Next months (=1)	0.18	0.38	0	1	640
Next year (=1)	0.05	0.21	0	1	640
Next two to five years (=1)	0.02	0.14	0	1	640
Next five to ten years (=1)	0.01	0.11	0	1	640

Note: [a]Figures reported in 100,000 COP.

Table 4: Summary Statistics: Income Risk – Colombian Data

	Mean	s.d.	Median	Min	Max	Obs.
Local Unemployment Risk (in pc)	24.7	6.0	25.4	15.2	36.5	640

Note: The unemployment risk results from self-reported individual figures in our survey that are averaged at the *UPZ* level; see Section 4 for details.

Table 5: Summary Statistics of Experimental Measures – Colombian Data

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Kahneman-Tversky (KT) $-\mu(-x)/\mu(x)$	1.1	2.7	0.4	0.1, 1.1	579
Neilson (N) $(\mu(-x)/-x)\,/\,(\mu(y)/y)$	0.2	0.5	0.0	0.0, 0.1	640
Wakker-Tversky (WT) $\mu'(-x)/\mu'(x)$	12.3	110.9	0.1	0.0, 0.3	564
Bowman (B) $\inf \mu'(-x)/\sup \mu'(x)$	0.1	0.2	0.0	0.0, 0.0	564
Köbberling-Wakker (KW) $\mu'(0_-)/\mu'(0_+)$	10.9	76.6	0.2	0.0, 1.0	640
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	1.1	4.6	0.1	0.0, 0.4	640
Meta Measure 2 (KT, KW, N)	1.0	3.1	0.1	0.0, 0.6	579
Meta Measure 3 (all)	0.3	1.0	0.1	0.0, 0.2	509
Impatience					
Near future impatience	29.6	15.2	22.0	16.0, 50.0	640
Increase in patience over time	0.3	16.3	0.9	-2.9, 0.9	640
Risk Preferences					
Utility Curvature: Gain Domain	6.0	29.9	0.7	0.2, 2.5	640
Utility Curvature: Loss Domain	8.0	16.0	1.1	0.5, 3.5	640
Probability Weighting: Gain Domain	41.5	32.9	40.6	9.4, 71.9	640
Probability Weighting: Loss Domain	68.5	28.5	78.1	46.9, 96.9	640

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail. Near future impatience is the mean annual interest rate, see Section 4. Utility curvature is the parameter of a power utility function and probability weighting is the probability that is perceived as 50%; see Section 4.

Table 6: Correlation Coefficients for Loss-Aversion Measures – Colombian Data

	KT	N	WT	B	KW	Meta 1	Meta 2	Meta 3
						KW, N	KT, KW, N	All
Kahneman-Tversky (KT)	1.000							
Neilson (N)	0.677***	1.000						
Wakker-Tversky (WT)	0.664***	0.281***	1.000					
Bowman (B)	0.020	0.088*	0.002	1.000				
Köbberling-Wakker (KW)	0.380***	0.536***	0.054	-0.019	1.000			
Meta Measure 1 (KW, N)	0.535***	0.780***	0.139***	-0.003	0.901***	1.000		
Meta Measure 2 (KT, KW, N)	0.720***	0.828***	0.276***	0.002	0.814***	0.955***	1.000	
Meta Measure 3 (all)	0.869***	0.714***	0.517***	0.140**	0.407***	0.575***	0.753***	1.000

Note: This table presents Pearson correlation coefficients for the different measures of loss aversion used in this study. See Section 4 in the main text for a brief overview, Section 5 for summary statistics, and Section B.1 in the Appendix for a more detailed account on these measures.

Table 7: Summary Statistics – Colombian Data – Sub-sample facing High Income Risk ($\geq$ 75% Quantile)

	Mean	s.d.	Min	Max	Obs.
Individual Information					
Age	49.3	12.6	24	87	217
Male (=1)	0.28	0.45	0	1	217
Relationship to head of HH					
Head of household (=1)	0.66	0.48	0	1	217
Partner (=1)	0.23	0.42	0	1	217
Son/Daughter or their partner (=1)	0.07	0.26	0	1	217
Other (=1)	0.04	0.19	0	1	217
Household Characteristics					
Number of adult household members	3.0	1.5	1	12	217
Number of adolescents	1.1	1.2	0	6	217
Father still alive (=1)	0.31	0.46	0	1	217
Mother still alive (=1)	0.53	0.50	0	1	217
Exercising					
Every day (=1)	0.19	0.39	0	1	217
At least once a week (=1)	0.16	0.36	0	1	217
At least once a month (=1)	0.06	0.23	0	1	217
Never or hardly ever (=1)	0.60	0.49	0	1	217
Other Health Indicators					
BMI	25.3	4.3	12.9	40.4	217
Education					
Highest year passed	5.4	3.2	0	11	217
Financial literacy score (max. 18)	9.1	3.2	0	15	217
Financial Situation of the Household					
SISBEN Level 2 (=1)	0.50	0.50	0	1	217
Size of safety net (# Persons)	2.8	4.9	0	60	217
Monthly HH income per capita[a]	3.02	2.02	0.13	15.00	217
Market price of house[a]	195.35	387.72	0.00	2000.00	217
Debt[a]	15.34	64.63	0.00	541.71	217
Savings[a]	3.69	19.29	0.00	200.00	217
Engaging in saving (=1)	0.13	0.34	0	1	217
Conditional savings	27.59	46.78	1.00	200.00	29
Planning Horizon					
Day to day (=1)	0.69	0.46	0	1	217
Next months (=1)	0.24	0.42	0	1	217
Next year (=1)	0.05	0.22	0	1	217
Next two to five years (=1)	0.01	0.12	0	1	217
Next five to ten years (=1)	0.01	0.12	0	1	217

Note: [a]Figures reported in 100,000 COP.

Table 8: Summary Statistics of Experimental Measures – Colombian Data – Sub-sample facing High Income Risk ($\geq 75\%$ Quantile)

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.0	0.1	0.0	0.0, 0.0	193
Kahneman-Tversky (KT)	1.2	3.1	0.5	0.1, 1.2	198
Köbberling-Wakker (KW)	11.6	71.6	0.2	0.0, 1.8	217
Neilson (N)	0.2	0.5	0.0	0.0, 0.2	217
Wakker-Tversky (WT)	14.1	155.1	0.1	0.0, 0.3	193
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	1.2	5.2	0.1	0.0, 0.6	217
Meta Measure 2 (KT, KW, N)	1.1	3.5	0.2	0.0, 0.8	198
Meta Measure 3 (all)	0.3	0.9	0.1	0.0, 0.2	178
Impatience					
Near future impatience	28.2	14.9	18.0	16.0, 42.0	217
Increase in patience over time	-0.7	15.7	0.9	-2.9, 0.9	217
Risk Preferences					
Utility Curvature: Gain Domain	6.2	30.3	0.7	0.2, 2.5	217
Utility Curvature: Loss Domain	7.8	16.0	1.1	0.4, 2.9	217
Probability Weighting: Gain Domain	46.4	33.4	46.9	15.6, 78.1	217
Probability Weighting: Loss Domain	68.4	29.4	78.1	46.9, 96.9	217

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail. Near future impatience is the mean annual interest rate, see Section 4. Utility curvature is the parameter of a power utility function and probability weighting is the probability that is perceived as 50%; see Section 4.

Table 9: Summary Statistics – Colombian Data – Loss-Averse Sub-sample (Loss Aversion Meta Measure 1 > 1)

	Mean	s.d.	Min	Max	Obs.
Individual Information					
Age	47.5	14.6	24	77	97
Male (=1)	0.30	0.46	0	1	97
Relationship to head of HH					
Head of household (=1)	0.63	0.49	0	1	97
Partner (=1)	0.23	0.42	0	1	97
Son/Daughter or their partner (=1)	0.07	0.26	0	1	97
Other (=1)	0.07	0.26	0	1	97
Household Characteristics					
Number of adult household members	2.8	1.5	1	8	97
Number of adolescents	1.2	1.2	0	5	97
Father still alive (=1)	0.39	0.49	0	1	97
Mother still alive (=1)	0.49	0.50	0	1	97
Exercising					
Every day (=1)	0.21	0.41	0	1	97
At least once a week (=1)	0.25	0.43	0	1	97
At least once a month (=1)	0.05	0.22	0	1	97
Never or hardly ever (=1)	0.49	0.50	0	1	97
Other Health Indicators					
BMI	25.9	4.3	15.9	39.0	97
Education					
Highest year passed	5.6	3.1	0	11	97
Financial literacy score (max. 18)	9.8	2.9	2	16	97
Financial Situation of the Household					
SISBEN Level 2 (=1)	0.55	0.50	0	1	97
Size of safety net (# Persons)	2.5	3.5	0	30	97
Monthly HH income per capita[a]	3.17	1.99	0.36	10.00	97
Market price of house[a]	203.09	426.03	0.00	2000.00	97
Debt[a]	11.13	52.99	0.00	476.71	97
Savings[a]	2.96	16.22	0.00	150.00	97
Engaging in saving (=1)	0.14	0.35	0	1	97
Conditional savings	20.48	39.46	1.50	150.00	14
Planning Horizon					
Day to day (=1)	0.75	0.43	0	1	97
Next months (=1)	0.16	0.37	0	1	97
Next year (=1)	0.06	0.24	0	1	97
Next two to five years (=1)	0.02	0.14	0	1	97
Next five to ten years (=1)	0.00	0.00	0	0	97

Note: [a]Figures reported in 100,000 COP.

Table 10: Summary Statistics: Income Risk – Colombian Data – Loss-Averse Sub-sample (Loss Aversion Meta Measure 1 > 1)

	Mean	s.d.	Median	Min	Max	Obs.
Local Unemployment Risk (in pc)	25.4	6.5	26.1	15.2	36.5	97

Note: The unemployment risk results from self-reported individual figures in our survey that are averaged at the *UPZ* level; see Section 4 for details.

Table 11: Summary Statistics of Experimental Measures – Colombian Data – Loss-Averse Sub-sample (Loss Aversion Meta Measure 1 > 1)

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.1	0.2	0.0	0.0, 0.1	89
Kahneman-Tversky (KT)	4.4	5.3	2.3	1.6, 4.6	96
Köbberling-Wakker (KW)	68.7	187.2	12.5	4.0, 36.9	97
Neilson (N)	1.0	0.9	0.7	0.5, 1.3	97
Wakker-Tversky (WT)	69.8	271.0	0.4	0.0, 3.3	89
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	6.2	10.5	2.6	1.5, 4.6	97
Meta Measure 2 (KT, KW, N)	4.8	6.5	2.7	1.7, 4.4	96
Meta Measure 3 (all)	1.4	2.0	0.6	0.3, 1.6	88
Impatience					
Near future impatience	31.1	14.9	26.0	16.0, 50.0	97
Increase in patience over time	0.9	17.6	0.9	-2.9, 3.7	97
Risk Preferences					
Utility Curvature: Gain Domain	13.7	19.5	2.7	1.2, 20.3	97
Utility Curvature: Loss Domain	2.1	6.3	0.6	0.4, 1.3	97
Probability Weighting: Gain Domain	61.4	27.8	65.6	46.9, 84.4	97
Probability Weighting: Loss Domain	68.3	27.2	71.9	46.9, 96.9	97

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail. Near future impatience is the mean annual interest rate, see Section 4. Utility curvature is the parameter of a power utility function and probability weighting is the probability that is perceived as 50%; see Section 4.

Table 12: Results from Estimating Model 2 Using a Negative Binomial Hurdle Model and Different Meta-Measures of Loss Aversion – Colombian Data

	Loss Aversion 2 Measures (KW, N)			Loss Aversion 3 Measures (KW,N,KT)		Loss Aversion 5 Measures (All)
	(1)	(2)	(3)	(4)	(5)	(6)
Likelihood of Saving						
Loss Aversion	0.042	0.043	0.030	0.079**	0.063	0.266**
	(1.61)	(1.58)	(1.09)	(1.96)	(1.51)	(2.48)
Income Risk (Survey)	0.027	0.055	0.068**	0.055	0.069**	0.067**
	(1.07)	(1.64)	(2.44)	(1.67)	(2.53)	(2.44)
Amount of Savings						
Loss Aversion	0.062**	0.068***	0.085***	0.092***	0.100***	0.228*
	(2.13)	(2.97)	(3.79)	(3.52)	(3.53)	(1.93)
Income Risk (Survey)	0.145*	0.165***	0.136*	0.166***	0.138*	0.139*
	(2.06)	(3.33)	(1.88)	(3.33)	(1.89)	(1.73)
AIC	1072	967	858	964	857	860
Controls	25	25	25	25	25	25
Region	Yes	Yes	Yes	Yes	Yes	Yes
Occupation	Yes	Yes	Yes	Yes	Yes	Yes
Observations	640	579	509	579	509	509

⋆ p < 0.1, ⋆⋆ p < 0.05, ⋆⋆⋆ p < 0.01. Wild cluster (score) bootstrapped t-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various savings devices; see Section 4. In this Negative Binomial Hurdle model, the participation equation estimates the likelihood to engage in savings, while the second equation estimates conditional savings – the amount of savings, given that a person is saving. Loss aversion is measured by continuous and experimentally elicited meta-measures. The meta-measure comprising two measures of loss aversion is the geometric mean of loss aversion coefficients according to the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005). The measure including three measures is the geometric mean of the former two loss aversion coefficients and, in addition, the one building on the definition of loss aversion by Kahneman and Tversky (1979). Finally, for the last measure, the coefficients based on definitions by Bowman et al. (1999) and Wakker and Tversky (1993) are also included. For more details on the applied measures of loss aversion, see Appendix B.1. Column 2 shows the results when restricting the sample to those for which the meta-measure combining three measures of loss aversion is available, and columns 3 and 5 show the results for similarly restricted samples, in order to be able to draw comparisons between the different meta-measures of loss aversion. We control for variables listed in Tables 3 and 5. Furthermore, we control for regional and occupational sectors at *localidad* level, as well as for the working sectors according to the ISIC classification of economic activities. We account for the cluster structure (at the UPZ level) and potential heteroskedasticity in our data by using wild cluster bootstrapping (Cameron et al., 2008).

Table 13: Results from Estimating Model 3 Using a Negative Binomial Hurdle Model and Different Meta-Measures of Loss Aversion – Colombian Data

	Measure 1 (KW, N)		Measure 2 (KT, KW, N)		Measure 3 (All)	
Likelihood of Saving						
Income Risk (IR)	0.025	(1.02)	0.055*	(1.72)	0.096***	(3.99)
Loss Aversion (LA)	0.032	(1.01)	0.064	(1.42)	0.131	(1.09)
LA × IR	0.012*	(1.70)	0.017**	(1.89)	0.046**	(2.57)
Amount of Savings						
Income Risk (IR)	0.150***	(2.67)	0.167***	(3.30)	0.186***	(2.70)
Loss Aversion (LA)	-0.177***	(-3.57)	-0.129	(-1.36)	-0.169	(-0.76)
LA × IR	0.022***	(5.29)	0.020***	(2.55)	0.047**	(1.96)
AIC	1055		955		853	
Controls	25		25		25	
Region	Yes		Yes		Yes	
Occupation	Yes		Yes		Yes	
Observations	640		579		509	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Wild cluster (score) bootstrapped t-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various savings devices; see Section 4. In this Negative Binomial Hurdle model, the participation equation estimates the likelihood to engage in savings, while the second equation estimates conditional savings – the amount of savings, given that a person is saving. Loss aversion is measured by continuous and experimentally elicited meta-measures. The meta-measure comprising two measures of loss aversion is the geometric mean of loss aversion coefficients according to the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005) (Measure 1). The measure including three measures is the geometric mean of the former two loss aversion coefficients and, in addition, the one building on the definition of loss aversion by Kahneman and Tversky (1979) (Measure 2). Finally, for the last measure, the coefficients based on definitions by Bowman et al. (1999) and Wakker and Tversky (1993) are also included (Measure 3). The coefficients of loss aversion are centered at 1; for more details on the applied measures of loss aversion, see Appendix B.1. Income risk is centered at the mean; see Section 4 for details. We control for variables listed in Tables 3 and 5. Furthermore, we control for regional and occupational sectors at *localidad* level as well as for the working sectors according to the ISIC classification of economic activities, if indicated. We account for the cluster structure (at the UPZ level) and potential heteroskedasticity in our data by using wild cluster bootstrapping (Cameron et al., 2008).

Table 14: Results from Estimating Model 3 Using a Random Effects Panel Regression and Different Meta-Measures of Loss Aversion – Laboratory Experiment

	Loss Aversion 2 Measures (KW, N)		Loss Aversion 3 Measures (KW, N, KT)		Loss Aversion 5 Measures (All)	
		Risk: High IR ($\geq$ 75% Quantile)		Risk: High IR ($\geq$ 75% Quantile)		Risk: High IR ($\geq$ 75% Quantile)
Panel A: Saving						
Income Risk (IR)	0.132***		0.130***		0.129***	
	(10.476)		(10.173)		(7.496)	
High IR (= 1)		3.525***		3.480***		3.511***
		(9.409)		(9.118)		(7.363)
Loss Aversion (LA)	-0.100	-0.464	-0.139	-0.495	-0.867	-1.045
	(-0.336)	(-1.533)	(-0.407)	(-1.423)	(-1.268)	(-1.461)
LA × IR	0.028***		0.027**		0.011	
	(2.775)		(2.252)		(0.445)	
LA × High IR (= 1)		0.851***		0.833**		0.467
		(3.093)		(2.527)		(0.780)
Panel B: Lin. Combination						
Loss Aversion		0.386		0.338		-0.578
		(0.284)		(0.419)		(0.459)
Observations	2562	2562	2562	2562	2562	2562

* p < 0.1, ** p < 0.05, *** p < 0.01.

Note: The dependent variable is the amount saved in a given saving decision task, and for every individual, we have 14 saving decisions, see Section 6.2. Following Xu et al. (2022), we employ a random effects panel regression to account for the multiple (dependent) observations from every individual. Loss aversion is measured by continuous and experimentally elicited meta-measures. The meta-measure comprising two measures of loss aversion is the geometric mean of loss aversion coefficients according to the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005) (Measure 1). The measure including three measures is the geometric mean of the former two loss aversion coefficients and, in addition, the one building on the definition of loss aversion by Kahneman and Tversky (1979) (Measure 2). Finally, for the last measure, the coefficients based on definitions by Bowman et al. (1999) and Wakker and Tversky (1993) are also included (Measure 3). The coefficients of loss aversion are centered at 1; for more details on the applied measures of loss aversion, see Appendix B.1. Income risk is centered at the mean.

Table 15: Summary Statistics of Experimental Measures – Laboratory Experiment

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.2	0.2	0.1	$0.0, 0.3$	183
Kahneman-Tversky (KT)	1.1	0.9	1.0	$0.4, 1.5$	183
Köbberling-Wakker (KW)	1.7	3.7	0.6	$0.1, 1.8$	183
Neilson (N)	0.4	0.4	0.2	$0.1, 0.6$	183
Wakker-Tversky (WT)	1.0	4.0	0.6	$0.3, 0.9$	183
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	0.7	1.0	0.4	$0.1, 0.9$	183
Meta Measure 2 (KT, KW, N)	0.7	0.9	0.5	$0.2, 1.0$	183
Meta Measure 3 (all)	0.5	0.5	0.3	$0.1, 0.7$	183

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail.

E. Instructions – Colombian Lab-in-the-Field Experiment (Translated from Spanish; for Online Publication)

[Instructions for reading to participants]

[Workshop leader] Make sure the person is not distracted by other matters. Read aloud from the script and always be alert to any questions. Be alert to participants' facial expressions to detect lack of understanding of the game.

Good morning. I am ____________ and I would first like to thank you for participating in this study "Saving for old age". The objective of the study is to learn more about the possibilities of saving for old age by Sisben level 1 and 2 households.

The study has two parts. One is the survey that we already did in the past days and the other part is this workshop.

In recognition of your collaboration we will give you a participation reward of 10,000 pesos that regardless of how you do in the game, you will take home safely.

In addition, during this workshop, you will have the possibility to earn more money. Out of 100 participants, 20 will be paid for their decisions in the activity. These 20 people will be selected at random. Once you finish the activity you will draw a ball from a bingo that has 100 balls numbered from 1 to 100. If the ball you draw has a number between 1 and 20, you will be paid for one decision. That decision will also be randomly selected so think very carefully about each decision as any one of them could be selected to be paid. The amount of money you win depends on your decisions as well as luck.

Because you can earn money for your decisions, it is very important that you pay attention to these instructions. In case there is anything you do not understand in the instructions, let me know and I will be happy to answer any questions. To avoid any distractions during the exercise I will ask you to turn off your cell phone.

What do you have to do during the workshop? During the workshop you have to choose between two investment options. In each of the investment options – as in real life – things can go right or they can go wrong. For example, when you start a business, it is possible that the business works out and you make a good profit, or it can be that conditions are unfavorable and you lose.

This happens because there are things that do not depend on one's control, but on chance or other factors.

You will have an endowment of twelve million five hundred thousand (12,500,000) pesos to start with. Depending on your decisions and luck, you can increase or decrease your endowment. This means that you will make each decision thinking that you have twelve million five hundred thousand (12,500,000) pesos in your pocket. In the end, for every 500 pesos you end up having, we will pay you 1 Colombian peso. Twelve million five hundred thousand (12,500,000) experimental pesos then translates into twenty-five thousand (25,000) Colombian pesos, since twelve million five hundred thousand divided by 500 gives twenty-five thousand. From this amount, you can win more money, or lose part of that money so think very well about each decision as any of the questions could be selected to be paid.

In this activity, what you have to do is to decide between two options to invest: Option A and Option B. You choose only one option, A or B, whichever you prefer. The options involve different payoffs that occur under different circumstances. The decisions you make will be kept private. Choose the option you like best and keep in mind that throughout the exercise there are no right or wrong decisions, it all depends on your preferences.

You will find that for each option, A and B there will be a roulette like this [SHOW ROULETTE] which will determine whether things go right or wrong, i.e., it will determine the amount of money you win on each decision. The roulette wheel has a red area and a blue area. If the roulette wheel lands on the red color, you will receive the payout corresponding to the red bar, and if it lands on blue you will receive the payout corresponding to the blue bar. The needle is spun and the color where the needle lands will then determine the payout you will receive.

Throughout the workshop, we will show you figures like this one:

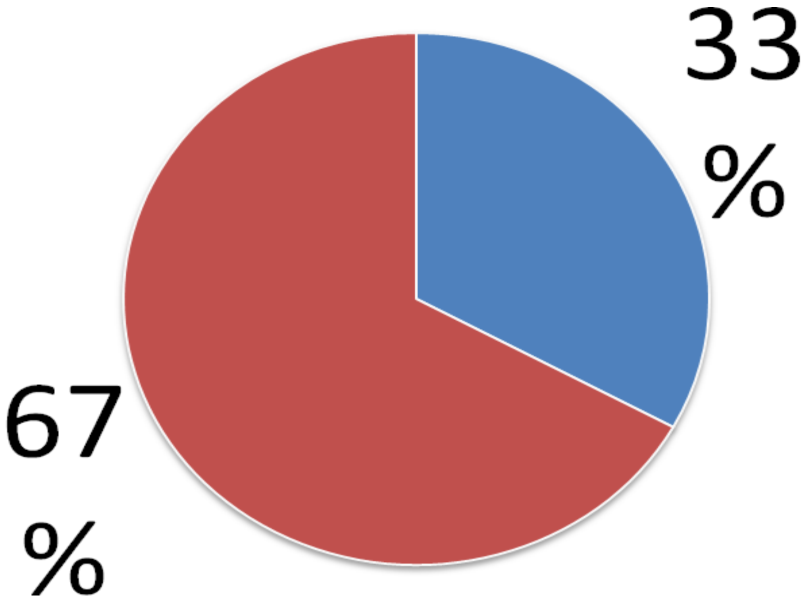

Figure 3: Roulette Wheel

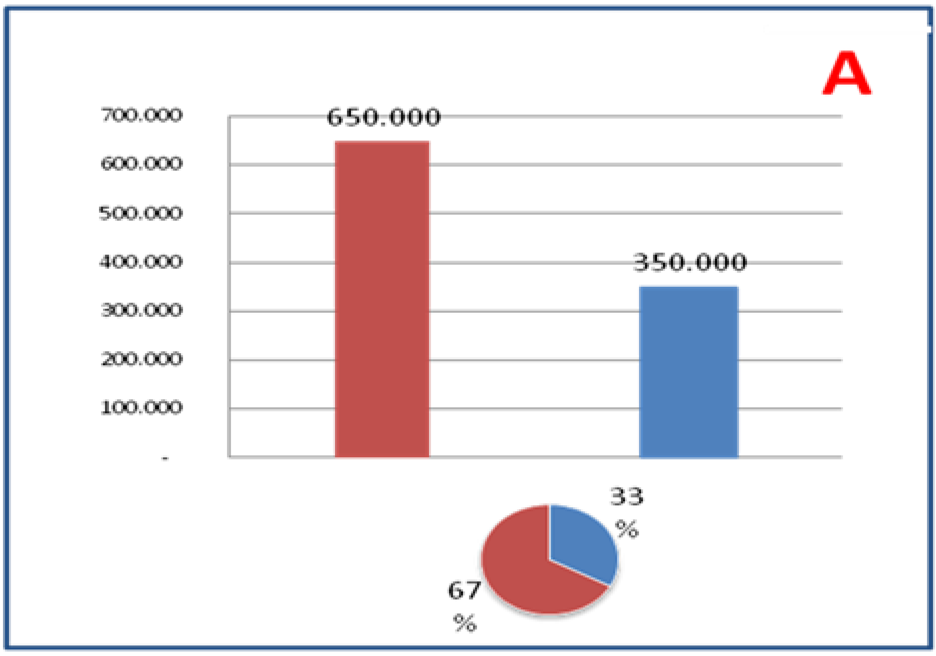

The bars indicate the possible values you will receive. We see two bars, one indicating 650 thousand pesos and the other indicating 350 thousand pesos. Below the bar, a roulette wheel appears, with a red part and a smaller blue part. This indicates that if the needle falls on the red color you receive 650 thousand experimental pesos, but if the roulette falls on the blue color, you receive 350 thousand experimental pesos.

Do you have any questions so far?

[ON THE NEXT PAGE THE CONTROL QUESTIONS START]

E.1. Control Questions 1

Now consider the following investment. [SHOW FIGURE]

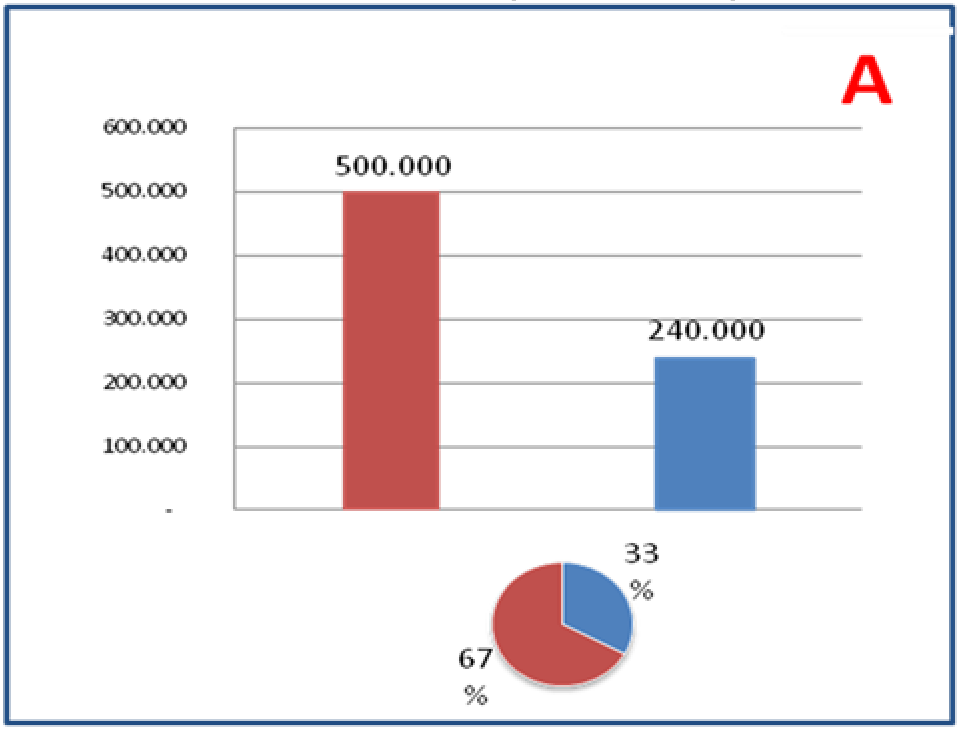

Tell us, in this decision, how much you would receive if things go well $ _________

And how much would you receive if things go wrong? $ ____________

Considering the colors of the roulette wheel, which payment are you most likely to receive? $______

If the roulette wheel falls on the blue color, what payment would you receive? $ ___________

E.2. Example Decision

Now let us look at an example of the decisions we will present in the exercise.

Consider this decision.

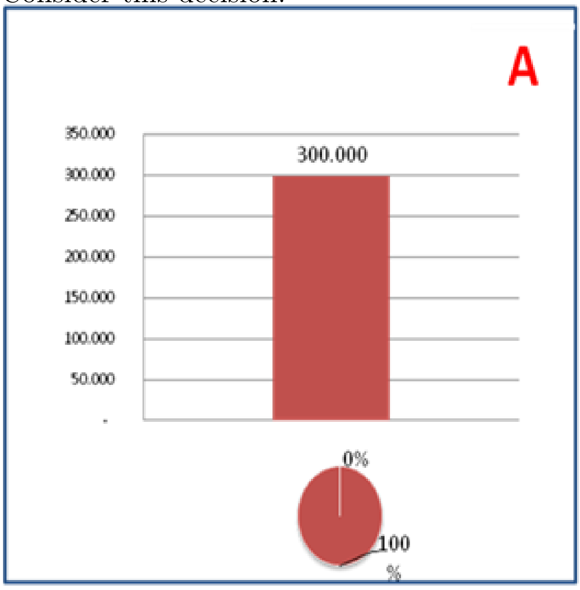

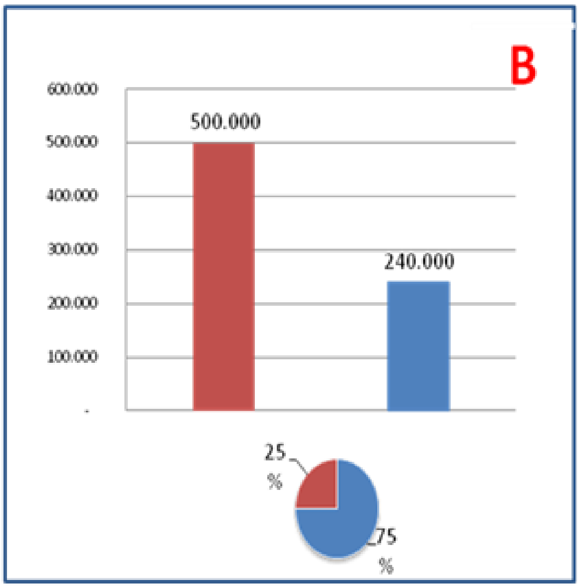

In option A you always win 300,000 pesos, because in this case the roulette wheel is all in red. In option B, if the roulette wheel falls on red you win 500,000 experimental pesos. But if by chance, the roulette wheel lands on blue you would receive 240,000 experimental pesos. The amount of money you win will depend on the area where the needle lands. Your task is to decide which of the two options you prefer, A or B. In this exercise there are no right or wrong decisions, it all depends on what you prefer.

Before we continue, we would like to ask you some questions to verify that you have understood the exercise correctly.

[ON THE NEXT PAGE THE CONTROL QUESTIONS START]

E.3. Control Questions 2

1. Between roulette A and roulette B, where is the needle most likely to land on the blue color?

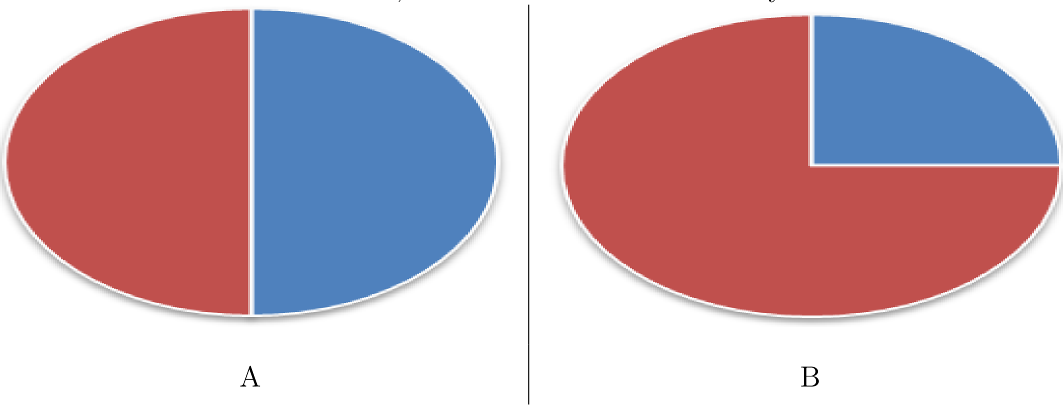

2. In which of these options is there a greater probability of earning more?

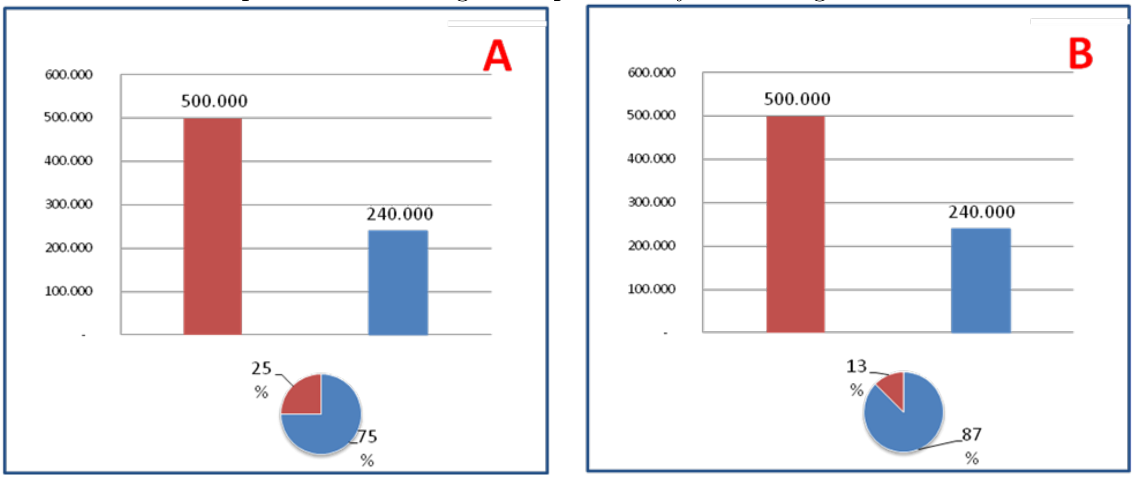

3. And between these two options, when is there a better chance of earning more?

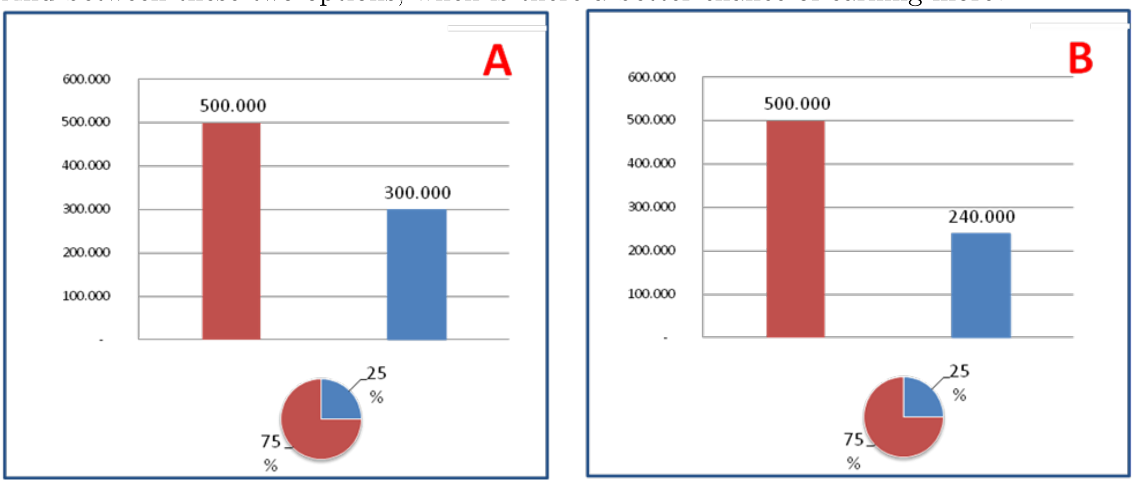

Throughout the workshop you will have to make 100 decisions like this one. Once you finish making the decisions, we will randomly determine which of the 100 decisions is selected to be paid. Therefore, it is important that you make each decision carefully as you will earn the payout for only one of them.

Once the decision to pay is selected, we will proceed to determine whether you receive a high or low payout by spinning the roulette wheel. Depending on your decision and the color the needle lands on, your payout is determined. The payout value is divided by 500, as each 500 experimental pesos represent 1 peso that you are paid. Finally you will choose a bingo ball with balls numbered from 1 to 100. If you draw a ball with a number between 1 and 20, you will receive the 10,000 pesos for having participated, plus or minus the results you obtained in the game. Otherwise you will receive 10,000 pesos for having participated in the workshop.

Finally, keep in mind that although some of the decisions involve the possibility of a negative payout or loss, it will never be more than the twelve million five hundred (12,500,000) endowment you receive at the start of the game. Are you ready to start the game?

[Workshop Leader] Start the game. Do not forget to go back to the instructions in question 106 where the time preference questions start, and hence, new additional instructions are required.

[Workshop leader] Keep in mind when explaining that for each question the person has an endowment of 12,500,000 experimental pesos – especially for the losses.

E.4. Part 2

[Workshop leader] Read it once you get to question 106.

In this part of the activity you will be asked a different set of questions than the ones you just have seen. In this part you will have to decide between two options A and B. However, the options shown do not depend on the roulette wheel as in the previous questions; they give you the option to choose in how much time you would like to receive how much money. That means that you will be shown two different amounts of money that you could receive at two different points in time. Your task is to decide which of the options A or B you prefer depending on the amount of money and the time in which you would receive it. In case one of these questions is selected to be paid, you would receive the 10 thousand pesos for participating today, and you would receive the payment corresponding to the decision **at the time indicated in the question**. For this we would take your data and you would be given a receipt to claim your money at the time indicated in the question. Do you have any questions?

[Workshop leader] Once the game is over, when the question to be paid is selected, explain with this circle the roulette result that will define whether the payout is high or low.

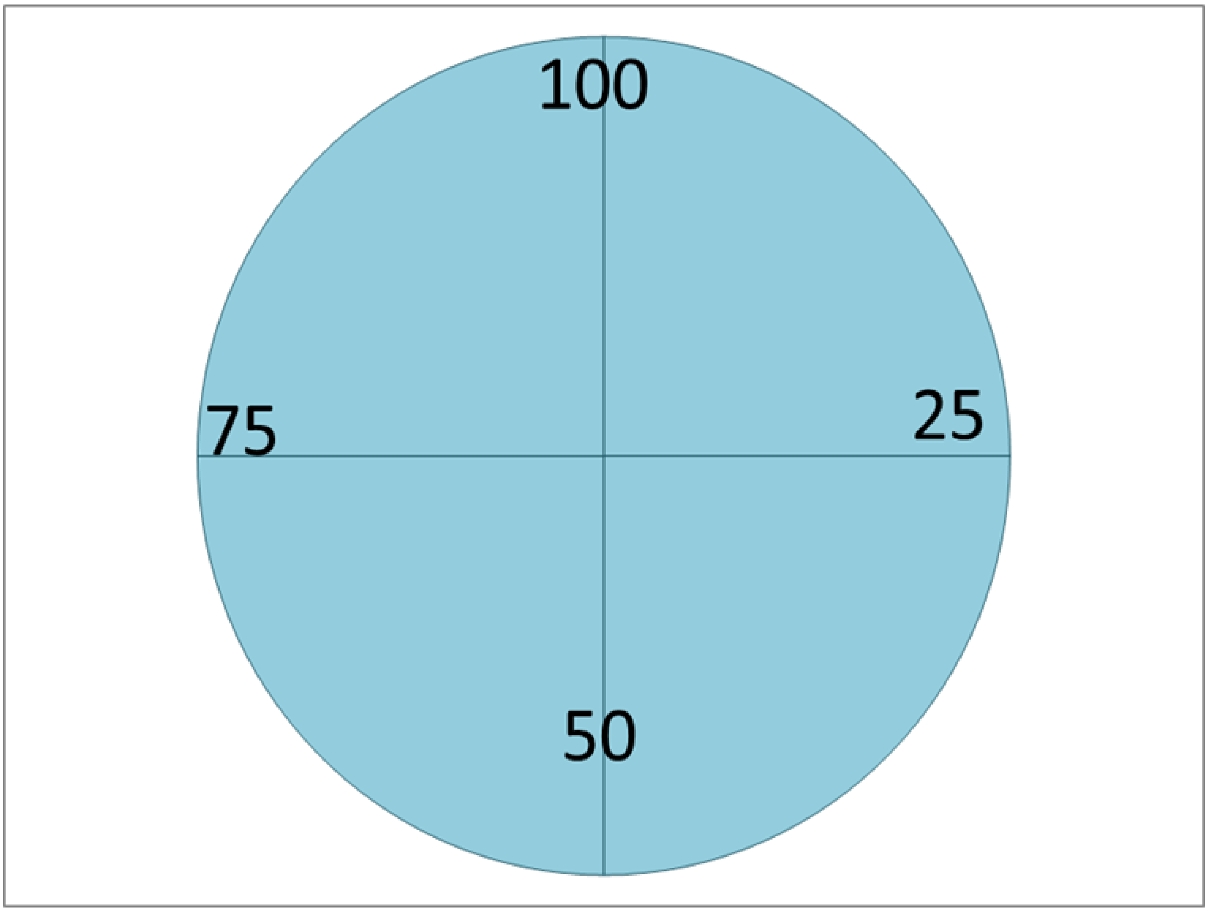
100
75
25
50

F. Instructions – Lab-Experiment (Translated; for Online Publication)

Welcome to the experiment!

Please read these instructions carefully. If you have any questions, please raise your hand.

As you know, you can earn money in this experiment. The exact amount also depends on your decisions.

The experiment consists of two sub-experiments. In each of the two parts, you will have to make a number of money-related decisions. To determine your payoff at the end of the experiment, one part (each with equal probability) is randomly selected as payoff-relevant. From the payoff-relevant part, a decision is again randomly determined, with each decision within that part being equally likely.

Thus, any decision you make could potentially become payout relevant. Since you do not know which one this will be, please process all tasks carefully.

We have also included attention tests to make the data quality measurable for us, but also to ensure fair payoffs. Throughout the experiment we have built in several tests. If your decision in all tests is that of an attentive participant, you get the full payout. If, on the other hand, all tests are answered against all reason, you will receive the minimum payout of EUR 7. If some – but not all – tests imply careful decision making, you will receive the minimum payout plus the difference between the minimum payout and your actual payout, multiplied by the proportion of "correct" test questions. (For example: half of all tests "passed" with an actual payout of EUR 27. Then the difference from the minimum payout is EUR 20, multiplied by 1/2 gives 10, plus EUR 7 minimum payout gives a final payout of EUR 17).

Continue with the instructions for the first part of the experiment, the wheel of fortune game.

[Continue]

F.1. The Wheel-of-Fortune Game

In the wheel-of-fortune game, a wheel of fortune is spun (or the random spinning of a wheel of fortune is simulated by the computer). This determines the payout in this game. Here you can see an example wheel of fortune:

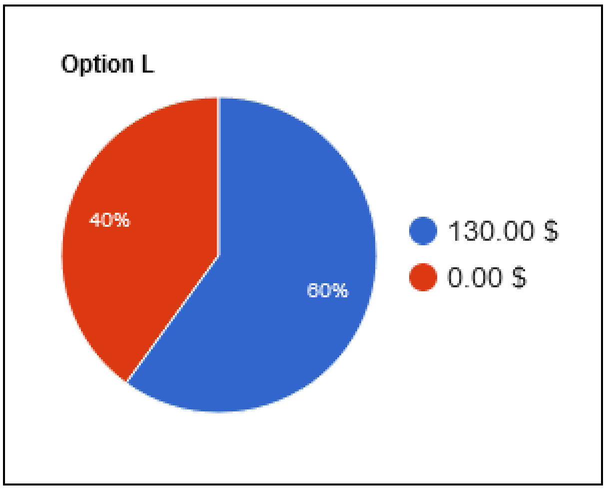

If the wheel of fortune comes to a stop on the **blue piece**, you will receive **$130.00**, as can be seen from the legend to the right of the circle of Option L. This happens with a **probability of 60%** – you can see this number from the corresponding indication on the blue piece. Also, the size of the blue piece takes 60% of the circle.

However, if the wheel of fortune comes to rest on the **red piece**, you will receive **$0.00**. Again, this can be seen in the legend to the right of the circle. This happens with a **probability of 40%** - you can also see this from the indication on the red piece. At the same time, the red area takes up 40% of the circle.

If you want, you can also move the mouse over the circle pieces – you will then be shown the information in a small "text-it".

If the circle consists of only one piece, chance naturally plays no role, because no matter where the wheel of fortune comes to a stop – the amount of money is always the same.

[Continue]

Screen Two In this part of the experiment, we will present you with different wheels of fortune at each step. Your task is to choose the wheel of fortune that you prefer. You can always choose between two options, "Option L" and "Option R".

Here you can see an exemplary decision situation:

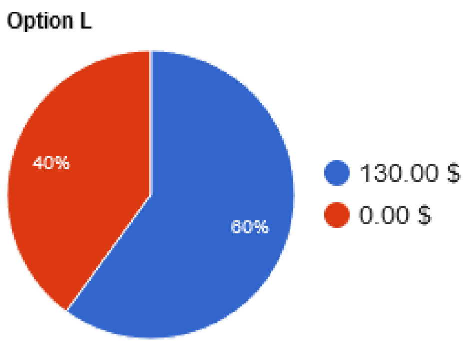

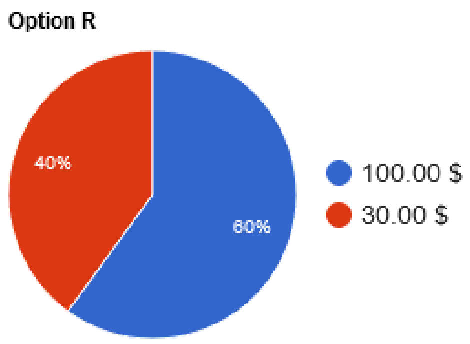

In Wheel of Fortune “Option L” you can get the following payouts:

- $ 130.00 with a probability of 60%, or
- $ 0.00 with a probability of 40%

In the wheel of fortune "Option R" you can get the following payouts:

- $ 100.00 with a probability of 60%, or
- $ 30.00 with a probability of 40%

Your task is to choose your preferred option. To enter your choice, first click the corresponding button under the options (labeled “L” or “R”), and then click “Next”.

In the decision situations, both the pieces of the circle will be different in terms of their size, as well as the amounts paid out. Each decision situation is different.

As mentioned earlier, you should always choose the option you prefer and make each decision as if it will determine your payoff – because it might. Also, some later decisions will depend on earlier decisions; the later decisions will only make sense to you if the earlier ones were made wisely.

The amounts shown here are in experimental dollars ($). They are multiplied by 0.001 in the Wheel of Fortune game in case of payout and this amount is then paid out (rounded to 10 cents) in Euros (that is: $10,000 equals 10 Euros).

For this part of the experiment ("Wheel of Fortune" game) we will give you an endowment of $15,000. Your payouts will be added to this endowment, and the total will be paid out if this sub-experiment is selected for payout. Some decisions in this first part of the experiment also include negative values; this may also reduce your endowment accordingly (in extreme cases it will be used up; however, we will pay out at least EUR 7 to you today – if this game is selected for payout).

If you have no questions and would like to participate in the experiment, you can now click “Continue” to check your understanding and start a short training round for the experiment. If you have any questions, or do not wish to continue participating, please raise your hand.

[Continue]

Comprehension Questions

Here you can see an exemplary decision situation:

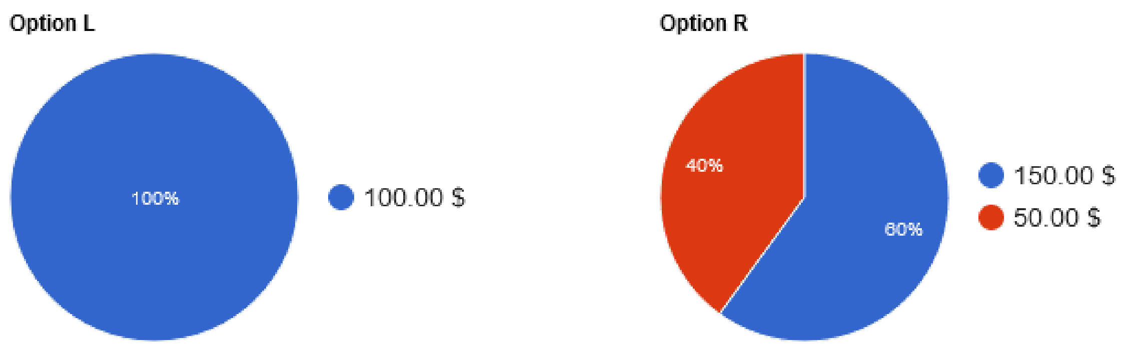

What is the minimum payout in Option L?

○ 0 ○ 50 ○ 100

With what probability (in %) is the payoff equal to the higher amount in option R?

○ 40 ○ 50 ○ 60

[Continue]

Training round

Please choose between option L and option R

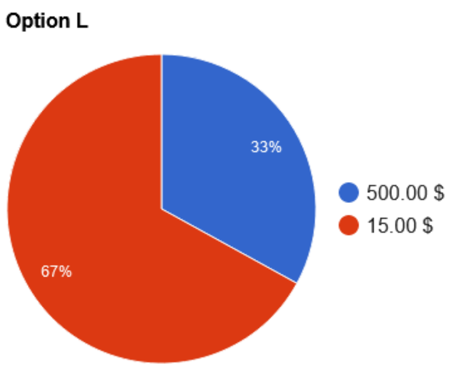

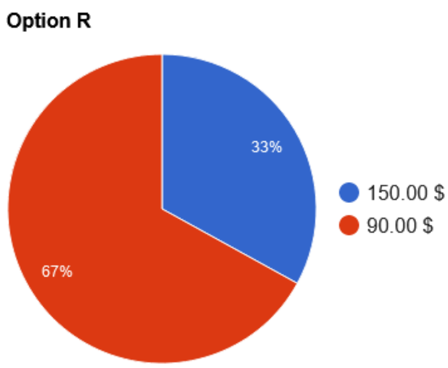

L

R

Take a good look at the two options: a payout of \$150.00 with a probability of 33% means that if you spin 100 times, the wheel of fortune will usually come to a stop on the blue piece around 33 times. The same is true for the other payout. With 100 spins, we would expect an average payout of about \$110.00 = 1/3 * \$150.00 + 2/3 * \$90.00 for Option R.

Alternatively, we could look at the differences in amounts, as the probabilities here are the same in both options (33% for the larger amount in both options). The lowest payoff in this decision situation is in one option – Option L – combined with the highest payoff. While the difference between the low payoffs of both options is \$75.00, it is \$350.00 for the high payoff, in the other direction – that is more than 4 times as much.

However, the probability for the high payout is half as high as for the low payout. In pure mathematical terms ("Better" amount more than 4 times the difference in the lower amounts; probability for the high amount in each case 1/2 as high as for the low amount), you could say that in Option L the maximum amount more than "doubles" the difference in the lower payout compared to Option R (4 * 1/2 = 2).

Note: For negative amounts, the "reverse" calculation applies accordingly (if all amounts in the example had a negative sign, Option R would be "half as bad").

Over the course of a few decisions, one option will always remain constant while the other changes. The amounts are chosen in such a way that the vast majority of participants will switch from one option to the other at some point.

[Continue]

Please choose between option L and option R

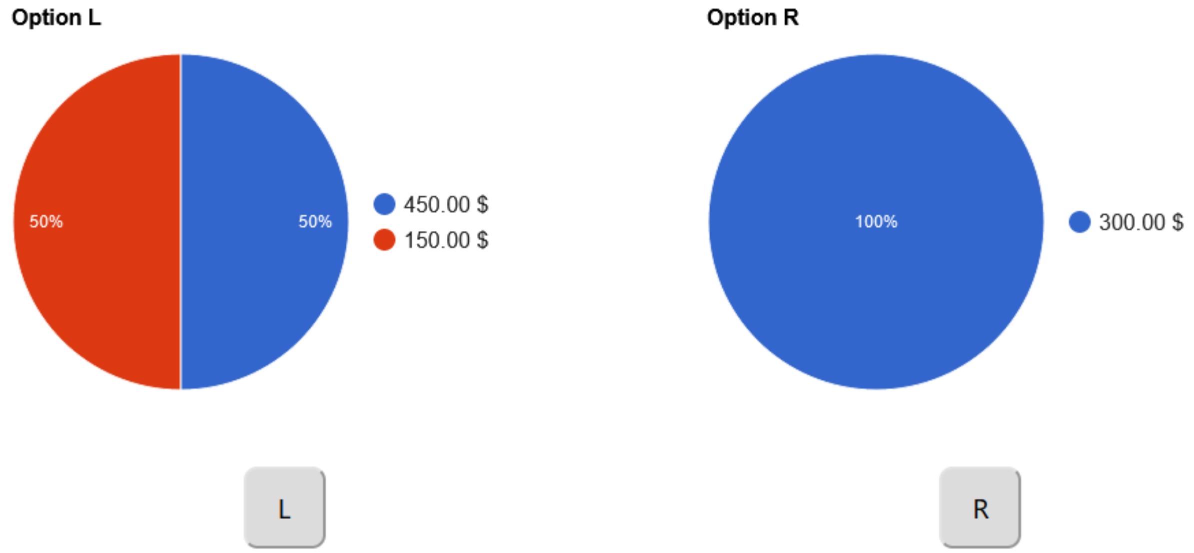

In these two options you can see that one option – Option R – is a "safe choice". It is higher here than the lowest payout of Option L, by $150.00, however there is also no chance of getting the high payout of Option L, which is also $150.00 "away" from the safe payout.

If we spun 100 times, we would get the same amount in Option L and Option R on average.

Again, the amounts are such that the vast majority of participants will eventually switch from one option to the other at least once (possibly more than once here).

F.2. Savings Decisions

In this part of the experiment, we ask for some choices with respect to saving preferences. There are no right or wrong decisions. The explanations in the following paragraph will become clear in the following training round – nevertheless, please read this paragraph as well as the following ones carefully; otherwise you may not understand the training round.

We will show you payouts at two different points in time, where you can "move" money between the earlier point in time ("Time 1") and the later point in time ("Time 2"). Apart from your savings decision, the payout at the earlier point in time is already fixed, while there are two possibilities for the payout at the later point in time, each with equal probability.

In the following decision situations, you can either move money from the earlier to the later point in time, i.e., save, or move from the later to the earlier point in time, i.e., borrow money. In doing so, you can shift a maximum of 20 experimental dollars. Don't worry, as you will see in a moment in the training round, the tasks are very intuitive and you can "try" different options before making your decision.

In this game, 1 Experimental Dollar is equal to EUR 0.80, or 80 cents. Again, the result is paid in Euros rounded to the nearest 10 cents.

Should one of the savings decisions be selected for payment, **chance will** decide whether the payment will **be paid at the earlier or later point in time** – and for the **sake of simplicity, it will be paid today, even if the payment should be selected at the later point in time**. If the payment is selected at the later point in time, chance will additionally decide whether the high or the low amount will be paid out. Accordingly, chance again plays a role in determining your payout.

Note: We did not include attention tests in this part of the experiment. However, the amounts are quite high – so choose carefully.

If you have no further questions, you can proceed to the training round by clicking "Continue".

[Continue]

Training Round 1

Screen 1

Click on the gray bar to make the slider visible.

Screen 2 (after having clicked on the slider bar)

Date 1: $ 38	Date 2: $ 36 or $ 36
With a 100% probability	Each with a 50% probability

Your choice: you save $0

The slider is set to $0 [VALUE IS ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER]. A positive amount means that you are moving money from the earlier point in time (Date 1) to the later point in time (Date 2); thus, you are saving money that will be subtracted from the payout at date 1 and added to the payouts at date 2. Conversely, a negative amount means that you are moving money from the later point in time to the earlier point in time; thus, you are borrowing money in date 1 from date 2.

Take a look at the payout options with this setting at both times:

- At the earlier point in time, “Date 1”, there is only one payoff option, since the amount listed under “Date 1” at the top left, $38 [VALUE IS ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER], is paid out with a probability of 100% (if Date 1 were randomly selected for today’s payoff in this decision situation).
- At the later time, “Date 2” there are always two payout options: The two amounts mentioned under “Date 2” on the top right, $36 and $36 [VALUES ARE ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER]. They will both be paid out with a probability of 50% each, should date 2 be randomly selected for today’s payout in this decision situation. In this decision situation, these amounts are identical, therefore there is practically only one amount at date 2; but this can change in the following situations.

If you haven’t tried it yet: Move the slider and watch how the amounts change to the earlier and later point in time – which choice do you like best?

[Continue]

Training Round 2

Screen 1

Click on the gray bar to make the slider visible.

Screen 2 (after having clicked on the slider bar)

Date 1: $ 38 | Date 2: $ 28 or $ 44

With a 100% probability | Each with a 50% probability

Your choice: you save $0

The slider is set to $0 [VALUE IS ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER].

Now take a good look again at the payout possibilities in this setting at both times.

- You can see that the amounts at time 2 (i.e., the amounts mentioned under date 2 above right) are now different. Both are selected for payment with a probability of 50% – should this decision situation as well as date 2 be selected for payment at random.

- The amount at time 1 (i.e., the amount mentioned under time 1 above left) will be selected for payment unchanged with a probability of 100% – should this decision situation as well as date 1 be randomly selected for payment.

Which choice do you like best here?

A. Theoretical Framework: Details (For Online Publication)

A.1. General Version of the Two-period Model by Kőszegi and Rabin (2009)

As in Section 2, we assume that an individual has to distribute wealth, W, for consumption across two periods such that $W = c_1 + c_2$, where c_t denotes consumption in period t for $t = 1, 2$. As in the main text, we consider the case in which wealth is stochastic and uncertainty is resolved in the second period.

Consumption in the first period (and thus saving) is determined by maximizing the expectation of the sum of instantaneous utilities u_t in both periods, where no discounting is assumed, i.e.,

$$U = u_1(c_1) + \mathbb{E}[u_2(c_2)]. \tag{3}$$

As in the simplified version of the model introduced in the main text, individuals are assumed to choose their favorite credible consumption plan before the first period starts (i.e., in period $t = 0$). Credible means that they anticipate whether or not they would be able to stick to the plan, and only consider those plans where they do not see an incentive to deviate from later on.[37] Favourite means that there are possibly several such credible plans, and the decision-maker chooses his or her preferred one according to the maximization principle. This plan is called preferred personal equilibrium (PPE) by Kőszegi and Rabin (2009) and at the time of planning in period $t = 0$, it leads to possibly stochastic 'rational beliefs' $F_{0,1}$ and $F_{0,2}$ about consumption in Period 1 and Period 2. Mathematically, these beliefs are simply probability distributions assigning a probability to any possible consumption level. Plans about consumption in period t that are made in the same period (i.e., $F_{t,t}$) assign a probability of 1 to the actual consumption level c_t. When uncertainty is resolved and consumption decisions are implemented, plans are updated and lead to new beliefs.

Instantaneous utility in periods $t = 1, 2$ is given by

$$u_t = m(c_t) + \sum_{\tau=t}^{2} \gamma_{t,\tau} N(F_{t,\tau}|F_{t-1,\tau}),$$

where $m(\cdot)$ is consumption utility that is three times differentiable, increasing and strictly concave, and corresponds to a "classical utility function". The 'gain-loss utility', $N(F_{t,\tau}|F_{t-1,\tau})$, reflects utility gains or losses due to changes in current 'beliefs' $F_{t,\tau}$ compared to former 'beliefs' $F_{t-1,\tau}$ about contemporaneous ($\tau = t$) and future ($\tau > t$) consumption. Depending on the

[37] Details about how these plans are formed are given in Appendix A.2 or in Kőszegi and Rabin (2009).

distance of a period $\tau \geq t$ in the future, the impact of changes in beliefs about consumption in that period via the 'gain-loss utility' differs, which is reflected by weights $\gamma_{t,\tau} \geq 0$ with $\gamma_{t,t} = 1$. For simplicity, we use the notation $\gamma_{1,2} = \gamma$. The weight $\gamma_{1,2} = \gamma$ is decisive for an individual to adhere to her plan, i.e., to resist overconsuming in the first period relative to the previously set consumption level, as explained below.

'Gain-loss utility' N compares every percentile of the distributions of consumption according to 'beliefs' $F_{t,\tau}$ and $F_{t-1,\tau}$, using a "universal gain-loss utility function" μ. More specifically, for a possibly discrete distribution F_d, $c_{F_d}(p/100)$ is a percentile for $0 \leq p \leq 100$ with $p \in \mathbb{N}$ if $F_d(c_{F_d}(p/100)) \geq p/100$ and $F_d(c) < p/100$ for all $c < c_{F_d}(p/100)$. Then, gain-loss utility from the change in beliefs from $F_{t-1,\tau}$ to $F_{t,\tau}$ is defined as

$$N(F_{t,\tau}|F_{t-1,\tau}) = \sum_{p=1}^{100} \mu(c_{F_{t,\tau}}(p/100), c_{F_{t-1,\tau}}(p/100)),$$

where

$$\mu(\hat{c}, \tilde{c}) = \begin{cases} \eta(m(\hat{c}) - m(\tilde{c})) & \text{if } \hat{c} \geq \tilde{c} \\ -\lambda\eta(m(\tilde{c}) - m(\hat{c})) & \text{if } \hat{c} < \tilde{c}. \end{cases}$$

for two consumption levels $\hat{c}$ and $\tilde{c}$, m as defined above and parameters $\eta > 0$ and $\lambda > 0$.[38]

The parameter $\eta > 0$ simply scales the difference in consumption utility, and $\lambda > 0$ may account for loss-averse ($\lambda > 1$) or gain-seeking ($\lambda < 1$) behavior.

The parameter $\gamma \geq 0$ 'discounts' anticipated future gains or losses in 'gain-loss' utility that affect utility already in period 1. For $\gamma > 1/\lambda$, the anticipated future loss is weighted high enough to prevent the consumer from deviating from the optimal ex-ante plan, i.e., they resist overconsuming; see Proposition 5 in Kőszegi and Rabin (2009). When $\lambda > 1$, following Kőszegi and Rabin (2009), we can assume $\gamma < 1$. As we allow for gain-seeking behavior, i.e., $\lambda < 1$, we leave γ unrestricted, to allow for $\gamma > 1/\lambda$. Then, the proof of Proposition 8 in Kőszegi and Rabin (2009) holds for $\lambda < 1$, although they do not consider this case.

If the agent resists deviating from the plan, instantaneous utility in Period 1 is given by

$$u_1 = m(c_1) + N(F_{1,1}|F_{0,1}) + \gamma N(F_{1,2}|F_{0,2}) = m(c_1),$$

[38]This choice of the "gain-loss utility function" fulfills certain desirable characteristics of a reference-dependent utility function for $\lambda > 1$; see Kőszegi and Rabin (2009), p. 914. In particular, it fulfills "the explicit or implicit assumptions" about the 'value function' by Kahneman and Tversky (1979), as formulated by Bowman et al. (1999).

as beliefs do not change in the first period (i.e., $F_{0,t} = F_{1,t}$ for $t = 1, 2$), since in addition to adherence to the plan, no uncertainty is resolved. In Period 2, utility is given by

$$u_2 = m(c_2) + N(F_{2,2}|F_{1,2}).$$

With that, the optimization problem can be solved by equalizing the marginal utility of saving and consumption in the first period.

If the agent cannot resist deviating from the ex-ante optimal plan, their PPE specifies a higher consumption level in Period 1 compared to the optimal one; see Proposition 5 in Kőszegi and Rabin (2009).

A.2. Rational Beliefs

In this Appendix, we explain the intuition behind 'rational beliefs'. We refer to Kőszegi and Rabin (2009) for a precise definition.

'Beliefs' are the result of a plan: They "must be rationally based on credible plans for state-contingent behavior".[39] One concept of what a credible plan could be was termed 'preferred personal equilibrium (PPE)' by Kőszegi and Rabin (2009) and was used in their text, although they note that other theories of forming beliefs could also be combined with their model. Roughly speaking, a plan is a PPE if it is the preferred "plan among those that are credible". A plan is credible if it maximizes the mathematical expectation of the reference-dependent utility in every period given the beliefs which the plan induced *and* if continuation plans are consistent. That is: If an individual plans for very low consumption in Period 1 in order to save for Period 2, but would not make the same choice if solving the maximization problem in Period 1 – e.g., because they are present-biased or cannot live with such a low level of consumption –, this would not be a credible plan, and it is not a PPE. Using backwards induction, they would anticipate their behavior in Period 1 and consume more in Period 1 from the beginning until their entire plan is consistent with solutions evolving from a similar maximization process in Period 1. This PPE reflects the idea that individuals anticipate the implications of their plans and only make plans they know they would adhere to them.

[39] The most simple example of a state-contingent plan could be: "If things go well, I will spend x\$ for consumption in Period 1. If things do not work out well, I will only spent y\$ in this period" (where $x > y > 0$).

A.3. Proofs

Proof of Proposition 1. This proof follows the rationale of the proof of Proposition 8 in Kőszegi and Rabin (2009).

We prove that the derivative of the marginal utility of increasing savings with respect to λ is positive. Equivalent to the argument in the proof of Kőszegi and Rabin's Proposition 8, this implies that $dc_1/d\lambda < 0$ for both $\gamma > 1/\lambda$ and $\gamma \leq 1/\lambda$, since in the first case, the ex-ante optimal plan involves a lower c_1 and the person adheres to this plan. In the latter case, a higher marginal utility in Period 2 makes a lower c_1 become consistent. Furthermore, since, for $\gamma \leq 1/\lambda$, the chosen c_1 will be higher than for $\gamma > 1/\lambda$, see Kőszegi and Rabin (2009), a lower c_1 will become consistent, as the agent adheres to the ex-ante optimal plan for a lower γ.

The derivation of marginal utility of increasing savings is due to Kőszegi and Rabin (2009): Let F be the cumulative distribution function of the (mean-zero) random variable y. The expected utility in Period 2 is

$$\begin{aligned}
&\int m(c_2 + sy)\, dF(y) + \iint \mu(m(c_2 + sy) - m(c_2 + sy'))\, dF(y')\, dF(y) \\
&= \int m(c_2 + sy)\, dF(y) \\
&\qquad - \frac{1}{2}\eta(\lambda - 1) \iint m(c_2 + s\max\{y, y'\}) - m(c_2 + s\min\{y, y'\})\, dF(y')\, dF(y).
\end{aligned}$$

Hence, the derivative of the expected utility in Period 2 with respect to c_2, i.e., the marginal utility from increasing savings is

$$\begin{aligned}
&\int m'(c_2 + sy)\, dF(y) \\
&\qquad + \frac{1}{2}\eta(\lambda - 1) \iint m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\})\, dF(y')\, dF(y).
\end{aligned}$$

Now, unlike in the proof or Proposition 8 in Kőszegi and Rabin (2009), we take the derivative of the expression above with respect to λ:

$$\frac{1}{2}\eta \iint m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\})\, dF(y')\, dF(y).$$

This derivative is positive for any strictly concave m, any $s > 0$, $\eta > 0$, and any non-degenerate random variable y. Thus, the marginal utility from increasing savings is an increasing function of λ. □

Proof of Corollary 1. As in the proof of Proposition 1, the derivative of the marginal utility from increasing savings with respect to λ is given by

$$\frac{1}{2}\eta \iint m'(c_2 + s\min\{y, y'\}) - m'(c_2 + s\max\{y, y'\})\, dF(y')\, dF(y).$$

The derivative of this expression with respect to s evaluated at $s = 0$ is

$$\frac{1}{2}\eta(-m''(c_2)) \iint |y' - y|\, dF(y')\, dF(y),$$

which is positive for any strictly concave consumption utility function m, $\eta > 0$ and any non-degenerate random variable y. □

B. Data: Details (For Online Publication)

B.1. Details on the Measures of Loss Aversion

In this section we describe how we operationalized the different measures of loss aversion with our data, following Abdellaoui et al. (2007).

Kahneman-Tversky (KT) Kahneman and Tversky (1979) define an individual as loss-averse, if for all amounts of money x the utility μ of receiving this amount is lower than the disutility of losing that same amount, i.e., if $\forall\ x > 0 : -\mu(-x) > \mu(x)$. A natural coefficient of loss aversion emerging from this definition is $-\mu(-x)/\mu(x)$ for every elicited amount $x > 0$. If $\mu(-x)$ for any of these eight elicited amounts of money $x > 0$ was not elicited, it was linearly interpolated. As the coefficient of loss aversion, we took the median of the computed coefficients.

Neilson (N) Neilson (2002) proposes computing the ratio of 'relative steepness', which is the utility value $\mu(x)$ divided by the corresponding x-value. This figure incorporates information about steep parts of the utility function at any point of the interval of interest – even in flat regions. If the relative steepness of the utility function over the loss domain is bigger than the one on the gain domain at any point, the individual is classified as loss averse, i.e., $\mu(-x)/x \geq \mu(y)/y$, $\forall\ x, y > 0$. For this definition, we computed the coefficient of loss aversion as the ratio of the infinum of $\mu(-x)/(-x)$ over the supremum of $\mu(y)/y$.

The remaining definitions rely on the steepness of the utility function as expressed by the derivative of the latter on both domains.

Wakker-Tversky (WT) Wakker and Tversky (1993) suggest applying the concept of Kahneman and Tversky (1979) to the derivative of utility, i.e., to compare the value of the derivative of the utility function for gains and losses 'point-wise' at certain absolute values: $\mu'(-x) > \mu'(x), \forall\, x > 0$. At every elicited utility point $x > 0$ on the gain domain, the derivative $\mu'(x)$ was operationalized as the mean of the two connecting slopes to the left-hand side and to the right-hand side. $\mu'(-x)$ was operationalized as the slope of the linearly interpolated utility function at the point $-x$. Similar to the case for the definition by Kahneman and Tversky (1979), a natural coefficient emerging from the definition $\mu'(-x) > \mu'(x), \forall\, x > 0$, is $\mu'(-x)/\mu'(x)$ for $x > 0$. In this case, we also took the median of the coefficients thus computed.

Bowman (B) Bowman et al. (1999) propose performing this comparison 'domain-wise', that is, $\mu'(-x) > \mu'(y), \forall\, x, y > 0$. As in the case for the definition by Neilson (2002), the definition $\mu'(-x) > \mu'(y),\ \forall\, x, y > 0$ can be transformed into a coefficient of loss aversion by computing $\inf \mu'(-x)/\sup \mu'(y)$ for $x, y > 0$, where the derivatives where operationalized as just described.

Köbberling-Wakker (KW) Finally, Köbberling and Wakker (2005) define an individual as loss-averse if the slope of the utility function on the left-hand side of the reference point is steeper than the slope of the utility function on the right-hand side of the reference point: $\mu'(0_-) > \mu'(0_+)$. The natural coefficient of loss aversion resulting from this definition, $\mu'(0_-)/\mu'(0_+)$, was computed as the ratio of slopes connecting 0 with the elicited utility points that are closest to 0 on both domains.

B.2. Parametric Estimation of a Power Utility Function

General Form for Positive Arguments Usually, the power family is defined for $x > 0$ by

$$m(x) = \begin{cases} x^b & \text{for } b > 0 \\ \ln(x) & \text{for } b = 0 \\ -x^b & \text{for } b < 0. \end{cases}$$

Considering Non-Positive Arguments Since $\ln(x)$ is not defined for $x < 0$, the case $b = 0$ must be excluded, if negative arguments are of interest. Furthermore, $b < 0$ has to be excluded as well, if the point $x = 0$ is to be considered.[40] Thus, when allowing for gains and losses, the *power family* reduces to

$$m(x) = \begin{cases} -(-x)^a & \text{for } a > 0,\ x < 0 \\ x^b & \text{for } b > 0,\ x \geq 0. \end{cases}$$

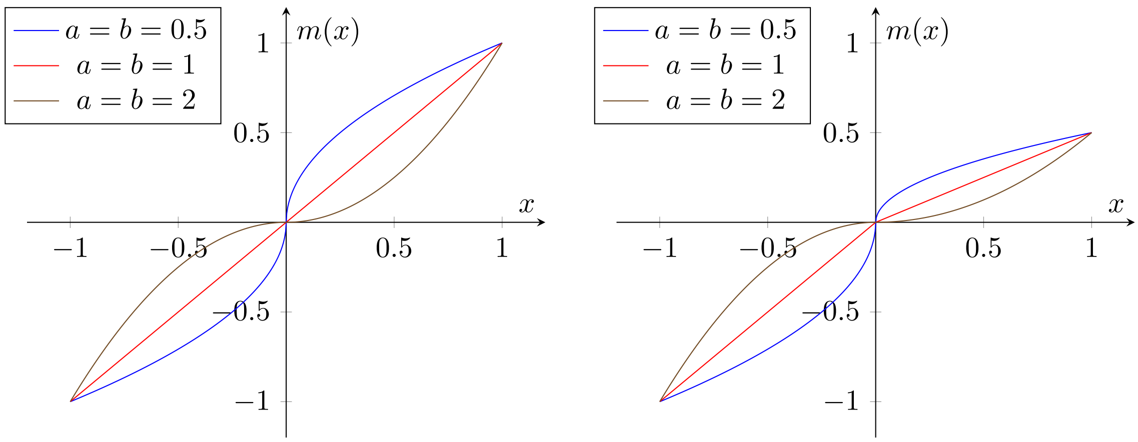

(a) Curvature of the power family for different values of a and b.

(b) Estimated power utility functions plotted for different values of a and b.

Figure 2: Illustration of the Power Family Utility Function with Different Values of a and b

Figure 2(a) illustrates the curvature of the power family for different values of a and b.

Rescaling Arguments Arguments x of the utility function must be rescaled in order to lie within the interval $[-1, 1]$ for all the subjects in the study in order to be able to compare estimated parameters.

Due to the method used, the minimal x-value observed is $L_1 = -5{,}000{,}000$. Thus, for losses, we need a transformation $x \mapsto -\frac{x}{L_1}$, where $x \in [L_1, 0]$.

For Gains, $G_{0.5}$ is the maximum x-value for any individual, we therefore transform $x \mapsto \frac{x}{G_{0.5}}$, where $x \in [0, G_{0.5}]$.

Rescaling Outputs By the method chosen, we need to have $m(L_1) = -1$, $m(0) = 0$ and $m(G_{.5}) = .5$. We check this: For the negative domain, we have

$$m(L_1) = -\left(\frac{L_1}{L_1}\right)^a = -(1)^a = -1,$$

[40] Wakker (2008, p.1336) gives a less technical explanation: "With both positive and negative x present, a negative power a or b generates an infinite distance between gains and losses. Such a phenomenon is not empirically plausible, so that negative a and b should then not be expected to occur."

independent of $a > 0$, so there is no need to rescale outputs. However, for the positive domain,

$$m(G_{0.5}) = \left(\frac{G_{0.5}}{G_{0.5}}\right)^b = 1^b = 1,$$

independent of $b > 0$. Therefore, and also to have estimates comparable for the negative and the positive domain, we rescale $m(x)$ for $x \geq 0$ and set:

$$m(x) = 0.5 \cdot \left(\frac{x}{G_{0.5}}\right)^b \qquad \text{for} \qquad x \geq 0.$$

Note that we could also leave the estimation formula untouched and multiply our outcomes by the factor 2, making them lie within the interval $[0, 1]$ instead of $[0, .5]$.

Estimation Equation The final estimation equation is thus

$$m(x) = \begin{cases} -\left(\frac{x}{L_1}\right)^a & \text{for } a > 0,\ x < 0 \\ 0.5 \cdot \left(\frac{x}{G_{0.5}}\right)^b & \text{for } b > 0,\ x \geq 0. \end{cases}$$

This equation is illustrated in Figure 2(b).

Curvature In order to classify a utility function as convex or concave based on the estimated values of the parameters a or b, we can deduct the curvature of the utility function from Figure 2 for the given values of a and b. Analytically, for classifying an individual's utility function, we calculate the second derivative of the estimated utility function.

$$m''(x) = \begin{cases} -\left(\frac{x}{L_1}\right)^a \cdot \frac{1}{x^2} \cdot a(a-1) & \text{for } a > 0,\ x < 0 \\ 0.5 \cdot \left(\frac{x}{G_{0.5}}\right)^b \cdot \frac{1}{x^2} \cdot b(b-1) & \text{for } b > 0,\ x > 0, \end{cases}$$

where $x = 0$ has to be excluded from the domain.

We immediately see that for $x > 0$,

$$m''(x) \begin{cases} < 0 \text{ thus } m \text{ strictly concave} & \text{if } 0 < b < 1 \\ = 0 \text{ thus } m \text{ linear} & \text{if } b = 1 \\ > 0 \text{ thus } m \text{ strictly convex} & \text{if } b > 1, \end{cases}$$

and for $x < 0$ we have

$$m''(x) \begin{cases} < 0 \text{ thus } m \text{ strictly concave} & \text{if } a > 1 \\ = 0 \text{ thus } m \text{ linear} & \text{if } a = 1 \\ > 0 \text{ thus } m \text{ strictly convex} & \text{if } 0 < a < 1. \end{cases}$$

C. Results: Details (For Online Publication)

C.1. Econometric Model

The outcome variable used in our analysis – savings (in 100,000 COP) – does not include negative values and is therefore a limited dependent variable according to the definition in Wooldridge (2013, Chapter 17). Furthermore, the empirical frequency of zeros in the distribution of the amount of savings in our sample exceeds the frequency of zeros according to any commonly used theoretical distribution in such cases (e.g., the Poisson distribution or the Negative Binomial distribution). This is to be expected, since not everybody actually engages in saving. Thus, the outcome variable is a so-called Corner Solution Response.[41]

The distribution of the value of saving in our sample is skewed, and values are reported repeatedly and are usually divisible by 100,000 COP. Therefore, we should assume a discrete rather than a continuous dependent variable. Given these characteristics of the outcome variable, we apply a Negative Binomial Hurdle model to study the relationship between income risk, loss aversion, and savings. The Poisson Hurdle model is nested in the Negative Binomial Hurdle model we fit and differences between the log-likelihoods of both models mostly exceed 100 by far. This indicates that a likelihood ratio test (conservatively assuming the test statistic to follow a chi-square distribution with one degree of freedom) would reject the hypothesis of no overdispersion.

This model is a so-called two-part model, where the probability of engaging in savings and the amount of savings is estimated separately by different models. For the Hurdle models applied here, the likelihood of both equations can be calculated separately. Using a logit-model, the probability 'that the hurdle is passed' and that a person engages in savings is estimated. The second model estimates the amount of savings once the hurdle is passed, using a Truncated Negative Binomial model. In Appendix C.2, we discuss alternative models and their suitability in this context.

[41] The options to deny the response or to indicate that they did not know about the amount of savings were allowed and treated separately. Four respondents denied answering and five respondents did not know the amount of savings they held at the time of the interview. Together, this corresponds to about 1% of the respondents whose savings amount we could not observe. These cases were excluded from the analysis.

Following Grogger and Carson (1991), we compute marginal effects of loss aversion and income risk on the predicted amount of unconditional savings using the estimates resulting from fitting Model 3 with a Negative Binomial Hurdle model. Denoting savings for individual i with Y_i, the overall marginal effect of X_{ih}, i.e., of covariate h for individual i, on his or her predicted savings can be computed as

$$\frac{\partial \mathbb{E}(Y_i|X_i)}{\partial X_{ih}} = \frac{\partial}{\partial X_{ih}}[\mathbb{E}(Y_i|X_i, Y_i > 0)][1 - F(0)] + \mathbb{E}(Y_i|X_i, Y_i > 0)\frac{\partial}{\partial X_{ih}}[1 - F(0)], \qquad (4)$$

where $1 - F(0)$ is the share of the population for which we observe $Y_i > 0$. This means the overall effect can be decomposed into two effects: The effect on those who are saving, weighted by the probability of saving, plus the effect on the proportion that 'passes the hurdle' and is saving, weighted by the mean amount of savings in the saving population. We compute marginal effects using mean values of covariates, unless otherwise indicated.

C.2. Discussion: Model Choice

In this part, we briefly discuss alternatives to the model chosen and assess their appropriateness in the setting of this paper.

Usually, OLS regression is a suitable starting point for modelling empirical relationships. However, a large share of the non-savers with zero COP of savings could mask relationships observed for the fraction of participants that actually saves. It seems appropriate to take the large share of the non-savers observed in our data into account when selecting a suitable model.

A Tobit model is frequently used in similar situations. Here, it is not suitable. A central assumption of the Tobit model is that the process determining participation is the same as the process determining the amount of saving. The signs of the coefficients of the independent variables in Table 1 differ in the two equations where many are significantly different from zero, showing that this assumption is violated. Second, normality and homoscedasticity of the dependent variable model are prerequisites for using a Tobit model. In contrast to OLS, where departures from these assumptions still lead to unbiased and consistent estimates, it is less clear how sensitive the Tobit model is to departures from these assumptions. The empirical distribution of the outcome variable we observe in our data is discrete. This observed empirical distribution is a rather bad approximation of any continuous probability distribution, so the assumption of normality is not likely to hold.

More flexible models for corner solution responses that can model the participation process and the savings process separately are – in addition to the Hurdle model applied in this study – so-called inflated models. For example, the Zero-Inflated Poisson model or the Zero-Inflated Negative Binomial model for the case of a discrete dependent variable.

Zero-inflated models rely on the assumption that a zero COP value of savings can be the result of two cases: In the first case, an individual would decide to save and then chooses a saving amount of zero. In the second case, an individual would decide not to save at all. We believe that the first case is rather unrealistic, since we did not ask for changes in savings in a given limited time, but rather look at the stock of savings. We therefore conclude that these models are not appropriate in our setting.

It is noteworthy that the excess zeros in the distribution of the outcome variable are not a problem of data observability, where models for censored data or sample correction models (e.g., the Heckman model) would be adequate. Only for around 1 percent of the participants are data actually missing, and these cases were excluded.

When only focusing on the positive amount of savings, no special care is needed to account for excessive zeros in the distribution of the outcome variable. In such cases, a traditional OLS model could be applied, or a log OLS model, if we expect the relationship to be proportional to the response.

Given the discrete character of the outcome variable, and its heavily non-symmetric empirical distribution, a model that accounts for this characteristic should be applied, such as the Zero-Truncated Poisson or the Zero-Truncated Negative Binomial model. The latter is the second part of the two-part model we apply, the Negative Binomial Hurdle model. Thus, if not accounting for excess zeroes, we would model conditional savings in the same way that we do in this study, while accounting for a large proportion of non-savers.

D. Further Results and Robustness (For Online Publication)

Table 3: Summary Statistics – Colombian Data

	Mean	s.d.	Min	Max	Obs.
Individual Information					
Age	49.0	13.4	24	87	640
Male (=1)	0.28	0.45	0	1	640
Relationship to head of HH					
Head of household (=1)	0.64	0.48	0	1	640
Partner (=1)	0.23	0.42	0	1	640
Son/Daughter or their partner (=1)	0.07	0.25	0	1	640
Other (=1)	0.06	0.24	0	1	640
Household Characteristics					
Number of adult household members	2.8	1.4	1	12	640
Number of adolescents	1.2	1.3	0	7	640
Father still alive (=1)	0.31	0.46	0	1	640
Mother still alive (=1)	0.51	0.50	0	1	640
Exercising					
Every day (=1)	0.17	0.37	0	1	640
At least once a week (=1)	0.18	0.38	0	1	640
At least once a month (=1)	0.09	0.28	0	1	640
Never or hardly ever (=1)	0.57	0.50	0	1	640
Other Health Indicators					
BMI	25.7	4.3	12.9	43.0	640
Education					
Highest year passed	5.8	3.3	0	11	640
Financial literacy score (max. 18)	9.3	3.4	0	16	640
Financial Situation of the Household					
SISBEN Level 2 (=1)	0.50	0.50	0	1	640
Size of safety net (# Persons)	2.5	3.5	0	60	640
Monthly HH income per capita[a]	3.19	2.26	0.01	18.00	640
Market price of house[a]	180.10	408.86	0.00	3000.00	640
Debt[a]	17.24	65.68	0.00	588.04	640
Savings[a]	2.56	13.91	0.00	200.00	640
Engaging in saving (=1)	0.15	0.35	0	1	640
Conditional savings	17.61	32.82	0.20	200.00	93
Planning Horizon					
Day to day (=1)	0.74	0.44	0	1	640
Next months (=1)	0.18	0.38	0	1	640
Next year (=1)	0.05	0.21	0	1	640
Next two to five years (=1)	0.02	0.14	0	1	640
Next five to ten years (=1)	0.01	0.11	0	1	640

Note: [a]Figures reported in 100,000 COP.

Table 4: Summary Statistics: Income Risk – Colombian Data

	Mean	s.d.	Median	Min	Max	Obs.
Local Unemployment Risk (in pc)	24.7	6.0	25.4	15.2	36.5	640

Note: The unemployment risk results from self-reported individual figures in our survey that are averaged at the *UPZ* level; see Section 4 for details.

Table 5: Summary Statistics of Experimental Measures – Colombian Data

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Kahneman-Tversky (KT) $-\mu(-x)/\mu(x)$	1.1	2.7	0.4	0.1, 1.1	579
Neilson (N) $(\mu(-x)/-x)\,/\,(\mu(y)/y)$	0.2	0.5	0.0	0.0, 0.1	640
Wakker-Tversky (WT) $\mu'(-x)/\mu'(x)$	12.3	110.9	0.1	0.0, 0.3	564
Bowman (B) $\inf \mu'(-x)/\sup \mu'(x)$	0.1	0.2	0.0	0.0, 0.0	564
Köbberling-Wakker (KW) $\mu'(0_-)/\mu'(0_+)$	10.9	76.6	0.2	0.0, 1.0	640
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	1.1	4.6	0.1	0.0, 0.4	640
Meta Measure 2 (KT, KW, N)	1.0	3.1	0.1	0.0, 0.6	579
Meta Measure 3 (all)	0.3	1.0	0.1	0.0, 0.2	509
Impatience					
Near future impatience	29.6	15.2	22.0	16.0, 50.0	640
Increase in patience over time	0.3	16.3	0.9	-2.9, 0.9	640
Risk Preferences					
Utility Curvature: Gain Domain	6.0	29.9	0.7	0.2, 2.5	640
Utility Curvature: Loss Domain	8.0	16.0	1.1	0.5, 3.5	640
Probability Weighting: Gain Domain	41.5	32.9	40.6	9.4, 71.9	640
Probability Weighting: Loss Domain	68.5	28.5	78.1	46.9, 96.9	640

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail. Near future impatience is the mean annual interest rate, see Section 4. Utility curvature is the parameter of a power utility function and probability weighting is the probability that is perceived as 50%; see Section 4.

Table 6: Correlation Coefficients for Loss-Aversion Measures – Colombian Data

	KT	N	WT	B	KW	Meta 1	Meta 2	Meta 3
						KW, N	KT, KW, N	All
Kahneman-Tversky (KT)	1.000							
Neilson (N)	0.677***	1.000						
Wakker-Tversky (WT)	0.664***	0.281***	1.000					
Bowman (B)	0.020	0.088*	0.002	1.000				
Köbberling-Wakker (KW)	0.380***	0.536***	0.054	-0.019	1.000			
Meta Measure 1 (KW, N)	0.535***	0.780***	0.139***	-0.003	0.901***	1.000		
Meta Measure 2 (KT, KW, N)	0.720***	0.828***	0.276***	0.002	0.814***	0.955***	1.000	
Meta Measure 3 (all)	0.869***	0.714***	0.517***	0.140**	0.407***	0.575***	0.753***	1.000

Note: This table presents Pearson correlation coefficients for the different measures of loss aversion used in this study. See Section 4 in the main text for a brief overview, Section 5 for summary statistics, and Section B.1 in the Appendix for a more detailed account on these measures.

Table 7: Summary Statistics – Colombian Data – Sub-sample facing High Income Risk ($\geq$ 75% Quantile)

	Mean	s.d.	Min	Max	Obs.
Individual Information					
Age	49.3	12.6	24	87	217
Male (=1)	0.28	0.45	0	1	217
Relationship to head of HH					
Head of household (=1)	0.66	0.48	0	1	217
Partner (=1)	0.23	0.42	0	1	217
Son/Daughter or their partner (=1)	0.07	0.26	0	1	217
Other (=1)	0.04	0.19	0	1	217
Household Characteristics					
Number of adult household members	3.0	1.5	1	12	217
Number of adolescents	1.1	1.2	0	6	217
Father still alive (=1)	0.31	0.46	0	1	217
Mother still alive (=1)	0.53	0.50	0	1	217
Exercising					
Every day (=1)	0.19	0.39	0	1	217
At least once a week (=1)	0.16	0.36	0	1	217
At least once a month (=1)	0.06	0.23	0	1	217
Never or hardly ever (=1)	0.60	0.49	0	1	217
Other Health Indicators					
BMI	25.3	4.3	12.9	40.4	217
Education					
Highest year passed	5.4	3.2	0	11	217
Financial literacy score (max. 18)	9.1	3.2	0	15	217
Financial Situation of the Household					
SISBEN Level 2 (=1)	0.50	0.50	0	1	217
Size of safety net (# Persons)	2.8	4.9	0	60	217
Monthly HH income per capita[a]	3.02	2.02	0.13	15.00	217
Market price of house[a]	195.35	387.72	0.00	2000.00	217
Debt[a]	15.34	64.63	0.00	541.71	217
Savings[a]	3.69	19.29	0.00	200.00	217
Engaging in saving (=1)	0.13	0.34	0	1	217
Conditional savings	27.59	46.78	1.00	200.00	29
Planning Horizon					
Day to day (=1)	0.69	0.46	0	1	217
Next months (=1)	0.24	0.42	0	1	217
Next year (=1)	0.05	0.22	0	1	217
Next two to five years (=1)	0.01	0.12	0	1	217
Next five to ten years (=1)	0.01	0.12	0	1	217

Note: [a]Figures reported in 100,000 COP.

Table 8: Summary Statistics of Experimental Measures – Colombian Data – Sub-sample facing High Income Risk ($\geq$ 75% Quantile)

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.0	0.1	0.0	0.0, 0.0	193
Kahneman-Tversky (KT)	1.2	3.1	0.5	0.1, 1.2	198
Köbberling-Wakker (KW)	11.6	71.6	0.2	0.0, 1.8	217
Neilson (N)	0.2	0.5	0.0	0.0, 0.2	217
Wakker-Tversky (WT)	14.1	155.1	0.1	0.0, 0.3	193
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	1.2	5.2	0.1	0.0, 0.6	217
Meta Measure 2 (KT, KW, N)	1.1	3.5	0.2	0.0, 0.8	198
Meta Measure 3 (all)	0.3	0.9	0.1	0.0, 0.2	178
Impatience					
Near future impatience	28.2	14.9	18.0	16.0, 42.0	217
Increase in patience over time	-0.7	15.7	0.9	-2.9, 0.9	217
Risk Preferences					
Utility Curvature: Gain Domain	6.2	30.3	0.7	0.2, 2.5	217
Utility Curvature: Loss Domain	7.8	16.0	1.1	0.4, 2.9	217
Probability Weighting: Gain Domain	46.4	33.4	46.9	15.6, 78.1	217
Probability Weighting: Loss Domain	68.4	29.4	78.1	46.9, 96.9	217

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail. Near future impatience is the mean annual interest rate, see Section 4. Utility curvature is the parameter of a power utility function and probability weighting is the probability that is perceived as 50%; see Section 4.

Table 9: Summary Statistics – Colombian Data – Loss-Averse Sub-sample (Loss Aversion Meta Measure 1 > 1)

	Mean	s.d.	Min	Max	Obs.
Individual Information					
Age	47.5	14.6	24	77	97
Male (=1)	0.30	0.46	0	1	97
Relationship to head of HH					
Head of household (=1)	0.63	0.49	0	1	97
Partner (=1)	0.23	0.42	0	1	97
Son/Daughter or their partner (=1)	0.07	0.26	0	1	97
Other (=1)	0.07	0.26	0	1	97
Household Characteristics					
Number of adult household members	2.8	1.5	1	8	97
Number of adolescents	1.2	1.2	0	5	97
Father still alive (=1)	0.39	0.49	0	1	97
Mother still alive (=1)	0.49	0.50	0	1	97
Exercising					
Every day (=1)	0.21	0.41	0	1	97
At least once a week (=1)	0.25	0.43	0	1	97
At least once a month (=1)	0.05	0.22	0	1	97
Never or hardly ever (=1)	0.49	0.50	0	1	97
Other Health Indicators					
BMI	25.9	4.3	15.9	39.0	97
Education					
Highest year passed	5.6	3.1	0	11	97
Financial literacy score (max. 18)	9.8	2.9	2	16	97
Financial Situation of the Household					
SISBEN Level 2 (=1)	0.55	0.50	0	1	97
Size of safety net (# Persons)	2.5	3.5	0	30	97
Monthly HH income per capita[a]	3.17	1.99	0.36	10.00	97
Market price of house[a]	203.09	426.03	0.00	2000.00	97
Debt[a]	11.13	52.99	0.00	476.71	97
Savings[a]	2.96	16.22	0.00	150.00	97
Engaging in saving (=1)	0.14	0.35	0	1	97
Conditional savings	20.48	39.46	1.50	150.00	14
Planning Horizon					
Day to day (=1)	0.75	0.43	0	1	97
Next months (=1)	0.16	0.37	0	1	97
Next year (=1)	0.06	0.24	0	1	97
Next two to five years (=1)	0.02	0.14	0	1	97
Next five to ten years (=1)	0.00	0.00	0	0	97

Note: [a]Figures reported in 100,000 COP.

Table 10: Summary Statistics: Income Risk – Colombian Data – Loss-Averse Sub-sample (Loss Aversion Meta Measure 1 > 1)

	Mean	s.d.	Median	Min	Max	Obs.
Local Unemployment Risk (in pc)	25.4	6.5	26.1	15.2	36.5	97

Note: The unemployment risk results from self-reported individual figures in our survey that are averaged at the *UPZ* level; see Section 4 for details.

Table 11: Summary Statistics of Experimental Measures – Colombian Data – Loss-Averse Sub-sample (Loss Aversion Meta Measure 1 > 1)

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.1	0.2	0.0	0.0, 0.1	89
Kahneman-Tversky (KT)	4.4	5.3	2.3	1.6, 4.6	96
Köbberling-Wakker (KW)	68.7	187.2	12.5	4.0, 36.9	97
Neilson (N)	1.0	0.9	0.7	0.5, 1.3	97
Wakker-Tversky (WT)	69.8	271.0	0.4	0.0, 3.3	89
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	6.2	10.5	2.6	1.5, 4.6	97
Meta Measure 2 (KT, KW, N)	4.8	6.5	2.7	1.7, 4.4	96
Meta Measure 3 (all)	1.4	2.0	0.6	0.3, 1.6	88
Impatience					
Near future impatience	31.1	14.9	26.0	16.0, 50.0	97
Increase in patience over time	0.9	17.6	0.9	-2.9, 3.7	97
Risk Preferences					
Utility Curvature: Gain Domain	13.7	19.5	2.7	1.2, 20.3	97
Utility Curvature: Loss Domain	2.1	6.3	0.6	0.4, 1.3	97
Probability Weighting: Gain Domain	61.4	27.8	65.6	46.9, 84.4	97
Probability Weighting: Loss Domain	68.3	27.2	71.9	46.9, 96.9	97

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail. Near future impatience is the mean annual interest rate, see Section 4. Utility curvature is the parameter of a power utility function and probability weighting is the probability that is perceived as 50%; see Section 4.

Table 12: Results from Estimating Model 2 Using a Negative Binomial Hurdle Model and Different Meta-Measures of Loss Aversion – Colombian Data

	Loss Aversion 2 Measures (KW, N)			Loss Aversion 3 Measures (KW,N,KT)		Loss Aversion 5 Measures (All)
	(1)	(2)	(3)	(4)	(5)	(6)
Likelihood of Saving						
Loss Aversion	0.042	0.043	0.030	0.079**	0.063	0.266**
	(1.61)	(1.58)	(1.09)	(1.96)	(1.51)	(2.48)
Income Risk (Survey)	0.027	0.055	0.068**	0.055	0.069**	0.067**
	(1.07)	(1.64)	(2.44)	(1.67)	(2.53)	(2.44)
Amount of Savings						
Loss Aversion	0.062**	0.068***	0.085***	0.092***	0.100***	0.228*
	(2.13)	(2.97)	(3.79)	(3.52)	(3.53)	(1.93)
Income Risk (Survey)	0.145*	0.165***	0.136*	0.166***	0.138*	0.139*
	(2.06)	(3.33)	(1.88)	(3.33)	(1.89)	(1.73)
AIC	1072	967	858	964	857	860
Controls	25	25	25	25	25	25
Region	Yes	Yes	Yes	Yes	Yes	Yes
Occupation	Yes	Yes	Yes	Yes	Yes	Yes
Observations	640	579	509	579	509	509

⋆ $p < 0.1$, ⋆⋆ $p < 0.05$, ⋆⋆⋆ $p < 0.01$. Wild cluster (score) bootstrapped t-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various savings devices; see Section 4. In this Negative Binomial Hurdle model, the participation equation estimates the likelihood to engage in savings, while the second equation estimates conditional savings – the amount of savings, given that a person is saving. Loss aversion is measured by continuous and experimentally elicited meta-measures. The meta-measure comprising two measures of loss aversion is the geometric mean of loss aversion coefficients according to the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005). The measure including three measures is the geometric mean of the former two loss aversion coefficients and, in addition, the one building on the definition of loss aversion by Kahneman and Tversky (1979). Finally, for the last measure, the coefficients based on definitions by Bowman et al. (1999) and Wakker and Tversky (1993) are also included. For more details on the applied measures of loss aversion, see Appendix B.1. Column 2 shows the results when restricting the sample to those for which the meta-measure combining three measures of loss aversion is available, and columns 3 and 5 show the results for similarly restricted samples, in order to be able to draw comparisons between the different meta-measures of loss aversion. We control for variables listed in Tables 3 and 5. Furthermore, we control for regional and occupational sectors at *localidad* level, as well as for the working sectors according to the ISIC classification of economic activities. We account for the cluster structure (at the UPZ level) and potential heteroskedasticity in our data by using wild cluster bootstrapping (Cameron et al., 2008).

Table 13: Results from Estimating Model 3 Using a Negative Binomial Hurdle Model and Different Meta-Measures of Loss Aversion – Colombian Data

	Measure 1 (KW, N)		Measure 2 (KT, KW, N)		Measure 3 (All)	
Likelihood of Saving						
Income Risk (IR)	0.025	(1.02)	0.055*	(1.72)	0.096***	(3.99)
Loss Aversion (LA)	0.032	(1.01)	0.064	(1.42)	0.131	(1.09)
LA × IR	0.012*	(1.70)	0.017**	(1.89)	0.046**	(2.57)
Amount of Savings						
Income Risk (IR)	0.150***	(2.67)	0.167***	(3.30)	0.186***	(2.70)
Loss Aversion (LA)	-0.177***	(-3.57)	-0.129	(-1.36)	-0.169	(-0.76)
LA × IR	0.022***	(5.29)	0.020***	(2.55)	0.047**	(1.96)
AIC	1055		955		853	
Controls	25		25		25	
Region	Yes		Yes		Yes	
Occupation	Yes		Yes		Yes	
Observations	640		579		509	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Wild cluster (score) bootstrapped t-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various savings devices; see Section 4. In this Negative Binomial Hurdle model, the participation equation estimates the likelihood to engage in savings, while the second equation estimates conditional savings – the amount of savings, given that a person is saving. Loss aversion is measured by continuous and experimentally elicited meta-measures. The meta-measure comprising two measures of loss aversion is the geometric mean of loss aversion coefficients according to the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005) (Measure 1). The measure including three measures is the geometric mean of the former two loss aversion coefficients and, in addition, the one building on the definition of loss aversion by Kahneman and Tversky (1979) (Measure 2). Finally, for the last measure, the coefficients based on definitions by Bowman et al. (1999) and Wakker and Tversky (1993) are also included (Measure 3). The coefficients of loss aversion are centered at 1; for more details on the applied measures of loss aversion, see Appendix B.1. Income risk is centered at the mean; see Section 4 for details. We control for variables listed in Tables 3 and 5. Furthermore, we control for regional and occupational sectors at *localidad* level as well as for the working sectors according to the ISIC classification of economic activities, if indicated. We account for the cluster structure (at the UPZ level) and potential heteroskedasticity in our data by using wild cluster bootstrapping (Cameron et al., 2008).

Table 14: Results from Estimating Model 3 Using a Random Effects Panel Regression and Different Meta-Measures of Loss Aversion – Laboratory Experiment

	Loss Aversion 2 Measures (KW, N)		Loss Aversion 3 Measures (KW, N, KT)		Loss Aversion 5 Measures (All)	
		Risk: High IR ($\geq$ 75% Quantile)		Risk: High IR ($\geq$ 75% Quantile)		Risk: High IR ($\geq$ 75% Quantile)
Panel A: Saving						
Income Risk (IR)	0.132***		0.130***		0.129***	
	(10.476)		(10.173)		(7.496)	
High IR (= 1)		3.525***		3.480***		3.511***
		(9.409)		(9.118)		(7.363)
Loss Aversion (LA)	-0.100	-0.464	-0.139	-0.495	-0.867	-1.045
	(-0.336)	(-1.533)	(-0.407)	(-1.423)	(-1.268)	(-1.461)
LA $\times$ IR	0.028***		0.027**		0.011	
	(2.775)		(2.252)		(0.445)	
LA $\times$ High IR (= 1)		0.851***		0.833**		0.467
		(3.093)		(2.527)		(0.780)
Panel B: Lin. Combination						
Loss Aversion		0.386		0.338		-0.578
		(0.284)		(0.419)		(0.459)
Observations	2562	2562	2562	2562	2562	2562

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: The dependent variable is the amount saved in a given saving decision task, and for every individual, we have 14 saving decisions, see Section 6.2. Following Xu et al. (2022), we employ a random effects panel regression to account for the multiple (dependent) observations from every individual. Loss aversion is measured by continuous and experimentally elicited meta-measures. The meta-measure comprising two measures of loss aversion is the geometric mean of loss aversion coefficients according to the definitions of loss aversion by Neilson (2002) and Köbberling and Wakker (2005) (Measure 1). The measure including three measures is the geometric mean of the former two loss aversion coefficients and, in addition, the one building on the definition of loss aversion by Kahneman and Tversky (1979) (Measure 2). Finally, for the last measure, the coefficients based on definitions by Bowman et al. (1999) and Wakker and Tversky (1993) are also included (Measure 3). The coefficients of loss aversion are centered at 1; for more details on the applied measures of loss aversion, see Appendix B.1. Income risk is centered at the mean.

Table 15: Summary Statistics of Experimental Measures – Laboratory Experiment

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.2	0.2	0.1	$0.0, 0.3$	183
Kahneman-Tversky (KT)	1.1	0.9	1.0	$0.4, 1.5$	183
Köbberling-Wakker (KW)	1.7	3.7	0.6	$0.1, 1.8$	183
Neilson (N)	0.4	0.4	0.2	$0.1, 0.6$	183
Wakker-Tversky (WT)	1.0	4.0	0.6	$0.3, 0.9$	183
Meta Measures of Loss Aversion					
Meta Measure 1 (KW, N)	0.7	1.0	0.4	$0.1, 0.9$	183
Meta Measure 2 (KT, KW, N)	0.7	0.9	0.5	$0.2, 1.0$	183
Meta Measure 3 (all)	0.5	0.5	0.3	$0.1, 0.7$	183

Note: The measures and meta-measures of loss aversion are described in Section 4 and in Appendix B.1 with greater detail.

E. Instructions – Colombian Lab-in-the-Field Experiment (Translated from Spanish; for Online Publication)

[Instructions for reading to participants]

[Workshop leader] Make sure the person is not distracted by other matters. Read aloud from the script and always be alert to any questions. Be alert to participants' facial expressions to detect lack of understanding of the game.

Good morning. I am ____________ and I would first like to thank you for participating in this study "Saving for old age". The objective of the study is to learn more about the possibilities of saving for old age by Sisben level 1 and 2 households.

The study has two parts. One is the survey that we already did in the past days and the other part is this workshop.

In recognition of your collaboration we will give you a participation reward of 10,000 pesos that regardless of how you do in the game, you will take home safely.

In addition, during this workshop, you will have the possibility to earn more money. Out of 100 participants, 20 will be paid for their decisions in the activity. These 20 people will be selected at random. Once you finish the activity you will draw a ball from a bingo that has 100 balls numbered from 1 to 100. If the ball you draw has a number between 1 and 20, you will be paid for one decision. That decision will also be randomly selected so think very carefully about each decision as any one of them could be selected to be paid. The amount of money you win depends on your decisions as well as luck.

Because you can earn money for your decisions, it is very important that you pay attention to these instructions. In case there is anything you do not understand in the instructions, let me know and I will be happy to answer any questions. To avoid any distractions during the exercise I will ask you to turn off your cell phone.

What do you have to do during the workshop? During the workshop you have to choose between two investment options. In each of the investment options – as in real life – things can go right or they can go wrong. For example, when you start a business, it is possible that the business works out and you make a good profit, or it can be that conditions are unfavorable and you lose.

This happens because there are things that do not depend on one's control, but on chance or other factors.

You will have an endowment of twelve million five hundred thousand (12,500,000) pesos to start with. Depending on your decisions and luck, you can increase or decrease your endowment. This means that you will make each decision thinking that you have twelve million five hundred thousand (12,500,000) pesos in your pocket. In the end, for every 500 pesos you end up having, we will pay you 1 Colombian peso. Twelve million five hundred thousand (12,500,000) experimental pesos then translates into twenty-five thousand (25,000) Colombian pesos, since twelve million five hundred thousand divided by 500 gives twenty-five thousand. From this amount, you can win more money, or lose part of that money so think very well about each decision as any of the questions could be selected to be paid.

In this activity, what you have to do is to decide between two options to invest: Option A and Option B. You choose only one option, A or B, whichever you prefer. The options involve different payoffs that occur under different circumstances. The decisions you make will be kept private. Choose the option you like best and keep in mind that throughout the exercise there are no right or wrong decisions, it all depends on your preferences.

You will find that for each option, A and B there will be a roulette like this [SHOW ROULETTE] which will determine whether things go right or wrong, i.e., it will determine the amount of money you win on each decision. The roulette wheel has a red area and a blue area. If the roulette wheel lands on the red color, you will receive the payout corresponding to the red bar, and if it lands on blue you will receive the payout corresponding to the blue bar. The needle is spun and the color where the needle lands will then determine the payout you will receive.

Throughout the workshop, we will show you figures like this one:

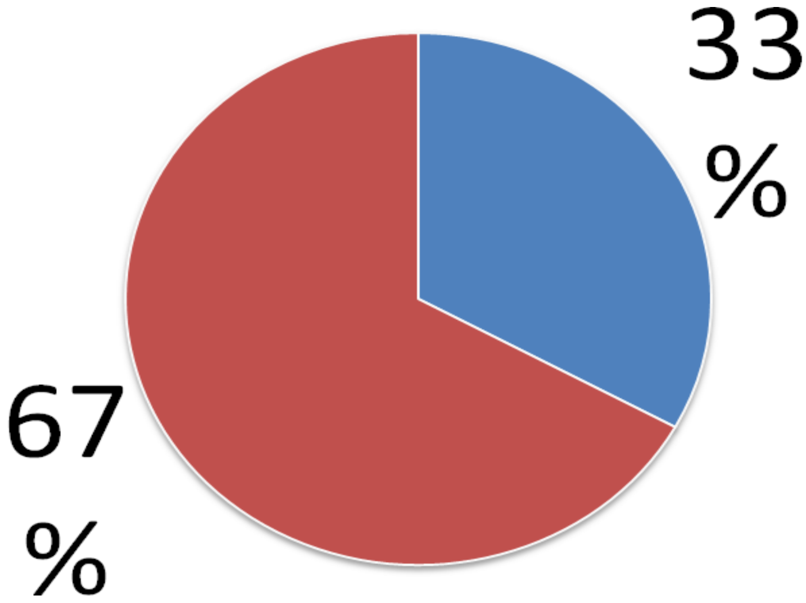

Figure 3: Roulette Wheel

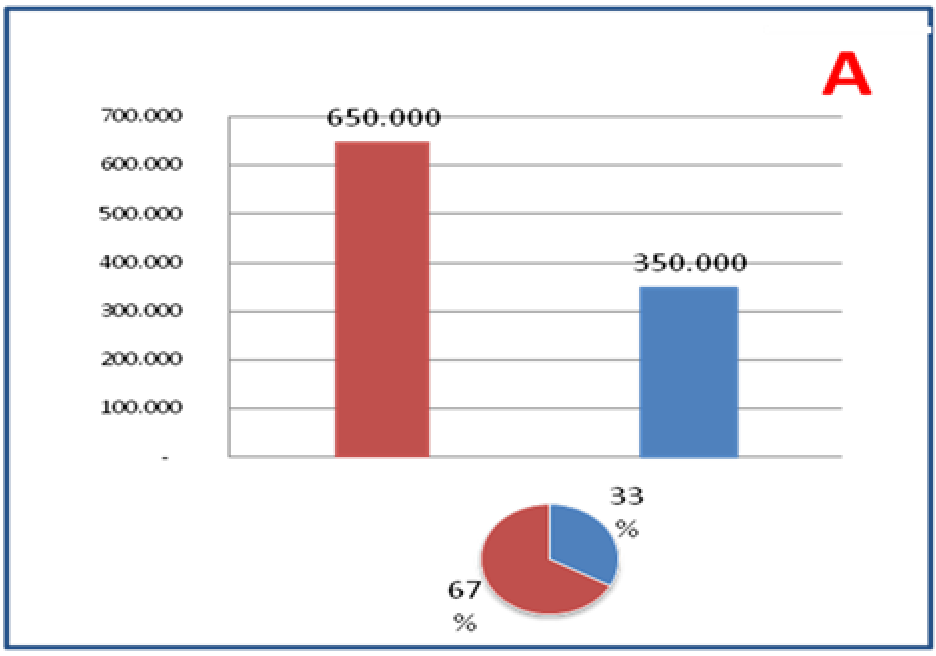

The bars indicate the possible values you will receive. We see two bars, one indicating 650 thousand pesos and the other indicating 350 thousand pesos. Below the bar, a roulette wheel appears, with a red part and a smaller blue part. This indicates that if the needle falls on the red color you receive 650 thousand experimental pesos, but if the roulette falls on the blue color, you receive 350 thousand experimental pesos.

Do you have any questions so far?

[ON THE NEXT PAGE THE CONTROL QUESTIONS START]

E.1. Control Questions 1

Now consider the following investment. [SHOW FIGURE]

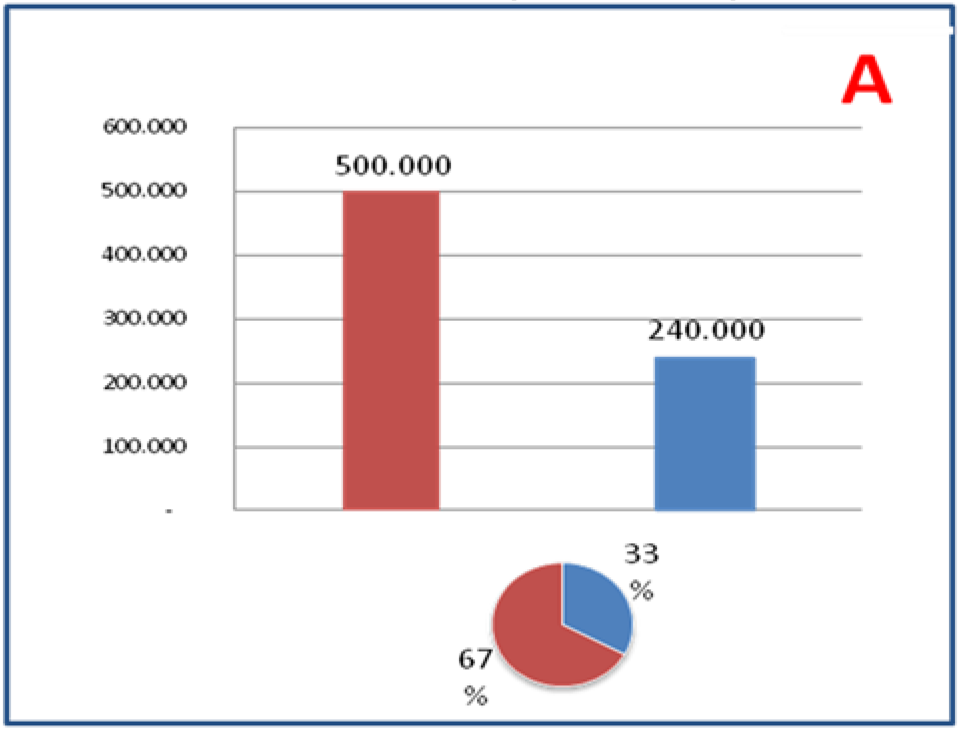

Tell us, in this decision, how much you would receive if things go well $ ________

And how much would you receive if things go wrong? $ ___________

Considering the colors of the roulette wheel, which payment are you most likely to receive? $______

If the roulette wheel falls on the blue color, what payment would you receive? $ __________

E.2. Example Decision

Now let us look at an example of the decisions we will present in the exercise.

Consider this decision.

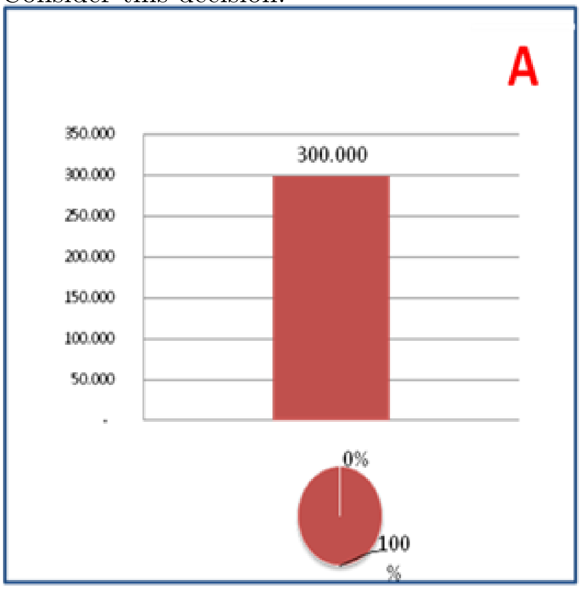

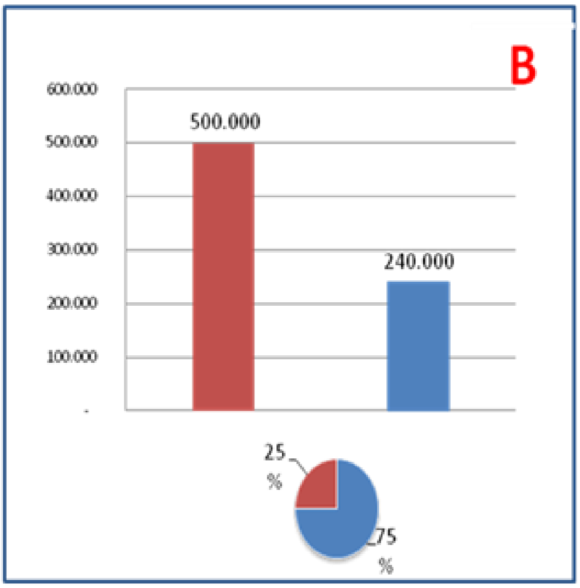

In option A you always win 300,000 pesos, because in this case the roulette wheel is all in red. In option B, if the roulette wheel falls on red you win 500,000 experimental pesos. But if by chance, the roulette wheel lands on blue you would receive 240,000 experimental pesos. The amount of money you win will depend on the area where the needle lands. Your task is to decide which of the two options you prefer, A or B. In this exercise there are no right or wrong decisions, it all depends on what you prefer.

Before we continue, we would like to ask you some questions to verify that you have understood the exercise correctly.

[ON THE NEXT PAGE THE CONTROL QUESTIONS START]

E.3. Control Questions 2

1. Between roulette A and roulette B, where is the needle most likely to land on the blue color?

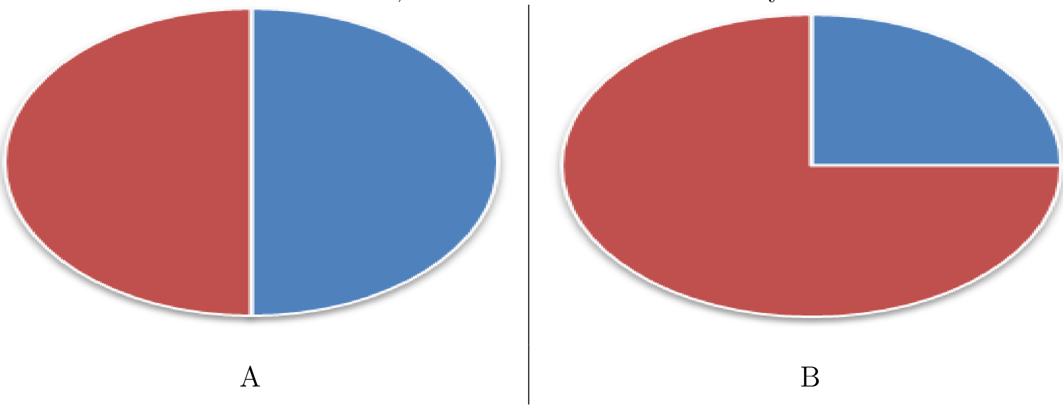

2. In which of these options is there a greater probability of earning more?

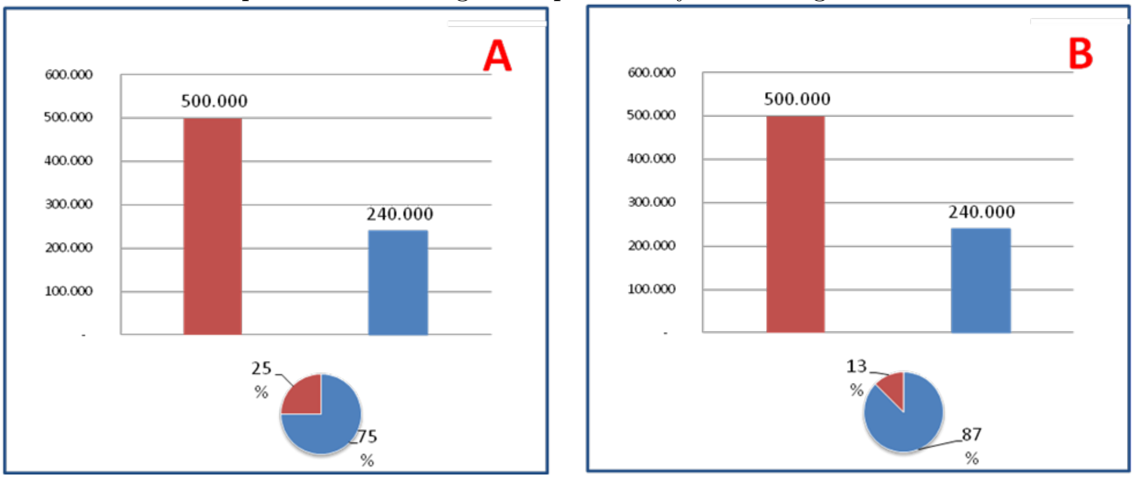

3. And between these two options, when is there a better chance of earning more?

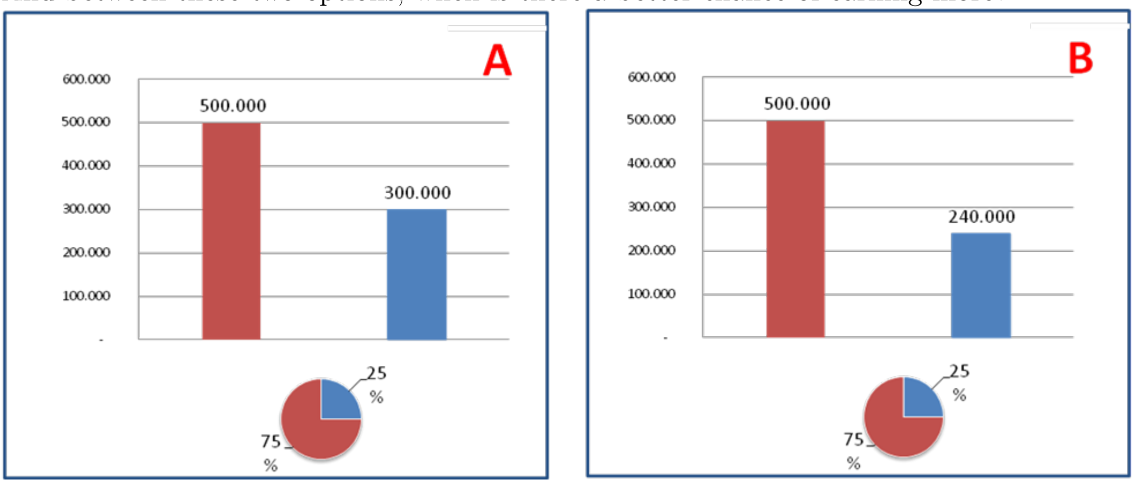

Throughout the workshop you will have to make 100 decisions like this one. Once you finish making the decisions, we will randomly determine which of the 100 decisions is selected to be paid. Therefore, it is important that you make each decision carefully as you will earn the payout for only one of them.

Once the decision to pay is selected, we will proceed to determine whether you receive a high or low payout by spinning the roulette wheel. Depending on your decision and the color the needle lands on, your payout is determined. The payout value is divided by 500, as each 500 experimental pesos represent 1 peso that you are paid. Finally you will choose a bingo ball with balls numbered from 1 to 100. If you draw a ball with a number between 1 and 20, you will receive the 10,000 pesos for having participated, plus or minus the results you obtained in the game. Otherwise you will receive 10,000 pesos for having participated in the workshop.

Finally, keep in mind that although some of the decisions involve the possibility of a negative payout or loss, it will never be more than the twelve million five hundred (12,500,000) endowment you receive at the start of the game. Are you ready to start the game?

[Workshop Leader] Start the game. Do not forget to go back to the instructions in question 106 where the time preference questions start, and hence, new additional instructions are required.

[Workshop leader] Keep in mind when explaining that for each question the person has an endowment of 12,500,000 experimental pesos – especially for the losses.

E.4. Part 2

[Workshop leader] Read it once you get to question 106.

In this part of the activity you will be asked a different set of questions than the ones you just have seen. In this part you will have to decide between two options A and B. However, the options shown do not depend on the roulette wheel as in the previous questions; they give you the option to choose in how much time you would like to receive how much money. That means that you will be shown two different amounts of money that you could receive at two different points in time. Your task is to decide which of the options A or B you prefer depending on the amount of money and the time in which you would receive it. In case one of these questions is selected to be paid, you would receive the 10 thousand pesos for participating today, and you would receive the payment corresponding to the decision **at the time indicated in the question**. For this we would take your data and you would be given a receipt to claim your money at the time indicated in the question. Do you have any questions?

[Workshop leader] Once the game is over, when the question to be paid is selected, explain with this circle the roulette result that will define whether the payout is high or low.

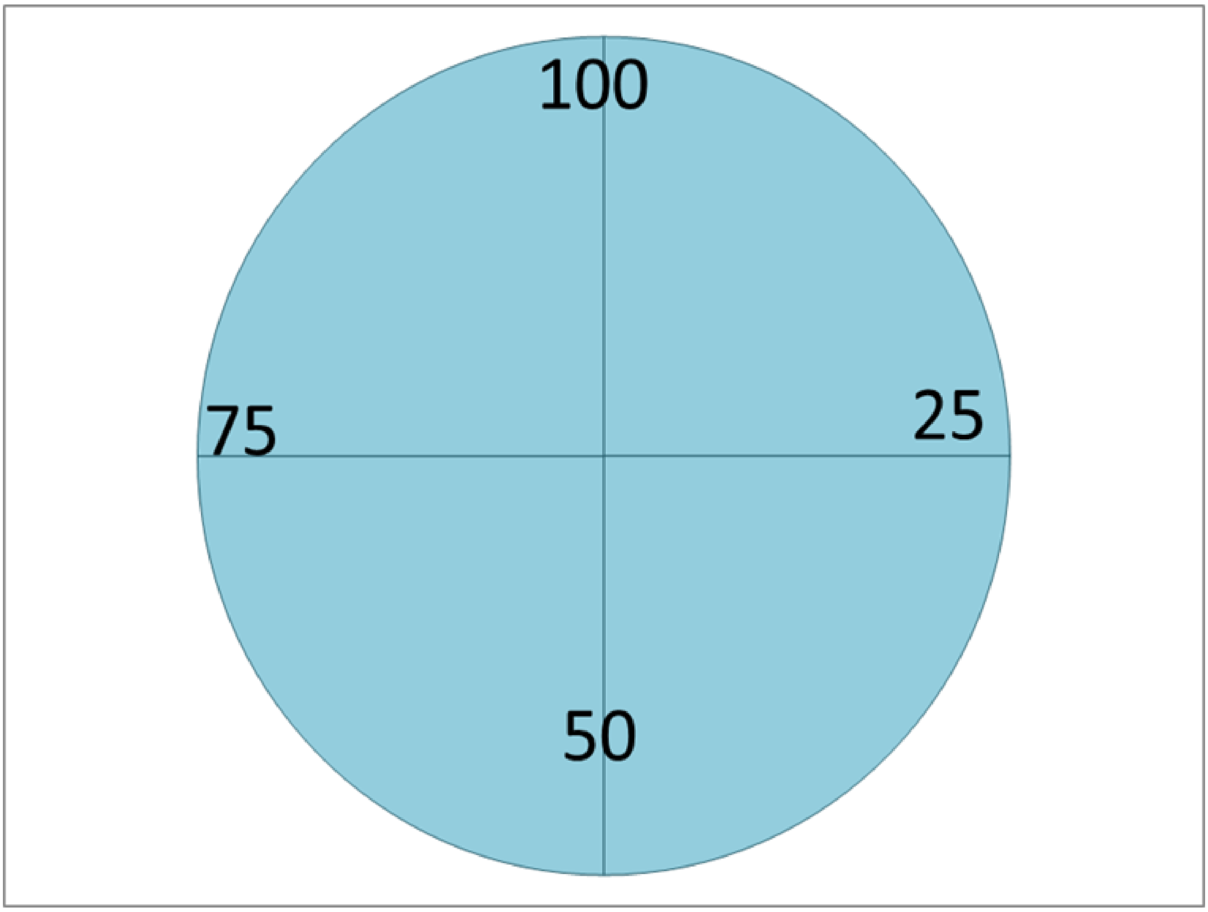
100
75
25
50

F. Instructions – Lab-Experiment (Translated; for Online Publication)

Welcome to the experiment!

Please read these instructions carefully. If you have any questions, please raise your hand.

As you know, you can earn money in this experiment. The exact amount also depends on your decisions.

The experiment consists of two sub-experiments. In each of the two parts, you will have to make a number of money-related decisions. To determine your payoff at the end of the experiment, one part (each with equal probability) is randomly selected as payoff-relevant. From the payoff-relevant part, a decision is again randomly determined, with each decision within that part being equally likely.

Thus, any decision you make could potentially become payout relevant. Since you do not know which one this will be, please process all tasks carefully.

We have also included attention tests to make the data quality measurable for us, but also to ensure fair payoffs. Throughout the experiment we have built in several tests. If your decision in all tests is that of an attentive participant, you get the full payout. If, on the other hand, all tests are answered against all reason, you will receive the minimum payout of EUR 7. If some – but not all – tests imply careful decision making, you will receive the minimum payout plus the difference between the minimum payout and your actual payout, multiplied by the proportion of "correct" test questions. (For example: half of all tests "passed" with an actual payout of EUR 27. Then the difference from the minimum payout is EUR 20, multiplied by 1/2 gives 10, plus EUR 7 minimum payout gives a final payout of EUR 17).

Continue with the instructions for the first part of the experiment, the wheel of fortune game.

[Continue]

F.1. The Wheel-of-Fortune Game

In the wheel-of-fortune game, a wheel of fortune is spun (or the random spinning of a wheel of fortune is simulated by the computer). This determines the payout in this game. Here you can see an example wheel of fortune:

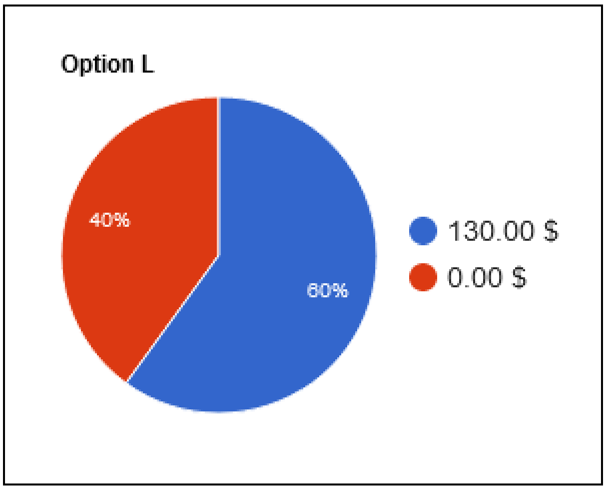

If the wheel of fortune comes to a stop on the **blue piece**, you will receive **$130.00**, as can be seen from the legend to the right of the circle of Option L. This happens with a **probability of 60%** – you can see this number from the corresponding indication on the blue piece. Also, the size of the blue piece takes 60% of the circle.

However, if the wheel of fortune comes to rest on the **red piece**, you will receive **$0.00**. Again, this can be seen in the legend to the right of the circle. This happens with a **probability of 40%** - you can also see this from the indication on the red piece. At the same time, the red area takes up 40% of the circle.

If you want, you can also move the mouse over the circle pieces – you will then be shown the information in a small "text-it".

If the circle consists of only one piece, chance naturally plays no role, because no matter where the wheel of fortune comes to a stop – the amount of money is always the same.

[Continue]

Screen Two In this part of the experiment, we will present you with different wheels of fortune at each step. Your task is to choose the wheel of fortune that you prefer. You can always choose between two options, "Option L" and "Option R".

Here you can see an exemplary decision situation:

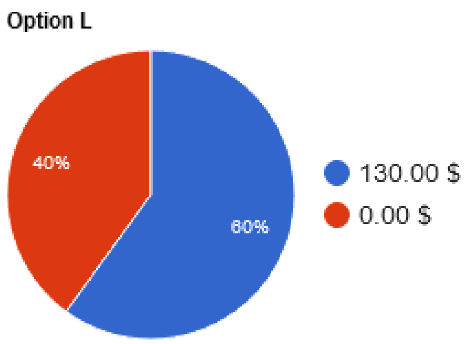

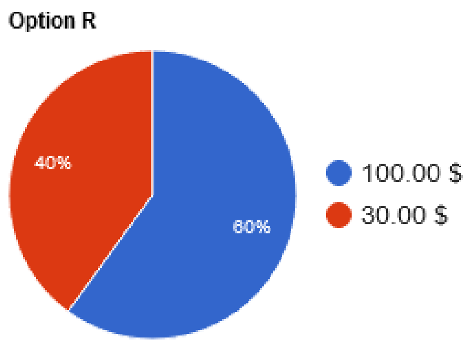

In Wheel of Fortune “Option L” you can get the following payouts:

- $ 130.00 with a probability of 60%, or
- $ 0.00 with a probability of 40%

In the wheel of fortune "Option R" you can get the following payouts:

- $ 100.00 with a probability of 60%, or
- $ 30.00 with a probability of 40%

Your task is to choose your preferred option. To enter your choice, first click the corresponding button under the options (labeled “L” or “R”), and then click “Next”.

In the decision situations, both the pieces of the circle will be different in terms of their size, as well as the amounts paid out. Each decision situation is different.

As mentioned earlier, you should always choose the option you prefer and make each decision as if it will determine your payoff – because it might. Also, some later decisions will depend on earlier decisions; the later decisions will only make sense to you if the earlier ones were made wisely.

The amounts shown here are in experimental dollars ($). They are multiplied by 0.001 in the Wheel of Fortune game in case of payout and this amount is then paid out (rounded to 10 cents) in Euros (that is: $10,000 equals 10 Euros).

For this part of the experiment ("Wheel of Fortune" game) we will give you an endowment of $15,000. Your payouts will be added to this endowment, and the total will be paid out if this sub-experiment is selected for payout. Some decisions in this first part of the experiment also include negative values; this may also reduce your endowment accordingly (in extreme cases it will be used up; however, we will pay out at least EUR 7 to you today – if this game is selected for payout).

If you have no questions and would like to participate in the experiment, you can now click "Continue" to check your understanding and start a short training round for the experiment. If you have any questions, or do not wish to continue participating, please raise your hand.

[Continue]

Comprehension Questions

Here you can see an exemplary decision situation:

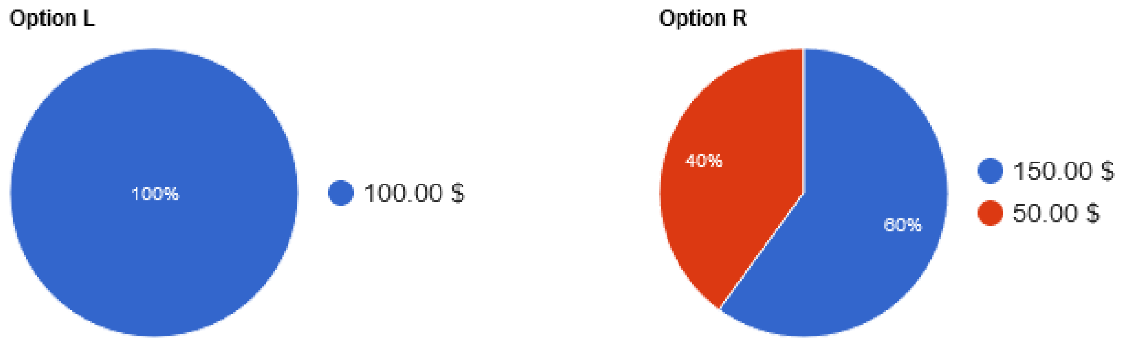

What is the minimum payout in Option L?

○ 0 ○ 50 ○ 100

With what probability (in %) is the payoff equal to the higher amount in option R?

○ 40 ○ 50 ○ 60

[Continue]

Training round

Please choose between option L and option R

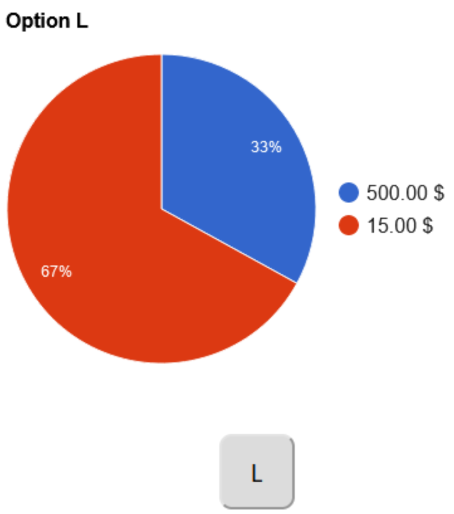

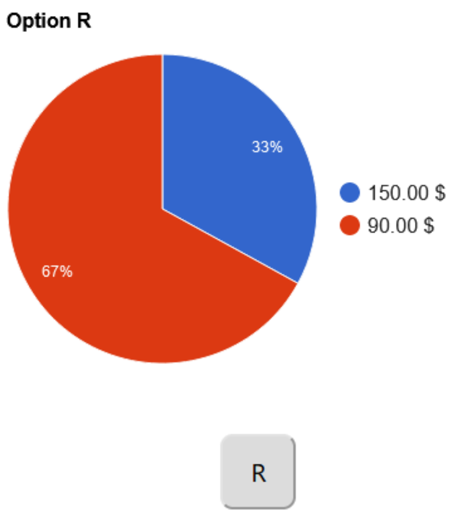

Take a good look at the two options: a payout of \$150.00 with a probability of 33% means that if you spin 100 times, the wheel of fortune will usually come to a stop on the blue piece around 33 times. The same is true for the other payout. With 100 spins, we would expect an average payout of about \$110.00 = 1/3 * \$150.00 + 2/3 * \$90.00 for Option R.

Alternatively, we could look at the differences in amounts, as the probabilities here are the same in both options (33% for the larger amount in both options). The lowest payoff in this decision situation is in one option – Option L – combined with the highest payoff. While the difference between the low payoffs of both options is \$75.00, it is \$350.00 for the high payoff, in the other direction – that is more than 4 times as much.

However, the probability for the high payout is half as high as for the low payout. In pure mathematical terms ("Better" amount more than 4 times the difference in the lower amounts; probability for the high amount in each case 1/2 as high as for the low amount), you could say that in Option L the maximum amount more than "doubles" the difference in the lower payout compared to Option R (4 * 1/2 = 2).

Note: For negative amounts, the "reverse" calculation applies accordingly (if all amounts in the example had a negative sign, Option R would be "half as bad").

Over the course of a few decisions, one option will always remain constant while the other changes. The amounts are chosen in such a way that the vast majority of participants will switch from one option to the other at some point.

[Continue]

Please choose between option L and option R

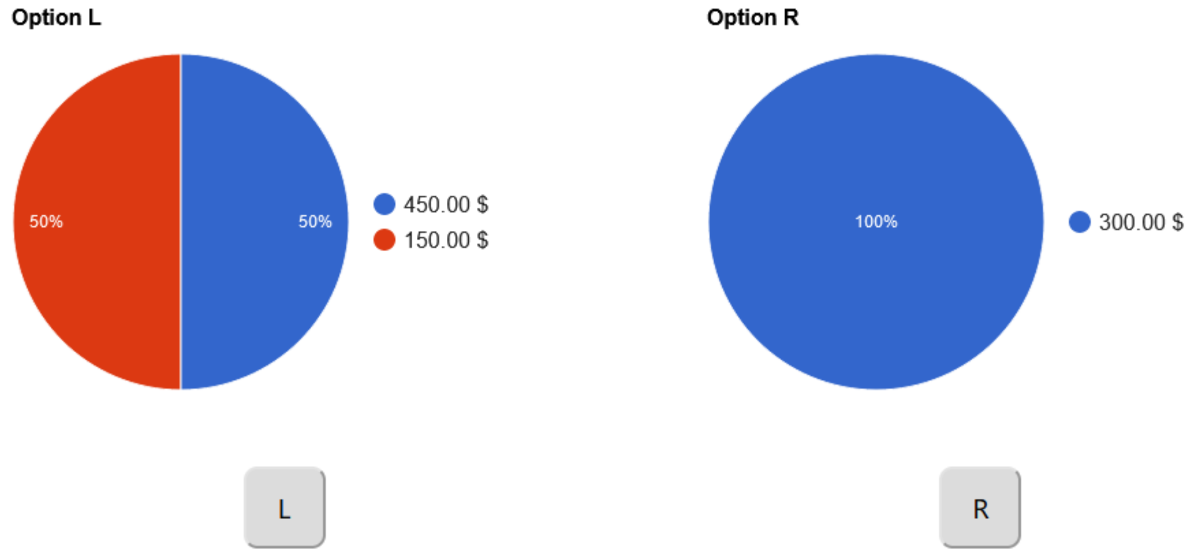

In these two options you can see that one option – Option R – is a “safe choice”. It is higher here than the lowest payout of Option L, by $150.00, however there is also no chance of getting the high payout of Option L, which is also $150.00 “away” from the safe payout.

If we spun 100 times, we would get the same amount in Option L and Option R on average.

Again, the amounts are such that the vast majority of participants will eventually switch from one option to the other at least once (possibly more than once here).

F.2. Savings Decisions

In this part of the experiment, we ask for some choices with respect to saving preferences. There are no right or wrong decisions. The explanations in the following paragraph will become clear in the following training round – nevertheless, please read this paragraph as well as the following ones carefully; otherwise you may not understand the training round.

We will show you payouts at two different points in time, where you can “move” money between the earlier point in time (“Time 1”) and the later point in time (“Time 2”). Apart from your savings decision, the payout at the earlier point in time is already fixed, while there are two possibilities for the payout at the later point in time, each with equal probability.

In the following decision situations, you can either move money from the earlier to the later point in time, i.e., save, or move from the later to the earlier point in time, i.e., borrow money. In doing so, you can shift a maximum of 20 experimental dollars. Don’t worry, as you will see in a moment in the training round, the tasks are very intuitive and you can “try” different options before making your decision.

In this game, 1 Experimental Dollar is equal to EUR 0.80, or 80 cents. Again, the result is paid in Euros rounded to the nearest 10 cents.

Should one of the savings decisions be selected for payment, **chance will** decide whether the payment will **be paid at the earlier or later point in time** – and for the **sake of simplicity, it will be paid today, even if the payment should be selected at the later point in time**. If the payment is selected at the later point in time, chance will additionally decide whether the high or the low amount will be paid out. Accordingly, chance again plays a role in determining your payout.

Note: We did not include attention tests in this part of the experiment. However, the amounts are quite high – so choose carefully.

If you have no further questions, you can proceed to the training round by clicking "Continue".

[Continue]

Training Round 1

Screen 1

Click on the gray bar to make the slider visible.

Screen 2 (after having clicked on the slider bar)

Date 1: $ 38 | Date 2: $ 36 or $ 36

With a 100% probability | Each with a 50% probability

Your choice: you save $0

The slider is set to $0 [VALUE IS ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER]. A positive amount means that you are moving money from the earlier point in time (Date 1) to the later point in time (Date 2); thus, you are saving money that will be subtracted from the payout at date 1 and added to the payouts at date 2. Conversely, a negative amount means that you are moving money from the later point in time to the earlier point in time; thus, you are borrowing money in date 1 from date 2.

Take a look at the payout options with this setting at both times:

- At the earlier point in time, “Date 1”, there is only one payoff option, since the amount listed under “Date 1” at the top left, $38 [VALUE IS ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER], is paid out with a probability of 100% (if Date 1 were randomly selected for today’s payoff in this decision situation).
- At the later time, “Date 2” there are always two payout options: The two amounts mentioned under “Date 2” on the top right, $36 and $36 [VALUES ARE ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER]. They will both be paid out with a probability of 50% each, should date 2 be randomly selected for today’s payout in this decision situation. In this decision situation, these amounts are identical, therefore there is practically only one amount at date 2; but this can change in the following situations.

If you haven’t tried it yet: Move the slider and watch how the amounts change to the earlier and later point in time – which choice do you like best?

[Continue]

Training Round 2

Screen 1

Click on the gray bar to make the slider visible.

Screen 2 (after having clicked on the slider bar)

Date 1: $ 38	Date 2: $ 28 or $ 44
With a 100% probability	Each with a 50% probability

Your choice: you save $0

The slider is set to $0 [VALUE IS ADJUSTED DYNAMICALLY DEPENDING ON THE SLIDER].

Now take a good look again at the payout possibilities in this setting at both times.

- You can see that the amounts at time 2 (i.e., the amounts mentioned under date 2 above right) are now different. Both are selected for payment with a probability of 50% – should this decision situation as well as date 2 be selected for payment at random.

- The amount at time 1 (i.e., the amount mentioned under time 1 above left) will be selected for payment unchanged with a probability of 100% – should this decision situation as well as date 1 be randomly selected for payment.

Which choice do you like best here?