UNIVERSITY OF GOETTINGEN

DOCTORAL THESIS

Advances and Applications of Experimental Measures to Test Behavioral Saving Theories and a Method to Increase Efficiency in Binary and Multiple Treatment Assignment

---

Author:
Sebastian O. SCHNEIDER

December 21, 2017

A thesis submitted in fulfillment of the requirements for the degree of Dr. rer. pol.

from the

Faculty of Economic Sciences

First Supervisor: Prof. Marcela Ibañez, Ph.D.
Second Supervisor: Prof. Dr. Martin Schlather
Third Supervisor: Prof. Stephan Klasen, Ph.D.

Versicherung gemäß Prüfungs- und Studienordnung

für den Promotionsstudiengang "Wirtschaftswissenschaften" der Georg-August-Universität Göttingen

1. Die Gelegenheit zum vorliegenden Promotionsvorhaben ist mir nicht kommerziell vermittelt worden. Insbesondere habe ich keine Organisation eingeschaltet, die gegen Entgelt Betreuerinnen und Betreuer für die Anfertigung von Dissertationen sucht oder die mir obliegenden Pflichten hinsichtlich der Prüfungsleistungen für mich ganz oder teilweise erledigt.
2. Ich versichere, dass ich die eingereichte Dissertation "Advances and Applications of Experimental Measures to Test Behavioral Saving Theories and a Method to Increase Efficiency in Binary and Multiple Treatment Assignment" selbstständig und ohne unerlaubte Hilfsmittel verfasst habe; fremde Hilfe habe ich dazu weder unentgeltlich noch entgeltlich entgegengenommen und werde dies auch zukünftig so halten. Anderer als der von mir angegebenen Hilfsmittel und Schriften habe ich mich nicht bedient. Alle wörtlich oder sinngemäß den Schriften anderer Autoren entnommenen Stellen habe ich kenntlich gemacht.
3. Die Richtlinien zur Sicherung der guten wissenschaftlichen Praxis an der Universität Göttingen werden von mir beachtet.
4. Eine entsprechende Promotion wurde an keiner anderen Hochschule im In- oder Ausland beantragt; die eingereichte Dissertation oder Teile von ihr wurden nicht für ein anderes Promotionsvorhaben verwendet.
5. Des Weiteren ist mir bekannt, dass Unwahrhaftigkeiten hinsichtlich der vorstehenden Erklärung die Zulassung zur Promotion ausschließen bzw. später zum Verfahrensabbruch oder zur Rücknahme des erlangten Titels berechtigen.

Datum und Unterschrift:

Acknowledgements

I would like to express my gratitude to my first supervisor Marcela Ibañez. She was eager to engage in this unusual project and introduced me to the field of behavioral economics. Marcela constantly provided inspiring and valuable advice, reacted flexible and open-minded to new ideas, and served as a critical sparring partner for all my work. She has greatly supported and motivated me over the last years.

I am also grateful to my long-term second supervisor Martin Schlather. His curiosity and sincere interest in combination with his ability to connect people and ideas lead to this doctoral project. Martin commented thoroughly and (up to incredibly) timely on countless draft versions and kept pushing, supporting, advising and motivating me ever since I have known him.

I want to thank Gerhard Riener, without whom three of the four chapters of this thesis would not exist. I appreciate his motivation, enthusiasm, his ideas and his support, not only by being a co-author, but also by being a valuable and pleasant contact person.

I thank Stephan Klasen for being part of my thesis committee. I acknowledge his tremendous contribution in establishing the research environment in Development Economics in Göttingen and the RTG Globalization and Development, from which I both profited substantially.

First, by serving as an inspiring platform for academic exchange, where Stephan Klasen provided valuable feedback in plenty research seminars, and second, by giving me the opportunity to meet many lovely persons in this (extended) environment. I am grateful for all these encounters and the time spent together—inside and outside the University.

I acknowledge funding from the Foundation of German Business for my PhD studies, as well as the Chair of Development Economics, the Faculty of Economic Sciences and the RTG Globalization and Development for further financial support and travel funding, which enabled me to participate in international academic exchange.

I want to express my warmest gratitude to Nicole for her constant empathy and invaluable support in every dimension of my life. I also want to give my dearest thanks to my family, Christoph Schneider, Eva Schneider with Imanuel and Karl-Heinz Schneider with Ursel for their encouragement and tireless support, where tireless also refers to the timing of my support requests. They shaped my personality, my values and attitudes which I deem the foundation of every success I have ever had and will ever have.

Contents

List of Figures

List of Tables

List of Abbreviations

AIC	**A**kaike **I**nformation **C**riterion
A&P	**A**rrow-**P**ratt, measure of risk aversion (Pratt, 1964)
B	**B**owman, Minehart, and Rabin (1999), definition of loss aversion
BMI	**B**ody **M**ass **I**ndex
C&E	**C**rainich and **E**eckhoudt (2008), measure of prudence
COP	**C**olombian **P**eso
CV	**C**ross **V**alidation
DANE	***D**epartamento **A**dministrativo **N**acional de **E**stadística*, Colombia's Statistical Department
EU	**E**xpected **U**tility **T**heory
HH	**H**ouse**h**old
KT	**K**ahneman-**T**versky (1979), definition of loss aversion
KW	**K**öbberling-**W**akker (2005), definition of loss aversion
MSE	**M**ean **S**quared **E**rror
N	**N**eilson (2002), definition of loss aversion
OLS	**O**rdinary **L**east **S**quares regression
SISBEN	***SIS**stema de Identificación de Potenciales **BEN**eficiarios de Programas Sociales*, Identification System for Social Health Insurance
UPZ	***U**nidades de **P**laneamiento **Z**onal*, Local Planning Unit in Bogota, Colombia
USD	**U**nited **S**tates **D**ollar
WT	**W**akker-**T**versky (1993), definition of loss aversion

Chapter 1

Introduction and Overview

1.1 Overview

The experimental approach has gained tremendous attention and importance in economic research over the last decades (Heckman and Smith, 1995; Harrison and List, 2004; Falk and Heckman, 2009; Levitt and List, 2009; Banerjee and Duflo, 2010; Viceisza, 2016). Its rise in popularity is tightly connected with its suitability to establish causality in a credible way (Viceisza, 2016). Experiments are conducted to test theories or new interventions, to estimate the size of a possible impact of an intervention or to elicit characteristics. They thus contribute to scientific advance, stimulate new research and inform policy[1] (Falk and Heckman, 2009; Levitt and List, 2009; Viceisza, 2016).

This dissertation presents four independent papers that advance experimental methods and test theories on saving using experimentally elicited characteristics.

The first two papers test two behavioral theories on precautionary saving based on the characteristics *loss aversion* and *prudence* by using a combination of survey data and experimental measures of these preferences from a sample of poor households in Bogota. In that context, the second paper presents a new method to elicit risk preferences with their intensities. The last two papers make a methodological contribution to binary and multiple treatment group assignment in experiments by presenting a new theoretically derived method with a simulation study and a software implementation for its convenient application.

1.2 Testing Theories on Precautionary Saving Using Experimental Measures

Uninsured risk is a cause of poverty, since without insurance, shocks can result in poverty (Dercon, 2010). One prevalent risk for many people, especially in low-income urban areas in Latin America, is income risk due to unemployment, as informality levels are high and formal working contracts are rather an exception (Loayza, Servén, and Sugawara, 2009). Saving is a risk coping strategy that is available even in absence of or with only restricted access to insurance or credit markets (Dercon, 1996, 2010); a setting often found in developing contexts. Yet, in low-income economies, the propensity to save is only about half as large as in high-income economies (World Bank, 2014). This is additionally consequential, since the lack of accumulated capital may lead people to apply risk management

1. Even for educational purposes, experiments have been used (McPeak, Chantarat, and Mude, 2010).

strategies such as engaging in low-risk low-return projects, which implies permanently low incomes: People are stuck in a poverty trap (Dercon, 1996). Thus, a failure in saving can result in vulnerability to poverty.

The first two papers address this issue by advancing our understanding of how individual preferences may affect the saving process. These insights may be used for policy interventions to increase the saving rate among those that would benefit from it, namely the poor and those exposed to otherwise uninsured risks.

Income Risk, Precautionary Saving, and Loss Aversion - An Empirical Test
This first paper is joint work with Marcela Ibañez and studies the empirical and theoretical connection between loss aversion, income uncertainty and saving. More specifically, we investigate predictions derived from the theoretical framework of Kőszegi and Rabin (2009) with a sample of the poor population in Bogota, Colombia, where loss aversion is elicited experimentally.

Kőszegi and Rabin (2009) derive from their model of reference-dependent consumption—which bases on prospect theory (Kahneman and Tversky, 1979)—that loss-averse individuals increase their savings as a response to income uncertainty. According to their model, for loss-averse individuals, the disutility from lowering consumption is larger than the utility from increasing consumption by the same amount relative to the expected level. As a result, uncertainty induces a precautionary saving motive, since for a higher level of saving, the decrease in consumption relative to the expected level connected with a negative income shock is lower and thus its impact is mitigated. We extend the analysis in Kőszegi and Rabin (2009) with respect to the degree of loss aversion and derive the hypothesis that the just explained increase in savings is larger, the larger the degree of loss aversion.

We test these hypotheses with a sample of the low-income population in Bogota, Colombia, as this group is exposed to informality and would profit from interventions to increase savings as outlined above. We apply different definitions of loss aversion to experimentally elicited utility points on the individual level using the procedure proposed by Abdellaoui, Bleichrodt, and Paraschiv (2007). Data on saving figures and household characteristics is obtained from a survey with the same population and income uncertainty is measured as the official unemployment rate or aggregated self-reported unemployment risk in the area an individual is living in.

We advance the literature by empirically and theoretically investigating the relation between loss aversion and saving in presence of income uncertainty, using individual and incentivized measures of loss aversion. In particular, our study is the first to test the loss-aversion based precautionary savings motive proposed by Kőszegi and Rabin (2009).

We find a substantially lower proportion of loss-averse individuals in our sample compared with student samples. The empirical analysis is consistent with the predictions of the model by indicating that income uncertainty is related to an increase in savings for loss-averse individuals. This precautionary motive for saving increases with the degree of uncertainty. Moreover, the stronger the degree of loss aversion, the higher the savings. These findings establish that the relation between loss aversion and saving can be positive and is thus more complex than previously assumed.

The findings in this chapter base on correlations, however, and with the data used, causality cannot be established without any doubt. Future research could fill this gap with the use of panel data or a lab-experiment. Moreover, the role of loss aversion with respect to saving deserves more attention.

A possible intervention to increase savings using the insights of this chapter could illustrate potential 'losses' connected with income shocks, for example by picturing a loss in the standard of living. Moreover, the uncertainty of income in informal settings could be highlighted.

Higher Order Risk: An Application To Savings of the Poor in Bogota

This chapter, which is joint work with Marcela Ibañez and Gerhard Riener, investigates the relationship between prudence, income uncertainty and saving. Prudence can be defined as the preference to accept risk at higher levels of wealth instead of lower levels of wealth.

Extending the work by Leland (1968), we show that—in addition to prudent risk-averse individuals—also prudent risk-loving individuals react to income uncertainty by increasing their saving rate proportionally to the degrees of income uncertainty and prudence. Previously, this precautionary motive for saving was established also for prudent risk-loving individuals (Crainich, Eeckhoudt, and Trannoy, 2013), but leading to the prediction that they would save their whole income as a reaction to income uncertainty. Moreover, we show that the measure of intensity of prudence advocated for by Crainich and Eeckhoudt (2008)—which is independent of risk aversion—can directly be linked to this precautionary motive for savings, whereas the commonly used measure by Kimball (1990) can only be used to indicate the degree of prudence for risk-averse individuals.

To test the predictions generated by this model, we present a new experimental method to elicit higher order risk preferences. This method uses a non-parametric estimation of the utility function using P-splines to connect utility points that are elicited with the procedure proposed by Abdellaoui, Bleichrodt, and Paraschiv (2007). By its non-parametric nature, it is flexible enough to reflect any combination of risk preferences that is observed empirically. This method allows the computation of the abovementioned, theoretically derived measures of the intensities of prudence and risk aversion, which is a novelty.

We apply this method to measure prudence in a low-income sample from Bogota, Colombia. Income uncertainty is measured by working sector dynamics as expressed in the ratio of closed to existing businesses in Bogota in 2013. This rate of firm closures of the working sector an individual usually is working in is assigned as individual income risk. Savings and other household characteristics are obtained from a survey with the same sample.

We find comparable results to earlier studies with respect to the classification of individuals as prudent or imprudent among a sample of poor households in Bogota. In addition, the results strongly support the theoretical prediction that uncertainty leads to increases in savings for prudent individuals—even when we pool risk-averse and risk-seeking individuals, which supports our extension.

It would be fruitful to see a comparison of our proposed method for the elicitation of prudence with the conventionally applied methods that base on risk-apportionment tasks with compound lotteries as introduced by Deck and Schlesinger (2010). Moreover, it would be interesting to investigate the predictions from our model in a laboratory experiment to establish causality and confirm the findings in this chapter. Lastly, it could be interesting to empirically compare the prudence-based precautionary motive for saving with the one that relies on loss aversion as in the first paper of this dissertation.

The results obtained in this paper suggest that the population group under study lacks alternative options to smooth consumption. Thus, one policy implication of this research is to provide consumption smoothing devices at an affordable cost without introducing an additional risk of indebtedness. Moreover, pension or saving campaigns might stress the uncertain nature of income and make the use of saving salient by illustrating that risk in a state of higher wealth is less consequential than when wealth is low.

1.3 A New Method for Binary and Multiple Treatment Assignment

A striking drawback of field experiments are their costs and the difficulty of implementation (Deaton, 2010; Al-Ubaydli and List, 2015; Viceisza, 2016). One of the difficulties in implementation that directly affects costs is connected with treatment assignment. For example, when assigning cluster to treatment groups using a matching approach, the drop-out of one single cluster from the experiment may leave the whole study underpowered and an effect that otherwise could have been found might not be verified anymore. Similarly, when units have to be assigned to multiple treatment groups, pure randomization might lead to some groups in which desired characteristics are not present at all. This is disadvantageous if the experiment's purpose lies in studying heterogeneous treatment effects connected with the missing characteristics.

Moreover, comparable experimental groups, which can be obtained through appropriate treatment assignment, are credited for allowing an efficient estimation of average treatment effects (Kallus, forthcoming (2017)). A failure in assigning units such that treatment groups are comparable might thus result in additional costs in terms of an increased sample size that is necessary to detect a given treatment effect.

Unbalanced treatment groups are also a threat to the validity of the outcome of any experiment—in the lab or in the field (Fisher, 1935). An experiment might on average yield the true result when performing random treatment assignment, but any given single realization might lead to wrong conclusions. This is of particular interest with the increasing attention on replicability of experiments (Camerer et al., 2016). The last two papers of this dissertation address these issues by providing a solution to the mentioned problems.

The min Mean Squared Error Treatment Assignment Method

The third paper of this dissertation is joint work with Martin Schlather. We present a new approach to treatment assignment in (field) experiments for the case of one or multiple treatment groups. This procedure—which we call the minimizing Mean Squared Error (MSE) Treatment Assignment method—uses sample characteristics to obtain balanced treatment groups and is particularly suited for multiple treatment assignment or when attrition might be a concern. The information used for treatment assignment can be multivariate, continuous and may consist of any number of variables.

We implement the idea for conducting treatment assignment of Kasy (2016), allowing for randomness and using a frequentist approach, thereby increasing its applicability by eliminating the need of specifying any parameter value. Moreover, as we show, this leads to 'more equal groups' as compared to the suggested implementation in Kasy (2016). We present a way to extend this method to assign multiple treatment groups, making the min MSE approach the first one to rely on a theoretically-derived decision statistic allowing for multiple treatment assignment.

We show theoretical properties of this method and compare it to competing mechanisms theoretically and by virtue of simulations. In the simulation study, our proposed procedure performs better than, or at least comparably to, competing approaches, such as matching. Furthermore, the min MSE approach is attrition tolerant, offers greater flexibility and is very fast. It can conveniently be implemented and aims at balancing a generalized second moment of the covariate distribution of the treatment groups, which is important whenever the treatment effect differs across individuals (and is a function of covariates).

A possible limitation is that p-values from conventional t-tests might not be adequate as the distribution of covariates can in general not be assumed to be normal. This is the case in many settings, such as in experiments with a small sample size or even when performing randomization several times until the result is satisfactory. It would therefore be of great interest to see a software implementation for the computation of exact p-values resulting from a permutation test, sometimes referred to as Fisher's exact test, as this might make the analysis more convenient.

A clear recommendation from the simulation study is that purely random treatment assignment should be avoided if information is available that is likely to affect the treatment, as group means in baseline variables in treatment and control groups might differ considerably in any given realization of treatment assignment.

Software for min MSE Treatment Assignment

The last chapter presents a software implementation of the min MSE Treatment Assignment mechanism derived in the third paper for the statistical software Stata and explains its usage.

Although most scholars might agree that controlled treatment assignment should be preferred to completely random treatment assignment, in practice, applying a sophisticated method is often timely and complicated. For multiple treatment assignment—to our knowledge—no software implementation of any sophisticated or theory-based approach is available.

This gap is filled by the min MSE ado-package for Stata. It implements the min MSE Treatment Assignment procedure for one and multiple treatment groups as derived in the third paper using the stochastic simulated annealing algorithm (Kirkpatrick, Gelatt, and Vecchi, 1983). When implementing an experiment, researchers can use this software to conveniently assign subjects to one or multiple treatment groups based on observed pre-treatment characteristics with a single command before treatment is executed. This method is particularly suited for treatment assignment in the field, as it is very fast compared to several competing methods for binary treatment assignment. The created treatment groups are *balanced* across multiple dimensions and, within the model laid out in the third paper, the mean squared error of the estimator for the conditional average treatment effect is minimized.

The min MSE ado-package is also suited for rebalancing to preserve the sample's balance, for example in case an intervention is implemented on several days and attrition has happened on the first day(s). With the same function, sequential treatment allocation may be realized.

Chapter 2

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

This chapter is joint work with Marcela Ibañez. We thank Sonia Triviño Munévar, Lukas Hermann and Tatiana Orozco Garcia for research assistance.

2.1 Introduction

Income shocks, such as unemployment, diseases or natural disasters, are pervasive in developing countries (World Bank, 2013). Due to limitations in access to credit markets and weak social protection systems, saving is one of the few alternatives that households have to mitigate the effect of income shocks (Dercon, 2010). However, saving rates are quite low in developing countries (World Bank, 2014). Hence, from a policy perspective it is important to understand which factors influence savings. This paper contributes to answering this question by investigating the impact of uncertainty and loss aversion on savings.

Our conceptual framework is based on Kőszegi and Rabin's (2009) reference-dependent model of intertemporal consumption. Building on prospect theory (Kahneman and Tversky, 1979), this model assumes that deviations in consumption from previous expectations induce a gain-loss sensation. When individuals are loss averse, bad news about consumption changes are more unpleasant than good news are pleasant. The prospective loss from lowering expected future consumption generates a precautionary savings motive. Individuals who are loss averse save to decrease the impact of any given wealth shock on future consumption utility. The larger the frequency and magnitude of the shocks and the larger the degree of loss aversion, the more individuals save.

To empirically test the above hypotheses, we take as a case study the low-income population in Bogota, Colombia. Similar to comparable population groups in other developing countries, this population is exposed to substantial uncertainty and is highly vulnerable to poverty (Loayza, Servén, and Sugawara, 2009). The unemployment rate is relatively high and a large share of the population depends on informal employment or works as domestic employees; naturally, the savings are quite low for this population. Among our sample of low-income individuals, we find that more than 80 percent do not have any savings while 75 percent report having a day-by-day financial plan. The problem of a low saving rate is even more worrisome given that less than 10 percent of the low-income population has access to a pension in the old age (Aguila, Attanasio, and Quintanilla, 2010).

To inform policy makers about the obstacles that this population group faces in saving, we launched an independent study on the financial situation of low-income households in Bogota, Colombia in 2013. The study comprised the survey entitled 'Savings for old age in Colombia' and complementary experimental measures of risk and time preferences for the same group of individuals. This survey allows us to estimate the total value of households' assets as well as households' perceived unemployment risk. The risk experiment is based on the experimental design by Abdellaoui, Bleichrodt, and Paraschiv (2007) and aims at obtaining individual estimates of loss aversion and risk attitudes. This method has the advantage that it provides flexibility in estimating the utility function for both gains and losses, while also correcting for probability weighting. Nonetheless, it is easy for participants to understand as it is based on comparisons of binary lotteries.

In line with Kőszegi and Rabin's (2009) hypothesis, we find that loss-averse individuals increase savings when they are exposed to greater uncertainty. This finding is robust to different measures of loss aversion and income uncertainty, to the inclusion of controls on a large set of socioeconomic characteristics and to different econometric specifications. The novelty of this result is that it demonstrates that the relationship between loss aversion and savings is more complex than previously proposed. While previous studies have considered that loss aversion creates an obstacle for savings (e.g. Thaler and Benartzi, 2004), our study shows that in the presence of uncertainty, loss aversion leads to an increase in savings. The larger the degree of loss aversion, the larger is the increase in savings.

Bowman, Minehart, and Rabin's (1999) theoretical model predicts that because of loss aversion there is an asymmetric response to learning about certain future income shocks on savings—i.e. good news regarding future income leads to an immediate upward adjustment in consumption, reducing the possibility of future consumption growth. Bad news results in less than proportional adjustments of consumption, leading to a larger decrease in future consumption when the shock is realized. Empirical evidence supports this hypothesis. Using macro data, Shea (1995) and Bowman, Minehart, and Rabin (1999), showed that consumption changed more in response to anticipated positive income shocks than in response to anticipated negative shocks. Using microdata, Fisher and Montalto (2011) find that there is an asymmetric effect of positive and negative income shocks on savings. In particular, unusually low income decreased savings more than unusually high income increased them.

We make various contributions to the literature on this topic. First, this is—to the best of our knowledge—the first paper to empirically test the precautionary savings motive predicted by the reference-dependent model of intertemporal consumption by Kőszegi and Rabin (2009). Second, unlike previous studies, we use a direct and incentivized measure of loss aversion at the individual level. Hence, we contribute to this research by explicitly considering the impact of loss aversion on savings. Third, we consider the impact of income risk that is resolved in the future, and hence, do not consider the asymmetric response of savings to certain future income shocks as predicted by the model of Bowman, Minehart, and Rabin (1999). Last, we propose a new measure to capture uncertainty about future income by considering the effect of unemployment.

Our paper also addresses the literature considering behavioral approaches to increase the savings rate among the poor (for a comprehensive overview of the research on saving among the poor see Karlan, Ratan, and Zinman, 2014). In particular, a couple of interventions have considered the effect of loss aversion on savings. One of the most well known interventions is Thaler and Benartzi's (2004) 'SMarT (Save More Tomorrow™) Program'. Considering that due to loss aversion, individuals would perceive saving a portion of their current income as a loss of income, this program proposes to make saving decisions on future income increases. Karlan et al. (2016) compares the effectiveness of reminders that are framed as a loss ("your dreams won't come true") versus as a gain ("your dreams will come true") and find no significant effects of the loss versus the gain frame on a household's savings rate. However, they do not consider heterogeneous effects of the degree of loss aversion. We contribute to this literature, exploring the relationship between uncertainty, loss aversion and savings. This research could suggest alternative interventions to increase savings.

This paper is structured as follows. The next section presents the model of intertemporal consumption by Kőszegi and Rabin (2009) from which the hypotheses of the study are derived. Section 2.3 explains the empirical strategy used to test the predictions of this model and Section 2.4 explains how the different measures were obained. Results are presented in Section 2.5 and the approaches and findings of the paper are discussed in Section 2.6. Section 2.7 concludes.

2.2 Theoretical Framework

The conceptual framework that we use in the analysis is based on the reference-dependent utility model of intertemporal 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 in Kőszegi and Rabin (2009) using this model to derive hypotheses relating the strength of the precautionary saving motive with the degree of loss aversion.

2.2.1 Kőszegi and Rabin (2009): Model of Consumption and Saving

The model considers a two-period consumption-savings decision problem where individuals face uncertainty regarding their future wealth. Here, we only present the model for the case in which wealth, W, is is a binary random variable and and uncertainty is resolved in the second period.[1] 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 uncertainty.[2]

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 not 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.[3]

The expected instantaneous utility in the second period, $\mathbb{E}[u_2(c_2)]$, depends on the expected utility of consumption in that period and on the so-called 'gain-loss utility'. The 'gain-loss utility' reflects utility gains or losses due to changes in 'beliefs' about consumption levels after uncertainty is resolved compared to

1. In Appendix A.1.1, we present the two-period model for a more general case.
2. Results generalize to non-binary random income; see the corresponding Proposition 8 in Kőszegi and Rabin (2009), as well as the Proposition and its Corollary in this study, for more general results.
3. 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).

the 'beliefs' before uncertainty was resolved.[4] If new beliefs imply a higher consumption level than previously planned, this results in a utility gain. On the other hand, if consumption according to new beliefs is lower than previously planned, the feeling of a loss lowers utility. Following Kahneman and Tversky (1979), it is assumed that individuals weight losses relative to the believed level of consumption by a factor λ that captures the degree of loss aversion. For an individual who is loss averse, we have $\lambda > 1$, whereas for a gain-seeking individual, $\lambda < 1$.

Hence, the expected instantaneous utility of 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 $c_2^- = W_0 - c_1 - s$. 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 $\eta > 0$. Contrarily, 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.[5]

For an interior solution, the optimal consumption path satisfies the following condition:[6]

$$m'(c_1) = \frac{1}{2}m'(c_2^+) + \frac{1}{2}m'(c_2^-) + \frac{1}{4}\eta(\lambda - 1)[m'(c_2^-) - m'(c_2^+)]. \tag{2.1}$$

To see whether or not $m'(c_1)$ increases when uncertainty s increases (then c_1 decreases if m is strictly concave), we apply a Taylor approximation of the right

4. Beliefs result from 'credible consumption plans' that specify possibly stochastic consumption levels for each period, which the agent forms before the first period starts. 'Credible' plans only feature consumption levels, from which the individual would not deviate later on. Details on this concept are given in Appendix A.1.2 or in Kőszegi and Rabin (2009).

5. 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 could not exclusively validate this assumption (e.g. Schmidt and Traub, 2002). Therefore, we only assume $\lambda > 0$ and allow for 'gain-seeking' behavior. See Appendix A.1.1 for further details.

6. See equation (10) in Kőszegi and Rabin (2009).

hand side of (2.1) around $s = 0$ to find[7]

$$m'(c_1) \approx m'(c_2) + \frac{1}{2}m'''(c_2)s^2 + \frac{1}{2}\eta(\lambda - 1)(-m''(c_2))s. \tag{2.2}$$

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 $m''' > 0$ or the last term dominates the second term in (2.2). The first condition corresponds to the classical theory of precautionary saving, as initiated by Leland (1968), where 'prudence', defined as the preference for allocating a zero-mean risk to a state of higher wealth instead of to a state of lower wealth, causes the individual to save. The latter condition is, according to Kőszegi and Rabin (2009), "technically speaking" true only for a small amount of risk. They argue, however, that m is a global utility function and therefore a small risk in the model can still be substantial in "practical terms". Using this interpretation of small risks, and generalizing to cases where second-period income has more than just two realizations and where people might overconsume in the first period, we can derive the first hypothesis from this model.[8]

Hypothesis 1. *With loss aversion, the effect of uncertainty on savings is unambiguously positive.*

2.2.2 The Effect of the Degree of Loss Aversion in the Two-period Model by Kőszegi and Rabin (2009)

We now extend the analysis in Kőszegi and Rabin (2009) regarding the degree of loss aversion. From both, (2.1) and (2.2), we see that savings increases in the degree of loss aversion.[9] This finding can be generalized to non-binary income risk and individuals overconsuming in the first period; the latter case is linked to a parameter $\gamma \geq 0$ in this more general framework (Appendix A.1.1), where an individual increases consumption in the first period relative to the ex-ante optimal level if $\gamma < 1/\lambda$; see Appendix A.1.1 for further details.

Similarly to Kőszegi and Rabin (2009), we call a 'personal preferred equilibrium' the preferred plan among those that are credible before the first period starts.

7. See equation (11) in Kőszegi and Rabin (2009).
8. See Proposition 8 in Kőszegi and Rabin (2009).
9. Although this effect is independent of the amount of risk on the right-hand side of (2.2), we have to keep in mind that this expression is a Taylor approximation around $s = 0$. Thus it is a good approximation only for small amounts of risk. Although the first-order condition (2.1) holds independent of the amount of risk, technically, for large amounts of risk, it cannot be certain that this condition yields a utility maximum. For small amounts of risk, however, the second-order condition for a utility maximum is satisfied.

Proposition 1. *Suppose wealth is equal to* $W_0 + sy$*, where* y *is a non-deterministic mean-zero lottery that is resolved in period 2. 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* $dc_1/d\lambda < 0$*.*

The proof of Proposition 1 is in Appendix A.1.3. From this proposition, we derive the following hypothesis:

Hypothesis 2. *With income uncertainty, the effect of the degree of loss aversion on savings is unambiguously positive. 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 having a positive effect on savings, from (2.1) and (2.2), we see that the effect of loss aversion on saving increases in uncertainty and that the effect of uncertainty on saving increases in loss aversion. As expected, this result generalizes to non-binary income lotteries and individuals overconsuming in the first period:

Corollary 1. *Suppose wealth is as specified in Proposition 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* $d^2 c_1/(ds d\lambda)|_{s=0} < 0$*.*

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

Hypothesis 3. *The effect of the degree of loss aversion is an increasing function of uncertainty. Equivalently, the effect of uncertainty is an increasing function of the degree of loss aversion. As in Hypothesis 2, this also includes coefficients of loss aversion* $\lambda \leq 1$*.*

2.3 Empirical Strategy

We want to test the hypotheses derived from Kőszegi and Rabin's (2009) model and investigate the relationship between income uncertainty, loss aversion and the joint effect of both on the individual level of savings using individual measures of savings, uncertainty and loss aversion.

Hypothesis 1 To test Hypothesis 1, i.e. if the effect of income uncertainty on the saving decision for loss averse individuals is unambiguously positive, we run the following regression:

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

where Savings_i is the first period income minus c_1 (first period consumption), 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 corresponding to income uncertainty and socioeconomic characteristics, respectively, and β_0 is the intercept of this model.

The data would support Hypothesis 1 if $\beta_1 > 0$.

Hypothesis 2 To test the second hypothesis, postulating that the higher loss aversion, the higher saving when facing income uncertainty, we run the following regression:

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

where λ_i is the degree of loss aversion of individual i with corresponding regression coefficient β_2 and Savings_i, X_i, ζ, β_0 and ε_i are defined as before. We run this regression for the complete sample considering that this population group is highly exposed to income uncertainty. A positive β_2 would support Hypothesis 2.

Hypothesis 3 Finally, we test Hypothesis 3, claiming that the effect of uncertainty on saving 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, and Savings_i, X_i, $\beta_0, \beta_1, \beta_2, \zeta$ and ε_i are defined as before. Hypothesis 3 is supported if $\beta_3 > 0$.

Note that we center income risk measures around mean values and loss aversion around 1. 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. We report marginal estimations following from this regression for the effects of loss aversion and income uncertainty at different levels of income uncertainty and loss aversion, respectively, which allows these results to be interpreted within the contexts of Hypotheses 1 and 2.

The definitions of loss aversion, savings and income uncertainty, as well as the methods used to estimate them are presented in Section 2.4.

2.4 Data and Definition of Variables

The data used in the study was collected between October and November 2013 as part of the project 'Savings for the Old Age'. The objective of the study was to investigate the expectations of the poor population living in Bogota, Colombia regarding their old age and their degree of preparedness for it. The study was comprised of an extensive survey of the financial situation of the households and incentivized economic experiments on risk and time preferences.

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 compromised of these low-income households. Enumerators visited the randomly selected households and confirmed that they belonged to the target group. If it was not the case, enumerators visited their neighbors, assuming that they shared similar socioeconomic characteristics. The criteria for selecting participants was that they should be beneficiaries of the social health insurance, SISBEN. This condition would guarantee that the participants were from a 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 few days later and took about 20 minutes.

In the analysis, we consider the relation between savings, income risk and loss aversion. Below, we define the measures used and explain how the variables were calculated.

2.4.1 Savings

We measure savings as the value of an individual's assets in the analysis. 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.

2.4.2 Income Uncertainty

In the analysis, income uncertainty is measured by the unemployment rate aggregated at the locality level or the neighborhood level. We use this measure since unemployment is one of the main sources of income risk that our population group is confronted with. In addition, unemployment is quite high in Colombia. DANE estimated the unemployment rate at 8.64 percent in Bogota for 2013.

The advantage of this measure is that this variable can be considered to be exogenous for a single individual who cannot affect the unemployment risk. The assumption that we use is that individuals observe when neighbors lose employment, which makes their individual risk of income loss salient.

In the analysis, we use two different sources of information to compute unemployment figures as explained in greater detail below.

Income Uncertainty Based on Data from DANE This measure corresponds to the unemployment rate aggregated at the locality level (*Localidad* in Spanish). We use the multipurpose survey from Bogota (*Encuesta Multipropósito para Bogotá*) for the source of information to construct the unemployment data at the level of locality, which was conducted by Colombia's Statistical Department (DANE) and the District Planning Department (SDP). The figures correspond to the 2011 unemployment rate[10] published in *'Índices de Ciudad'* by the *Secretaría de Planeación, Alcaldía Mayor de Bogotá*. The unemployment rate ranged from 6.86 percent in the *Localidad 'Suba'* to 11.31 percent in *'San Cristobal'*.

Average Unemployment Risk Based on Survey Data All individuals that took part in the experiment and that were working at the time of the interview answered the question, "What is the probability that you will lose your job next year". All those stating a positive probability where asked to also answer the question: "If you lose your job, what is the probability of finding a new one next year?" Since the vast majority were confident of being able to find a new job within a year in the case of unemployment, and since this new job could potentially be better paid than the last one, we interpret the unemployment risk as income uncertainty and not just as the risk of a negative income shock.

In order to reduce problems of endogeneity associated with this measure, we use the average probability at the local planning unit, *UPZ*. This measure has lower aggregation than the measure based on *Localidad* but a larger degree of

10. The figure coming closest to what is commonly referred to as 'unemployment' is the *'Tasa de Desocupadas'*, the rate of unoccupied persons.

aggregation than *Barrios*.[11] Typically, a *Localidad* consists of several *UPZ*; for example, the *Localidad Suba* consists of 12 *UPZs*.

Using the average unemployment risk has two practical advantages. First, self-employed individuals cannot lose their job but they can be exposed to income risk. By using averages, we can assign a level of income uncertainty that is likely to be close to the reality of those individuals. Furthermore, the income risk an individual is exposed to might not only be driven by his or her own income risk, but also by the income risk their partner is exposed to as an example.

2.4.3 Loss Aversion

For the experimental elicitation of loss aversion, we used the non-parametric method introduced by Abdellaoui, Bleichrodt, and Paraschiv (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. Hence, it is cognitively equally demanding as other methods, e.g. Holt and Laury (2005). The method applied corrects for the misperception of probability. In addition, using these choices, it estimates the utility points iteratively over a large range of values. The utility points can 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 mostly focus on the elicitation of preferences at just one or a few (arbitrarily) selected points in the interval of interest (e.g. Holt and Laury, 2005; Binswanger, 1980; Eckel and Grossman, 2002).

The definitions of loss aversion we apply build on utility functions. Following Abdellaoui, Bleichrodt, and Paraschiv (2007), we use five different definitions of loss aversion, as there is no agreed upon definition of loss aversion so far, and measures differ considerably. In the general version of the model introduced in Section 2.2, that is presented in Section A.1.1), loss aversion enters into the instantaneous utility via the "universal gain-loss utility" μ.

Kahneman-Tversky (KT) Kahneman and Tversky (1979) define an individual as loss averse, if for all amounts of money x the utility 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 $-\mu(-x)/\mu(x)$

11. *UPZs* with less than 25 observations were grouped with their neighboring *UPZ(s)*.

for every elicited amount $x > 0$. If $\mu(-x)$ for any of these 8 elicited amounts of money $x > 0$ was not elicited, it was linearly interpolated. Then, we took the median of these thus computed coefficients as the coefficient of loss aversion.

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 to apply 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 thus computed coefficients.

Bowman (B) Bowman, Minehart, and Rabin (1999) propose to perform 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 definition of loss aversion $\mu'(0_-)/\mu'(0_+)$ was transformed into a coefficent of loss aversion by computing the ratio of slopes connecting 0 with the elicited utility points that are closest to 0 on both domains.

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 based on data availability. 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. The second meta measure additionally includes the coefficient based on the definition by Kahneman and Tversky (1979), which we can compute for 579 participants of the experiment. Finally, we compute a measure relying on all five definitions. 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.

2.4.4 Risk Preferences

The experimental data was also used to estimate risk preferences corresponding to the curvature of consumption utility m and probability weighting that we may use as control variables. Although in the theoretical model, m is assumed to be a concave function, for empirical elicitation of curvature, we relax this assumption. We use similar procedures as Abdellaoui, Bleichrodt, and Paraschiv (2007). First, we re-scaled monetary amounts x to lie in the interval $[-1, 1]$. Second, we rescaled the utility function such that $m(-1) = -1$, $m(0) = 0$ and $m(1) = .5$, consistent with the elicitation method. Taking into account the rescaling of monetary amounts and the utility, in the third step we estimate the curvature of utility, estimating the following power utility function for the gain and loss domains:

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

Finally, the estimated parameters, referred to as coefficients of risk aversion in expected utility (EU) theory, are used to classify utility curvature. 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 $0 < a < 1$. For more details see Appendix A.2.1.

The experimental method proposed by Abdellaoui, Bleichrodt, and Paraschiv (2007) also allows a non-parametric estimation of subjective probabilities. Using this method, we estimate the objective probability corresponding to a subjective probability of 50%.

2.4.5 Time Preferences

To elicit the near future impatience or interest rate, which we also use as control variables, we followed the experimental design by Andersen et al. (2008). People were asked whether they would prefer receiving 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 people normally switched 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, but since people are making choices dealing 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. In that setting, participants had to choose between receiving the lower amount in 180 days or receiving the higher amount in 210 days. This allows us to consider consistency in intertemporal choice. We compute the difference between interest rates or impatience for the two time frames. If people behaved consistently, we would expect impatience to receive a monetary amount 30 days earlier to be unaffected by the timely distance in the future. If people became more patient, impatience in the more distant future should be lower.

2.4.6 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:

- Age has been found to affect savings positively (e.g. Conley and Ryvicker, 2004; Finke, Huston, and Sharpe, 2006) but also negatively (Devaney, Anong, and Whirl, 2007).
- Female headed households tend to accumulate less wealth (Conley and Ryvicker, 2004) and are less likely to have longer term saving motives such as retirement savings, instead of short term saving motives such as emergency savings (Devaney, Anong, and Whirl, 2007).

- The number of children in a household decreases the likelihood of holding assets (Sanders and Porterfield, 2010; Fisher and Montalto, 2011) and also the amount of wealth accumulated (Conley and Ryvicker, 2004), but not necessarily the amount of conditional investment in assets (Sanders and Porterfield, 2010).

- Family size has been found to increase the likelihood of saving for safety and security reasons, such as emergencies or retirement, instead of saving for basic needs and safety reasons, respectively, in a hierarchy model of saving motives (Devaney, Anong, and Whirl, 2007).

- Marriage and inheritance have been reported to positively affect wealth (e.g. Conley and Ryvicker, 2004).

- Homeownership increases the odds of saving and the amount of savings (Finke, Huston, and Sharpe, 2006; Fisher and Montalto, 2011).

- Education and income are positively associated with higher wealth (Finke, Huston, and Sharpe, 2006; Conley and Ryvicker, 2004; Fisher and Montalto, 2011).

- Van Rooij, Lusardi, and Alessie (2012) report a positive effect of financial literacy on accumulated savings. In the survey we asked 18 questions on financial literacy related to topics such as interest rate, asset classes, basic math and financial math. The variable included in the regression contains the number of correctly answered questions.

- Interestingly, health status has been reported to have a positive influence on the likelihood of saving (Fisher and Montalto, 2010, 2011).

- Short-term planning and saving horizons, sometimes referred to as time preference for the present, have been found to negatively affect the likelihood of saving, but also net wealth (Devaney, Anong, and Whirl, 2007; Fisher and Montalto, 2010, 2011).

- The size of the safety net according to the number of individuals available for financial help can indicate the possibilities of households to cope with income shocks. We expect that as more individuals have access to a safety net, they would decrease their savings rate.

TABLE 2.1: Summary Statistics

	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 adolescent household members	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^a	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 2.2: Summary Statistics of Income Risk Measures

	Mean	s.d.	Min	Max	Obs.
Regional Unemployment Rate (in pc)	8.5	1.7	6.9	11.3	640
Regional Unemployment Risk (in pc)	24.7	6.0	15.2	36.5	640

2.5 Results

2.5.1 Descriptive Statistics

Summary statistics are reported in Table 2.1. 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. The education level of the sample was quite low even in a developing context. 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 related topics correctly.

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.

Around 85 percent of the sample does not engage in saving money and the overall mean of savings is less than the per capita household average monthly income of 256,000 COP—approximately 130 USD. 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%) reported carrying out their financial planning on a day by day basis and more than half of the sample never, or hardly ever, exercises, which is reflected in a mean BMI of 25.7, corresponding to an overweight person.

Summary statistics on income uncertainty are reported in Table 2.2. Secondary data reveals an average unemployment rate of 8.5%. This is substantially lower than the mean perceived risk of unemployment based on subjective data (25%). This difference can be due to a high degree of pessimism for the future or simply the commonly observed overweighting of small probabilities. The Pearson's correlation coefficient for the two measures of unemployment is $r = 0.505$ with

TABLE 2.3: Summary Statistics of Experimental Measures

	Mean	s.d.	Median	IQR	Obs.
Single Measures of Loss Aversion					
Bowman (B)	0.1	0.2	0.0	0.0, 0.0	564
Kahneman-Tversky (KT)	1.1	2.7	0.4	0.1, 1.1	579
Köbberling-Wakker (KW)	10.9	76.6	0.2	0.0, 1.0	640
Neilson (N)	0.2	0.5	0.0	0.0, 0.1	640
Wakker-Tversky (WT)	12.3	110.9	0.1	0.0, 0.3	564
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

$p < .0001$, which indicates a positive and large correlation according to Cohen's classification (see e.g. Cohen, 1992). Since the number of *Localidades* is limited, one could interpret the regional unemployment rate as an ordinal variable. In that case, the more appropriate Spearman's rank correlation coefficient is $r_s = 0.4886\ (p < .0001)$.

Experimental measures of the time and risk preferences are reported in Table 2.3. The measure of time preference indicates the mean annual interest rate r demanded to receive an amount $x \cdot (1 + r/12)$ in 60 days instead of an amount x in 30 days. This mean annual interest rate is valued at 29.6 percent. On average, the interest rate, or mean impatience, stays approximately constant when the timing of receiving the monetary amounts changes from 180 to 210 days, as indicated by the increase in patience over time, expressed in percentage points. The observed impatience is in line with estimates from recent experiments that used the general population in Denmark, which seem best suited for comparison since our data also uses a general population (Harrison, Lau, and Williams, 2002). However, individual estimation of time preferences when accounting for risk preferences is not possible with our data due to the limited amount of data points for a single individual. Then, our estimates are to be compared to the estimates assuming 'risk-neutrality'.

In Prospect Theory, the attitude towards risk can be expressed by the curvature of the utility function and probability weighting. The reported measure of utility curvature corresponds to the parameter of a power utility function, as explained in Section 2.4.4. On the gain domain, the median subject exhibits a concave utility curvature, indicated by a median curvature parameter of 0.7, which corresponds to risk aversion in expected utility theory settings. On the loss domain, the median curvature parameter is 1.1 and again indicates concave utility curvature, which corresponds to slightly risk averse behavior in an expected utility framework. For gains, subjects tend to overweight probabilities around 40%, as reflected by a median value of $p_g = 40.6\%$ s.th. $w^+(p_g) = 1/2$. In lotteries involving losses, probabilities around $p_l = 80\%$ are underweighted by the median individual in our sample, since $w^-(p_l) = 1/2$.

It has generally been found that large probabilities are underweighted, whereas smaller probabilities, up to around 33%, are overweighted, which results in an S-shaped probability weighting function (e.g. Tversky and Kahneman, 1992; Camerer and Ho, 1994; Gonzalez and Wu, 1999; Abdellaoui, 2000). Probability weighting in our study, however, seems to be more pronounced than what has been found in the literature with comparable methodology (e.g. Abdellaoui, 2000): Probabilities around 40% for the gain domain are still overweighted in our study. Similarly, probabilities of about 80% for the loss domain are more underweighted than, for example, in the study by Abdellaoui (2000), who find $w^-(78\%) = 2/3$. These probabilities were elicited for the larger outcome in a lottery in absolute terms; therefore, these findings could be due to the optimism of the Colombian people. With respect to utility curvature, our findings are in the range of previous results: For the gain domain, Abdellaoui, Bleichrodt, and Paraschiv (2007) report a median coefficient of utility curvature of 0.71, although less heterogeneity. Others have found less pronounced curvature (Schunk and Betsch, 2006; Booij and Kuilen, 2009). Etchart-Vincent (2004) report a median coefficient of 0.97 for the loss domain, although most reported coefficients are lower (e.g. Schunk and Betsch, 2006; Abdellaoui, Bleichrodt, and Paraschiv, 2007; Booij and Kuilen, 2009).

Table 2.3 shows summary statistics for the different measures of loss aversion applied in this study. 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.

Putting the coefficients of loss aversion resulting from our experiment into context with other experimental results is less straightforward, since few studies use the same definition of loss aversion with the exception of Abdellaoui, Bleichrodt, and Paraschiv (2007). They find higher mean and median values for all definitions.[12] Other studies focusing on monetary or health outcomes have found mean or median values between 1.43 and 4.8, relying on different definitions of loss aversion (e.g. Fishburn and Kochenberger, 1979; Tversky and Kahneman, 1992; Bleichrodt and Pinto, 2002; Schmidt and Traub, 2002; Pennings and Smidts, 2003; Booij and Kuilen, 2009).

In summary, the coefficients of loss aversion in this study are considerably lower than in other studies. This can be explained by a lower share of loss aversion and a higher share of gain-seeking behavior in our experiment. We elaborate on these characteristics in Section 2.6 and discuss how this could affect our results.

2.5.2 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

12. Loss aversion coefficients based on the definitions by Bowman, Minehart, and Rabin (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 even 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. This indicates that the loss aversion coefficients below 1 are not just the result of our sample, but are observed in other studies as well. Furthermore, and also in our experiment, the lowest mean and median values for the loss aversion coefficients are based on the definitions by Neilson (2002) and Bowman, Minehart, and Rabin (1999); and similarly, the mean and median values of the coefficients based on the definition of loss aversion by Köbberling and Wakker (2005) are amongst the highest in our experiment.

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.[13]

The distribution of the saving amounts in our sample is skewed, and values are reported repeatedly and usually are 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 uncertainty, 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 saving 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 saving is estimated. The second model estimates the amount of savings once the hurdle is passed, using a Truncated Negative Binomial model. In Appendix A.3.1, we discuss alternative models and their suitability in this context.

Following Grogger and Carson (1991) we compute marginal effects of loss aversion and uncertainty 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

$$\frac{\partial \mathbb{E}(Y_i|X_i)}{\partial X_{ih}} = \frac{\partial}{\partial X_{ih}}[\mathbb{E}(Y_i|X_i, Y_i > 0)][1 - F(0)] + \mathbb{E}(Y_i|X_i, Y_i > 0)\frac{\partial}{\partial X_{ih}}[1 - F(0)], \tag{2.3}$$

13. 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.

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 that 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.

2.5.3 Empirical Results

The estimated coefficients for Model 1, Model 2 and Model 3 are presented in Table 2.4. Model 1b reports coefficients for an estimation with the loss-averse subpopulation only. The columns labeled 'DANE' present the results for the measure of uncertainty based on secondary information obtained from DANE, while the columns labeled 'Survey' present the results for the measure that uses survey data. Given that Model 2 is estimated without controlling for the degree of uncertainty, this column has no label indicating a measure of income uncertainty. We present the results separately by the likelihood to save and the amount of savings, given that an individual is actually saving (i.e. the intensive margin of saving or conditional saving) for each of the models in the upper and lower panel of Table 2.4, respectively.

The results from estimating Model 3 are presented in columns 6 and 7 labeled 'Model 3'. We find that the likelihood to save does not depend on uncertainty but is slightly positively correlated with loss aversion in one of the models. With regard to the amount of conditional savings, we find a negative coefficient for loss aversion for both sources of information; the coefficient is significantly different from zero when using the survey measure for unemployment risk. As the coefficient of loss aversion is centered around 1, the coefficient of income uncertainty shows the correlation between uncertainty and the likelihood to save or savings for a loss-neutral agent. We find that this coefficient is positive and statistically significant in one of the models. The coefficients of the interaction terms between loss aversion and income uncertainty are positive and significantly different from zero for the two measures of uncertainty that we use in the analysis.

This result supports the prediction that we derived from the model of Kőszegi and Rabin (2009) in this respect and is consistent with Hypothesis 3. The first result is thus:

Result 1. *The relationship between uncertainty and savings is an increasing function of the degree of loss aversion. Equivalently, the relationship between loss aversion and* *savings is an increasing function of the degree of income uncertainty. This result strongly supports Hypothesis 3.*

Estimation results of a direct test of Hypothesis 1 using Model 1 are presented below 'Model (1)' and 'Model (1b)' in Table 2.4. The first two columns present the results for the entire sample, while the columns labeled 'Model (1b)' present the results for the subpopulation of loss-averse individuals in the sample using a subset of control variables. We find that the results partly support Hypothesis 1. The likelihood to save is uncorrelated with uncertainty in three of the models. Yet, uncertainty is positively related with the amount of savings, given that an individual is actually saving. This result holds both for the complete sample and for the subpopulation of loss-averse individuals.

In addition to the results obtained from estimating Model 1 and Model 1b, we may use the insights from Model 3 to evaluate Hypothesis 1: The coefficients of uncertainty resulting from estimating the likelihood to save become less negative or even positive, but are both statistically insignificant, once we control for the degree of loss aversion and include an interaction term of loss aversion and income uncertainty (see column 'Model 3'). The coefficients of uncertainty in the equation modelling conditional saving are positive and statistically significant for the survey measure.

To assess the overall effect, we compute the marginal effects (at mean values) on the predicted amount of (unconditional) savings, resulting from estimating Model 3. We compute these marginal effects for gain-seeking to loss-neutral behavior ($\lambda = 1$), moderate loss aversion ($\lambda = 1.5$) and high loss aversion ($\lambda = 2$).

Figure 2.1 displays the corresponding marginal effects of income uncertainty on predicted total savings.[14] We find that an increase in income uncertainty is always associated with an increase in total savings—independent of the degree of loss aversion and the level of uncertainty. Furthermore, we find that the increase associated with an increase in income uncertainty is larger, the higher uncertainty and the higher the degree of loss aversion. Hence these results also confirm Hypothesis 3. This result holds for both measures of income uncertainty. This effect is mainly driven by an increase in conditional savings rather than by changes in the likelihood to save.

14. The marginal effects on the intensive and extensive margin with confidence intervals are printed in Appendix A.4, see Figure A.2.

TABLE 2.4: Results from Estimating Model 1, Model 2 and Model 3 Using a Negative Binomial Hurdle Model and Different Measures of Income Uncertainty

	Model (1)		Model (1b)		Model (2)	Model (3)	
	DANE	Survey	DANE	Survey		DANE	Survey
Likelihood of Saving							
Uncertainty	-0.052	-0.004	-0.395*	-0.001		-0.041	0.025
	(-0.72)	(-0.18)	(-1.93)	(-0.01)		(-0.55)	(0.66)
Loss Aversion					0.041*	0.035*	0.032
					(1.68)	(1.76)	(1.15)
Loss Aversion						0.008	0.012
× Uncertainty						(0.83)	(1.55)
Amount of Savings							
Uncertainty	0.196*	0.057**	0.883***	0.130***		0.096	0.150***
	(1.89)	(1.96)	(3.37)	(6.02)		(1.12)	(3.35)
Loss Aversion					0.058**	-0.038	-0.177***
					(2.55)	(-0.88)	(-3.25)
Loss Aversion						0.031**	0.022***
× Uncertainty						(2.12)	(4.47)
AIC	1232	1231	194	195	1244	1240	1225
Controls	25	25	6	6	25	25	25
Region	No	Yes	No	No	Yes	No	Yes
Occupation	Yes	Yes	No	No	Yes	Yes	Yes
Observations	640	640	97	97	640	640	640

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. t-values in parentheses.

Note: The dependent variable is the sum of all self-reported savings data in various saving devices, see Section 2.4.1. In this Negative Binomial Hurdle model, the participation equation estimates the likelihood to engage in saving money, while the second equation estimates conditional savings—the amount of savings, given that a person is saving money. The coefficient of loss aversion is centered at 1 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 Section 2.4.3. Income uncertainty is centered at the mean and is based on different measures, partly building on secondary data; see Section 2.4.2 for details. We control for variables listed in Tables 2.1 and 2.3. Furthermore, we control for regional and occupational effects at the *Localidad* level and for the working sectors according to the ISIC classification of economic activities, if indicated. We account for potential heteroskedasticity by robust standard errors.

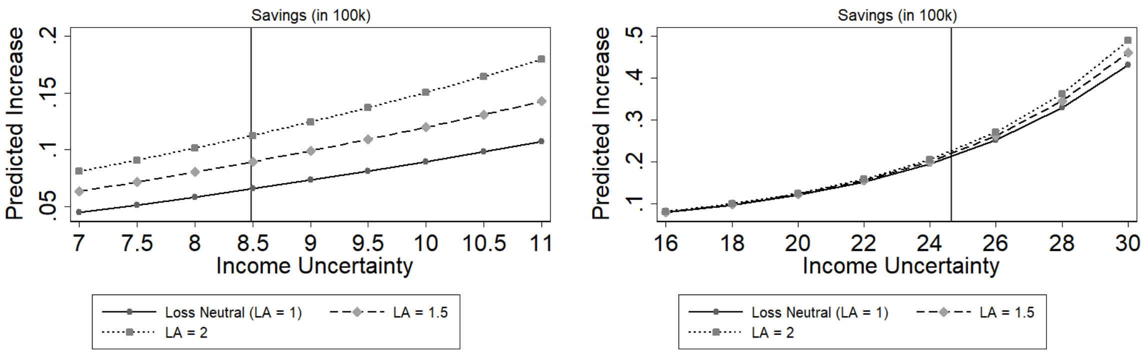

(a) Regional Unemployment Rate (DANE) (b) Regional Unemployment Risk (Survey)

FIGURE 2.1: Conditional Marginal Effect of Income Uncertainty on Predicted Savings for Different Degrees of Loss Aversion

Note: Marginal effects computed according to (2.3) using mean values for remaining covariates and estimates from fitting Model 3; see Table 2.4 for the corresponding coefficients and Figure A.2 for the marginal effects with confidence intervals for the intensive and extensive margins. Vertical lines indicate mean values of the corresponding measure of income uncertainty. Income Uncertainty is expressed in percent and the predicted increase in savings in 100,000 COP.

In terms of the magnitude of the effect, the result depends on the type of data we use. The estimations for the secondary data imply that when the average regional unemployment rate rises from 8.5 to 9.5 percent, for a moderately loss-averse sample, there is an increase in the average total savings of about 10,000 COP or roughly 5 USD. This corresponds to a relatively small increase in savings of 4 percent, which is due to the small proportion of individuals engaging in saving money. However, when the model is estimated using survey data on the average personal perception of unemployment risk, the predicted effect of a one percentage point increase at the mean level of income uncertainty has an effect that is nearly twice as large.

Summarizing, we find support for Hypothesis 1:

Result 2. *An increase in income uncertainty is associated with an increase in savings.*

Estimation results of a direct test of Hypothesis 2 using Model 2 are reported in column 5 ('Model 2') of Table 2.4. In this specification, no measure of income uncertainty is included, since it may be assumed that the entire sample is exposed to income uncertainty due to a lack of a social protection system (see Section 2.3). We deduce that an increase in loss aversion is associated with an increase in both the likelihood to save and the amount of conditional savings.

When controlling for the degree of uncertainty, including an interaction term between loss aversion and uncertainty (see column 'Model 3'), we see that this result is driven by those facing a high level of income risk. The coefficient of loss aversion in the equation for the likelihood to save is positive and so is the coefficient of the interaction term. However, in the equation for conditional savings, the coefficient of loss aversion is negative and significant for the survey measure. Yet, the interaction terms are both positive and significant. The coefficients of loss aversion correspond to an average level of income uncertainty, as uncertainty measures are centered around the mean. For an individual that is exposed to a high income uncertainty, the effect of loss aversion on savings is positive, as can be deduced from estimated marginal effects shown in Figure 2.2 (marginal effects on the intensive and extensive margin with confidence intervals are printed in Figure A.3 in the Appendix). We consider different levels of income uncertainty: The average level of income uncertainty, a high level of income uncertainty, defined as the third quartile of the distribution of income uncertainty, a very high and an extremely high income uncertainty corresponding to the 90% and 95% quantile of the income uncertainty distribution. The marginal plot relying on the secondary DANE data (Figure 2.2(a)) indicates a positive effect of loss aversion on predicted savings even for average levels of income uncertainty. Yet for the survey data (Figure 2.2(b)) the effect is negative even at a high level of income uncertainty. For very high levels of income uncertainty, the effect is positive. We summarize our findings with respect to Hypothesis 2:

Result 3. *An increase in loss aversion is associated with an increase in savings for individuals facing a high to a very high level of income uncertainty.*

2.5.4 Robustness Tests

In the literature, various definitions of loss aversion have been proposed. Yet, it is unclear which definition of loss aversion is best and should be considered standard (Abdellaoui, Bleichrodt, and Paraschiv, 2007). To address this issue, we have constructed different meta measures of loss aversion compromising two or more definitions. Our results are robust to using any of the different meta measures of loss aversion. This increases the validity of our results in comparison to using a single measure.

To test the robustness of our results in comparison to alternative measures of loss aversion, we reestimate Model 2 and Model 3 using different measures of loss aversion. Results of the estimated coefficients for Model 2 using different definitions of loss aversion are presented in Table 2.5. Column 2 shows the results when restricting the sample to those for which the meta measure compromised

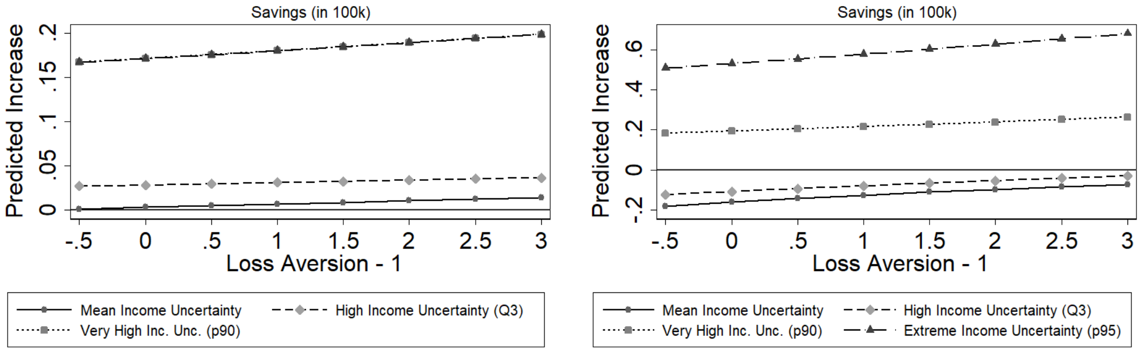

(a) Regional Unemployment Rate (DANE) (b) Regional Unemployment Risk (Survey)

FIGURE 2.2: Conditional Marginal Effect of Loss Aversion on Predicted Savings for Different Levels of Uncertainty

Note: Marginal effects computed according to Equation (2.3) using mean values for remaining covariates and estimates from fitting Model 3; see Tables 2.4 for the corresponding coefficients and Figure A.3 for marginal effects with confidence intervals for the intensive and extensive margins. The predicted increase in savings is expressed in 100,000 COP.

of 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. The results confirm our previous findings. Independent of the measure of loss aversion, we find that there is a significant increase in conditional savings associated with an increase in loss aversion. Loss aversion, however, does not affect the likelihood to save with one exception. The estimated coefficients for Model 3 using different meta measures of loss aversion are presented in Table 2.6. Although significance levels vary, the results are consistent with previous findings.

2.6 Discussion

Our empirical results support the predictions formulated by Kőszegi and Rabin (2009), as well as those derived in our paper, and we do find that savings are larger for loss-averse individuals who are exposed to income risk and that the larger the degree of loss aversion, the larger is the amount of savings for those exposed to income risk. This finding, however, does not unambiguously establish causality. Our intention here was to explore if asset accumulation provides supportive evidence of the predictions of this model. Future work should, for example, focus on establishing causal relationships through the use of panel data.

TABLE 2.5: Results from Estimating Model 2 Using a Negative Binomial Hurdle Model and Different Meta Measures of Loss Aversion

	Loss Aversion Meta Measure 2 Measures			Loss Aversion Meta Measure 3 Measures		Loss Aversion Meta Measure 5 Measures
	(1)	(2)	(3)	(4)	(5)	(6)
Likelihood of Saving						
Loss Aversion	0.0414	0.0415	0.0276	0.0777*	0.0606	0.267
	(1.68)	(1.65)	(1.00)	(2.00)	(1.42)	(1.93)
Amount of Savings						
Loss Aversion	0.0583*	0.0655**	0.0883***	0.0885***	0.103***	0.247*
	(2.55)	(2.84)	(3.62)	(3.44)	(3.43)	(2.09)
AIC	1244	1140	1019	1138	1019	1021
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

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. t-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various saving devices, see Section 2.4.1. 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. Loss aversion is measured by continuous and experimentally elicited meta measures. The meta measure compromising 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, Minehart, and Rabin (1999) and Wakker and Tversky (1993) are also included. For more details on the applied measures of loss aversion, see Section 2.4.3. We control for variables listed in Tables 2.1 and 2.3. Furthermore, we control for regional and occupational effects on the *Localidad* level and for the working sectors according to the ISIC classification of economic activities. We account for potential heteroskedasticity by robust standard errors.

TABLE 2.6: Results from Estimating Model 3 Using a Negative Binomial Hurdle Model and Different Meta Measures of Loss Aversion

	Regional Unemployment Rate			Regional Unemployment Risk		
	Measure 1	Measure 2	Measure 3	Measure 1	Measure 2	Measure 3
Likelihood of Saving						
Loss Aversion	0.00774	0.0120	0.0609	0.0121	0.0174*	0.0459**
× Uncertainty	(0.83)	(0.82)	(0.86)	(1.55)	(1.79)	(2.48)
Loss Aversion	0.0353*	0.0698**	0.273**	0.0319	0.0635	0.131
	(1.76)	(2.12)	(2.00)	(1.15)	(1.43)	(0.80)
Uncertainty	-0.0410	-0.0316	0.0336	0.0253	0.0547	0.0959**
	(-0.55)	(-0.40)	(0.36)	(0.66)	(1.29)	(2.23)
Amount of Savings						
Loss Aversion	0.0315**	0.0315	0.0611	0.0224***	0.0199**	0.0465
× Uncertainty	(2.12)	(1.22)	(0.78)	(4.47)	(2.21)	(1.54)
Loss Aversion	-0.0382	-0.0204	0.0893	-0.177***	-0.129	-0.169
	(-0.88)	(-0.26)	(0.58)	(-3.25)	(-1.22)	(-0.64)
Uncertainty	0.0955	0.0573	0.108	0.150***	0.167***	0.186***
	(1.12)	(0.68)	(0.75)	(3.35)	(3.69)	(2.95)
AIC	1240	1137	1023	1225	1125	1015
Controls	25	25	25	25	25	25
Region	No	No	No	Yes	Yes	Yes
Occupation	Yes	Yes	Yes	Yes	Yes	Yes
Observations	640	579	509	640	579	509

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. t-values in parentheses.

Note: The dependent variable is the sum of self-reported savings data in various saving devices, see Section 2.4.1. 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. Loss aversion is measured by continuous and experimentally elicited meta measures. The meta measure compromising 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, Minehart, and Rabin (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 Section 2.4.3. Income uncertainty is centered at the mean and is based on different measures, partly building on secondary data; see Section 2.4.2 for details. We control for variables listed in Tables 2.1 and 2.3. Furthermore, we control for regional and occupational effects at the *Localidad* level and for the working sectors according to the ISIC classification of economic activities, if indicated. We account for potential heteroskedasticity by robust standard errors.

The measure for loss aversion that we use in this study is, on average, lower than that estimated by Abdellaoui, Bleichrodt, and Paraschiv (2007) using a similar approach. However, other studies also find considerably lower shares of loss aversion at the individual level. For example, Schmidt and Traub (2002) and Bleichrodt and Pinto (2002) report that between 33 percent and 5 to 30 percent of their participants could be classified as loss averse, respectively. In addition, Schmidt and Traub (2002) report that 24 percent of their participants exhibit gain-seeking behavior at the individual level. Against this background, our findings do not seem unrealistic and even less so when we consider that our experiment was conducted with a non-student sample in a development setting, where the mean household per capita income ranges near the poverty level and the prospect of losing money might be very familiar and not frightening enough to make an effort to avoid these losses at the price of forgoing gains.

We observe considerable heterogeneity in our measures of loss aversion, as expressed in standard deviations (see Table 2.3), which is to be expected in field settings. Should we assume that—for some reason—our sample is more loss averse in reality than our measures indicate, we should expect a higher coefficient of loss aversion for every individual in the sample. This would not affect our results with respect to Hypothesis 3, as they are independent of the absolute value of the coefficient of loss aversion. Furthermore, as our results regarding Hypothesis 1 based on marginal effects also hold for what we label loss-neutral individuals and slightly gain-seeking individuals, we also consider the results regarding Hypothesis 1 to be unaffected by a possible miscalibration of our measures of loss aversion. Lastly, for the direct test of Hypothesis 2, we can assert that the results are unaffected by such a transformation. Given similar findings in other studies, however, we are quite confident in our measures of loss aversion.

In this study, we use the official regional unemployment rate and the perceived regional unemployment risk as a measure of income risk. These measures can be interpreted as exogenous as one single individual cannot affect the labor market. Unfortunately, they are average measures and one of the measures relies on self-reported data. Therefore, they are likely to contain errors—be it because of their coarse nature or their incorrect assessment of unemployment risk. Therefore, the precautionary saving motive might be imperfectly captured in our data. Another limitation of this measure is that the labor market can generally be considered as integrated since individuals in urban areas can commute to other *Localidades* to find employment. On the other hand, we find substantial variability in unemployment levels for the different units of analysis, which suggests that the labor market is fragmented. While unemployment risk is an important source of income uncertainty for our population group, future work should also explore other sources of income risk affecting households such as health risk and price changes.

Although the theory in Section 2.2 focuses on future income uncertainty, contemporaneous unemployment figures are a good prediction of future unemployment rates and thus future income uncertainty. In addition, a higher unemployment rate in the past is likely to induce a certain feeling of uncertainty about future income. Lastly, labor income risk affects overall income risk, because retirement income depends on savings built up during an individual's working years or the income of family members in the absence of a formal social security system. The latter, however, can be assumed to be limited if the household has been in monetary need during their education years, reducing the priority of education and increasing the priority of earning money.

All of our results are based on one particular sample: People from low-income households in Bogota and approximately two-thirds are female. It would therefore be of great interest to validate this study's findings with other data from other countries which have other socio-demographic characteristics.

Despite the limitations, this study is to be considered as a first step in determining the role of loss aversion in financial decisions of the poor.

2.7 Conclusion

In this paper, we have tested if the theoretical predictions of the intertemporal model of consumption and saving by Kőszegi and Rabin (2009) can be empirically supported. More specifically, we have tested if loss averse individuals that face income uncertainty hold higher savings, which would be consistent with a loss-aversion based precautionary savings motive (Hypothesis 1). Our results support this hypothesis and the increase in savings for a moderately loss averse sample associated with a one percentage point increase in income uncertainty is modest and amounts to between 4 and 8 percent of total savings. Secondly, we have tested whether individuals that exhibit a higher degree of loss aversion hold more savings than individuals with a lower degree of loss aversion, given that they face income uncertainty (Hypothesis 2). Our results support this hypothesis, on average, and a detailed analysis shows that this is driven by individuals facing a high level of income uncertainty. Lastly, we have investigated whether the increase in savings associated with an increase in income uncertainty is larger, the higher the degree of loss aversion (Hypothesis 3). For this last hypothesis, we find very strong support.

These findings can be used to inform policy makers. Savings and pension campaigns could stress the uncertainty of future income to boost savings. Moreover, if future income is uncertain, a higher degree of loss aversion can be expected to induce an additional savings motive compared to a lower degree of loss aversion, so people should be reminded that it is unlikely they will be able to maintain their current standard of living when their income drops.

Chapter 3

Higher Order Risk: An Application To Savings of the Poor in Bogota

---

This chapter is joint work with Marcela Ibañez and Gerhard Riener.

3.1 Introduction

Understanding motivations for savings' decisions is fundamental for policy makers and academic economists alike. In the recent decades several behavioral channels that determined subjects' savings decisions have been explored, especially emphasizing potential short sightedness of agents resulting in a present bias. In this work, we reconsider a less explored channel, that only recently received more attention (Noussair, Trautmann, and Kuilen, 2014): the relation between uncertainty and higher order risk attitudes. In his seminal work, Leland (1968) shows that uncertainty of future income can generate a precautionary savings motive. As a consequence, prudent individuals, i.e. agents who prefer taking risks in a state of higher wealth,[1] have an incentive to accumulate wealth in order to be better prepared to confront uncertainty. Hence, this model predicts that savings do not depend on risk aversion but on the intensity of prudence.

We extend the model by Leland (1968) to show that precautionary savings are not incompatible with risk-seeking behavior when individuals are prudent. We show that the demand for savings is proportional to the intensity measure of prudence advocated for by Crainich and Eeckhoudt (2008). In previous empirical literature the standard measure for prudence relied on the definition introduced by Kimball (1990). As this measure depends on the second derivative of the utility function, it can never be interpreted as a measure for the strength of the precautionary saving motive independent of the level of risk aversion. Our simple extension allows to rationalize the empirical finding that risk-loving, prudent individuals save a fraction of their income while previous theoretical models (Crainich, Eeckhoudt, and Trannoy, 2013) predict that risk-loving individuals would save all of their income when facing uncertainty, given that they are prudent.

To test the prediction of the model we develop a method to obtain a parameter-free estimate of the intensity of prudence from binary lottery choices. This measure does not impose a relationship between risk aversion and prudence, as commonly used parametric forms do. The experimental procedure builds on the design by Abdellaoui, Bleichrodt, and Paraschiv (2007). Individuals made pair-wise comparisons of simple lotteries that vary by the size of the rewards. By finding the values of income that give mid-point utility levels, that method allows to obtain a mapping of income to utility levels. We advance these non-parametric methods of utility point elicitation by providing a suitable and flexible way of

1. In expected utility (EU) theory, this is equivalent to a positive third derivative of the utility function (Eeckhoudt and Schlesinger, 2006).

linking the income-utility points to a continuous and differentiable function. In particular, we introduce a way of incorporating value constraints in estimating a P-spline utility function. We propose a way of jointly penalizing different orders (i.e. jointly smoothing different derivatives) using a data-dependent weight. One advantage of this method is that it allows to compute intensity measures and classify individuals over a large variation of stakes. The method we propose is to our knowledge the first one to allow the non-parametric computation of theory-based intensity measures; in particular the intensity measures by Kimball (1990) and Crainich and Eeckhoudt (2008).

We empirically test the relation between risk preferences, prudence and savings among a representative sample of 693 individuals from poor households in Bogotá, Colombia. This population group is subject to a large degree of income uncertainty, hence precautionary motives are an important motivation for saving. We combine answers from an extensive financial household survey with lab-in-the-field experimental data conducted in October and November 2013.

To the best of our knowledge, we are the first to study the correlation of experimental measures of prudence and real saving in a development setting. Most closely related work to this part of our study is Noussair, Trautmann, and Kuilen (2014), who found a positive relation of prudence and savings. In the spirit of Cardenas and Carpenter (2013) and Noussair, Trautmann, and Kuilen (2014), we collect financial information and other socioeconomic characteristics of households using a survey and relate this behavior with our incentivized measures of risk aversion and prudence.

We find that between 37% to 58% of our sample can be classified as prudent and 21% to 48% as risk averse, depending on the measure used. For prudence, these numbers lie within the range of previous studies (Tarazona-Gomez, 2004; Deck and Schlesinger, 2010; Maier and Rüger, 2011). However, we find considerably more risk-loving behavior. The fraction of prudent individuals is similar among risk-averse and risk-loving individuals (60% and 55%, respectively). While planning horizon is positively correlated with prudence, other demographics seem to play a minor role in predicting prudence.

The empirical analysis strongly supports Leland's (1968) model. We find a strong positive correlation between saving and prudence. This relation becomes even stronger when we include a measure of income risk—based on secondary data measuring the ratio of shut down to existing businesses in 2013 in Bogotá in the sector an individual was usually employed. We also find some support for a precautionary saving demand for the subset of risk-loving individuals. However, probably due to a lack of power, this finding is less robust. Our study adds to the extended research on household savings by theoretically and empirically establishing the link between risk aversion, savings and prudence. Previous studies used survey data to determine the share of savings that is due to income uncertainty and in that way assess the importance of prudence (Guiso, Jappelli, and Terlizzese, 1992; Dynan, 1993; Lee and Sawada, 2010; Fagereng, Guiso, and Pistaferri, 2017). However, those papers face the problem that the data does not allow for separately identifying individual risk aversion and prudence (see Appendix B.1 for an illustration of this point and its consequences). To deal with that limitation another branch of studies uses experimental measures of prudence and risk aversion (Tarazona-Gomez, 2004; Deck and Schlesinger, 2010; Maier and Rüger, 2011; Ebert and Wiesen, 2011). While those studies make a methodological contribution in the measurement of prudence, they focus on a student population and do not examine saving decisions. In addition, these studies have the methodological problem that risk-apportionment tasks cannot inform about theory-based intensity measures of prudence but simply classify individuals as prudent or imprudent. We contribute to this research by extending the analysis to consider measures of intensity of prudence without relying on parametric utility functions. Unlike previous studies that use compound lotteries that involve gains and losses, we present our participants simple lottery choices that only involve positive outcomes. In this way we simplify the decision problem and avoid confounding prudence with loss aversion.

The study closest to ours is Noussair, Trautmann, and Kuilen (2014) who, using a representative sample of the Dutch population, find support for the relation between prudence and the precautionary motive for saving. We contribute to this research by examining savings for both risk-averse and risk-loving individuals. While the theoretical model by Leland (1968) predicts that this relation is independent of the level of risk aversion, the empirical evidence has not yet examined the relation between risk-loving behavior, prudence and savings.

Alternative theoretical models examining the relation between uncertainty and savings were proposed by Bowman, Minehart, and Rabin (1999) and Kőszegi and Rabin (2009). These models consider reference-dependent utility functions and show that due to loss aversion there is an asymmetric response to good and bad news. In addition, Kőszegi and Rabin (2009) show that in case of uncertain future income there is a precautionary saving motive among loss-averse individuals who have a higher utility loss from lowering expected consumption than from comparable increases of it. In Chapter 2, Ibanez and Schneider empirically validate this conjecture and find strong support for it.

3.2 Theoretical Framework

We build upon the model by Leland (1968). In the first part of this section we modify the assumption on the relationship between the first and second period consumption utility giving up strict additive separability and show the consequences on the savings decisions. In the second part we discuss how two well-known measures of intensity of prudence relate with the theoretical model. We show that the measure of prudence popularized by Crainich and Eeckhoudt (2008) directly indicates the strength of such a saving motive while being defined independently of the second derivative of the utility.

The assumption of first and second period consumption utility being perfect substitutes that is implied by additive, time-separable overall utility has been empirically debated at least since Loewenstein (1987). Different models of decision over time, starting with Gilboa (1989), have been axiomatizing a utility path dependence (see also Wakai, 2008, Axiom 3 for a more recent example). We introduce a 'linking function' of first and second period consumption. In the spirit of Gilboa (1989), we assume that individuals dislike variation in consumption and experience increasing disutility when consumption differs over time. This extension allows for risk-loving individuals to save a non-trivial fraction of their income as opposed to the predictions derived by Crainich, Eeckhoudt, and Trannoy (2013).

3.2.1 An Extension of Leland's Precautionary Saving Model

Following Leland (1968), we assume that an agent lives for two periods and receives income w_t for $t \in 1, 2$. The income in the first period is deterministic, whereas in the second period, w_2, is random with known expectation and variance given by $\mathbb{E}[w_2] = \bar{w}_2$ and $\mathbb{E}[(w_2 - \bar{w}_2)^2] = \sigma^2$, respectively.

The agent has access to financial markets, where the fraction of income saved in $t = 1$, $k \leq 1$, receives an interest rate $r > 0$.[2] Consumption in $t = 1$ is $c_1 = (1 - k)w_1$ and consumption in $t = 2$ is $c_2 = w_2 + (1 + r)kw_1$.

The agent's objective is to maximize the expected inter-temporal utility of consumption $\mathbb{E}[U(c_1, c_2)]$ by deciding on the fraction of income, k, that they would save in $t = 1$. The inter-temporal utility of consumption is given by

$$U(c_1, c_2) = u_1(c_1) + u_2(c_2) + g(c_2 - c_1), \tag{3.1}$$

2. In case of negative savings, k is restricted such that c_2 will always be non-negative.

where $u_t(c_t)$ is the three times differentiable utility function of consumption for $t \in 1, 2$ and g denotes a three times differentiable linking function of first and second period consumption. We assume that the linking function is concave. By incorporating this linking function, we modify the assumption of additive time-separable utility previously adopted by Kimball (1990). Such an additive term relating consumption of the two periods has been applied before in the context of intertemporal consumption and saving models by Bowman, Minehart, and Rabin (1999) or Kőszegi and Rabin (2009). Unlike those models we do not assume reference-dependent utility preferences. We assume positive marginal utility of consumption ($u_t' > 0$) but do not make any assumption about the second derivative of the utility function and allow u_t'' to be either positive indicating risk aversion or negative indicating risk-loving behavior.

The first- and second-order condition for an interior solution imply that in absence of uncertainty, the following conditions are satisfied for the optimal consumption bundle (c_1^*, c_2^*) resulting from the optimal saving rate k^*:

$$\frac{dU}{dk} = (1+r)\frac{\partial U}{\partial c_2} - \frac{\partial U}{\partial c_1} = 0 \tag{FOC}$$

$$\frac{d^2U}{dk^2} = \frac{d}{dk}\left[(1+r)\frac{\partial U}{\partial c_2} - \frac{\partial U}{\partial c_1}\right] < 0. \tag{SOC}$$

For an additive U with a linking function as in (3.1), (SOC) writes

$$u_1''(c_1) + u_2''(c_2)(1+r)^2 + g''(c_2 - c_1)(2+r)^2 < 0.$$

Clearly, in case u_1 and u_2 are strictly concave in c_t (corresponding to a *risk-averse* individual) and when $g'' \leq 0$, the second order condition is satisfied. Moreover, if g'' is negative and the absolute value of the last term is bigger than the first two in (SOC), then overall utility is concave in k. Hence, the second order condition (SOC) is satisfied, even if u_1 and u_2 are convex as in the case of a risk-loving individual.

To study how income uncertainty affects the optimal level of the saving decision, we follow Leland. Taylor series expansions of $\mathbb{E}\left[\partial U/\partial c_t\right]$ in the resulting first-order condition around the optimal consumption bundle in absence of uncertainty indicate that the effect of uncertainty on the optimal saving rate depends on the sign of

$$\left(\frac{\partial^3 U^*}{\partial c_1 \partial (c_2)^2} - (1+r)\frac{\partial^3 U^*}{\partial (c_2)^3}\right)\sigma^2, \tag{3.2}$$

where U^* is defined as U evaluated at the consumption bundle resulting from the optimal saving rate in absence of income uncertainty.[3] Leland (1968) infers that if (3.2) is negative, we may "[u]nder reasonable regularity assumptions [...] say that the optimal [saving rate] will be larger [...] when uncertainty is present, and the more uncertainty, the greater will be the optimal [saving rate]." This implies that the saving rate will be larger when income is uncertain if the following expression is negative:

$$-\left((2+r)\,g^{*\prime\prime\prime}(c_2-c_1)+(1+r)u_2^{*\prime\prime\prime}(c_2)\right)\sigma^2. \tag{3.3}$$

This corresponds to (3.2) for utility as given in (3.1). Thus, uncertainty results in a positive precautionary demand for saving for prudent individuals (those for which $u_2''' > 0$) when $g''' \geq 0$, provided that the necessary condition of a negative second derivative of overall utility with respect to the saving rate (SOC) is satisfied. We summarize this in the following proposition:

Proposition 2. *Let $U(c_1, c_2) = u_1(c_1) + u_2(c_2) + g(c_2 - c_1)$ be the overall intertemporal utility of an agent. Assume g such that $d^2U/dk^2 < 0$. Then a non-negative third derivative of the linking function g and a positive third derivative of the second period utility, i.e. $g''' \geq 0$ and $u_2''' > 0$, are a sufficient (though not necessary) condition for a positive precautionary demand for saving under the assumption of the model of precautionary saving by Leland (1968). This statement holds independently of the sign of the second derivative of the second period utility.*

3.2.2 Examplary Linking Function

So far the 'linking function' has been characterized as a concave function of the difference of consumption in the first and in the second period and hence as a concave function of the proportion k of income saved. We now provide an example of a particular linking function. We assume that

$$g(x) := -lx^2 \tag{3.4}$$

with $l \geq 0$; for l large enough, $d^2U/dk^2 < 0$ and the extreme of the utility function at the critical point satisfying the first-order condition is a utility maximum (for a risk-averse individual, $l = 0$ is large enough).

Introducing a function $g(c_2 - c_1) = -l(c_2 - c_1)^2$ for $l > 0$ decreases overall utility with an increasing difference between consumption in the first period

3. More generally, for a function $f(k)$, we define f^* as f evaluated at the optimal saving rate k^* in absence of uncertainty.

and consumption in the second period. This special choice of the linking function incorporates variation aversion or a preference for spreading good outcomes evenly over time (for $l > 0$) as observed empirically by Loewenstein and Prelec (1993), which challenged time-separable utility. Gilboa (1989) axiomatically derived a utility function with path-dependence, where overall utility is decreased as the difference between consumption in the first and consumption in the second period grows, holding any of the two constant. More specifically, Gilboa (1989) proposed a linking function of the form $|u(c_2) - u(c_1)|$. In a more recent work, Wakai (2008), translates the idea of variation aversion in a setting where negative variation is more unpleasant than positive variation is pleasant.

The way we incorporate variation aversion is an analytically convenient variation of the utility function proposed by Gilboa (1989), which also decreases overall utility as the difference between consumption in the first and consumption in the second period grows, but our function does so independently of the absolute levels of consumption. This simplification might seem strict, as for high consumption levels, a relatively small difference might loom less than the same absolute difference for low consumption levels or vice versa. As the scaling parameter l might capture these individual wealth levels, however, we argue it is appropriate.

With respect to life-time saving, our choice of the linking function is a way to incorporate the aim of consumption smoothing as predicted by the permanent income hypothesis (Friedman, 1957) or simply the commitment of living and consuming in the next period.

This choice of utility could, additionally to any risk-averse individual, also represent an individual that is willing to take risk in each period, but only as long as the difference between consumption in both periods does not get too large—impeding 'ruthless' over-consumption in any one of both periods.

If $g(c_2 - c_1) = -l(c_2 - c_1)^2$ as in the example above, then $g''' \equiv 0$ and we rewrite (3.3):

$$-(1 + r)u_2^{*'''}(c_2)\sigma^2. \tag{3.5}$$

From this we see that just as in the case of simple additive, time-separable utility, a positive precautionary saving demand results solely from a positive third derivative of second period utility with respect to k.

We summarize these findings in a corollary:

Corollary 2. *Let $U(c_1, c_2) = u_1(c_1) + u_2(c_2) - l(c_2 - c_1)^2$ be the overall intertemporal utility of an agent. Assume l large enough such that $d^2U/dk^2 < 0$. Then a positive third derivative of the second period utility, i.e. $u_2''' > 0$, is both a sufficient and a necessary condition for a positive precautionary demand for saving under the assumption of the model of precautionary saving by Leland (1968). This statement holds independently of the sign of the second derivative of the second period utility.*

Corollary 2 states that—given the linking function, and thus the overall utility, is of the most simple form allowing for a regular utility maximum—a positive third derivative of the second period utility alone causes a positive precautionary demand for saving, for risk-loving and risk-averse individuals, where for the latter, $l = 0$ is large enough.

Unfortunately, the existence of a linking function and its particular shape is not directly testable. This drawback is also inherent in the models by Kőszegi and Rabin (2009) and Bowman, Minehart, and Rabin (1999).

3.2.3 Measuring the Strength of the Precautionary Saving Motive

We now show, building on the model by Leland (1968), that the measure by Crainich and Eeckhoudt (2008) can also be directly interpreted as a measure of intensity of a precautionary savings demand whereas the measure by Kimball (1990) is restricted to the case of risk-averse individuals. This second measure is adequate when comparing risk-averse individuals only, but cannot generally be used in the framework of Leland's model or the extension of the model that we present here.

Crainich and Eeckhoudt measure Building on previous work on downside risk aversion,[4] Crainich and Eeckhoudt (2008) suggest to measure the degree of prudence by $\pi = u'''/u'$ (Modica and Scarsini, 2017; Keenan and Snow, 2002). The intuitive interpretation they give of this measure is the analog to the utility premium for compensating the pain of a zero-mean risk. When there is 'misapportionment of risk' (meaning risk added to a state of lower wealth instead of to the state of higher wealth), $\pi = u'''/u'$ is proportional to the money equivalent of pain induced by this misapportionment. One advantage of this measure is that it is independent of the sign of the second derivative. Hence, it is closely related to

4. Downside risk aversion is equivalent to prudence for three times differentiable utility functions.

the theoretical prediction of the model by Leland (1968) as it can be computed for both risk-averse and risk-loving individuals leading to a similar interpretation.

To see how the measure by Crainich and Eeckhoudt (2008) can be incorporated in the Leland (1968) framework, rewrite (3.2) as

$$\left(\frac{\partial^3 U^*}{\partial c_1 \partial (c_2)^2} - \frac{\partial U^*}{\partial c_1} \Big/ \frac{\partial U^*}{\partial c_2} \frac{\partial^3 U^*}{\partial (c_2)^3}\right) \sigma^2 = \left(\frac{\partial^3 U^*}{\partial c_1 \partial (c_2)^2} - \frac{\partial U^*}{\partial c_1} \frac{\partial^3 U^*}{\partial (c_2)^3} \Big/ \frac{\partial U^*}{\partial c_2}\right) \sigma^2, \quad (3.6)$$

where we used the first-order condition (FOC) and rearranged terms.

If the utility is additively time-separable (i.e. $g \equiv l \equiv 0$), (3.6) can be written as

$$-u_1^{*\prime} \pi^* \sigma^2. \quad (3.7)$$

Hence, it is clear that the savings rate increases with the intensity of prudence $\pi^* = \frac{\partial^3 U^*}{\partial (c_2)^3} \Big/ \frac{\partial U^*}{\partial c_2} = \frac{u_2^{*\prime\prime\prime}}{u_2^{*\prime}}$.

Let us now turn to the case of a utility with an additive linking function of the form $-l(c_2 - c_1)^2$ relating consumption in the two periods. (3.6) in this case equals (3.3) with $g''' \equiv 0$. We focus on the second term,

$$-(1+r)u_2^{*\prime\prime\prime}(c_2) \quad (3.8)$$

and find that the larger $u_2^{*\prime\prime\prime}$, the larger the precautionary savings demand under the conditions derived before. Also in this case, π is a good measure of the precautionary savings demand. First, because dividing $u_2^{*\prime\prime\prime}$ by $u_2^{*\prime}$ leaves the sign unchanged. Second, following the rationale by Pratt (1964) when justifying his measure, multiplying u with a positive constant does not change behavior, but it changes u'''. The measure π is unaffected by such a transformation.

We summarize these findings:

Proposition 3. *Let $U(c_1, c_2) = u_1(c_1) + u_2(c_2) - l(c_2 - c_1)^2$ be the overall intertemporal utility of an agent. Assume l large enough such that $d^2U/dk^2 < 0$. Then, all else equal, $m^* = \frac{\partial^3 U^*}{\partial (c_2)^3} \Big/ \frac{\partial U^*}{\partial c_2}$ indicates the strength of a precautionary demand for saving under the assumption of the model of precautionary saving by Leland (1968), including regularity assumptions. This statement holds independently of the sign of the second derivative of the second period utility.*

Note that for risk-averse individuals, Proposition 3 holds for $l = 0$, i.e. under the usual assumption of time-separable utility.

Kimball (1990) measure The first measure of the degree of intensity of prudence was proposed by Kimball (1990). Using a close analogy to the Arrow-Pratt measure of risk aversion, the intensity of prudence is defined as $-u'''/u''$. In a simple two-period model with additive time-separable utility, Kimball (1990) shows that the savings function of a globally more prudent individual at a given level of saving moves upward at a lower level of risk. This measure thus is directly related to the intensity of the precautionary saving motive.

The measure by Kimball (1990) has two shortcomings that are especially relevant when trying to apply the concept empirically, see Appendix B.1 for an example. First, since the measure depends on the second derivative of the (per-period) utility function, it implies that precautionary savings depend on risk aversion. However, Leland (1968) shows that the precautionary demand for savings is independent of the degree of risk aversion.

Second, when focusing on precautionary savings, Kimball (1990) neglects the possibility of a convex utility function. Hence, the proposed measure is only meaningful for risk-averse individuals, as for them, it shows a positive value when prudent, but not for risk-loving individuals, for whom its value is negative if prudent. Similarly, this measure yields a positive intensity of prudence for an imprudent individual (negative third derivative) that is risk loving. This is clearly a contradiction.

3.3 Methodology

In this section, we introduce the procedure for the joint and consistent elicitation of (higher order) risk preferences. This method is particularly suitable for the computation of intensity measures of (higher order) risk preferences. To this end, we combine well-established methods to elicit utility points with the statistical approach of P-spline regression. The experimental elicitation and estimation procedure consists of these three steps:

1. Elicitation of utility points using any suitable method, such as the certainty equivalent method or the trade-off method used in Wakker and Deneffe, 1996.

2. Estimation of a differentiable individual utility function and their derivatives based on elicited utility points using a penalized spline (P-spline) approach (P. H. Eilers and Marx, 1996).

3. Derivation of higher order risk preferences with intensity measures to individual utility functions

The preferred procedure to estimate continuous and differentiable utility functions is to specify a parametric utility function and then estimate these parameters by maximum likelihood or generalized method of moments. A limitation of this approach is that parametric functions often impose restrictions on the higher order derivatives linking them to lower order derivatives. In this respect, already Fuchs-Schündeln and Schündeln (2005) note that individuals are nearly always assumed to be prudent and specifications allowing for imprudence cannot be combined with risk aversion.[5] An alternative to this is approach is to apply a fully non-parametric procedure that does not restrict the relation between lower and higher order derivatives. P-splines offer such an alternative. This method is suited in situations where (potentially with error) elicited utility points should be interpolated and smoothed to yield a continuous and possibly p-times continuously differentiable utility function including smoothed derivatives. This allows us to derive intensity measures of prudence and risk aversion. For its non-parametric character, this procedure can be seen as the natural completion of non-parametric elicitation methods for utility points to a non-parametric elicitation method for utility functions.

3.3.1 Elicitation of Utility Points

In the expected utility (EU) framework, one established and accepted method to non-parametrically elicit utility points is the trade-off method (Wakker and Deneffe, 1996).[6] The method elicits certain payoff (x_i) that imply indifference between two outcome gambles. Denote a binary lottery with $(x, p; y)$, where x is the upside, occuring with probability $p > 0$ and y is the downsidewith corresponding probability $(1 - p)$. The participant first states the value x_1 that makes her indifferent between $(x_1, p; r)$ and $(x_0, p; R)$ where $x_1 > x_0$ and $R > r$. Then the participant is asked for the value x_2 that makes her indifferent between $(x_2, p; r)$ and $(x_1, p; R)$. Assuming EU, and denoting the utility[7] of a monetary outcome x with $U(x)$, these two indifferences imply that

$$pU(x_1) + (1 - p)U(r) = pU(x_0) + (1 - p)U(R) \quad \text{and}$$
$$pU(x_2) + (1 - p)U(r) = pU(x_1) + (1 - p)U(R).$$

5. We derive these limitations in Appendix B.2 for the exponential (CARA), the power (CRRA) and the expo-power utility family.
6. Note that it is also possible to use the less complex certainty equivalent method for the elicitation of utiltiy functions as introduced in this paper.
7. Note that for this method, no further specification of the utility function are needed.

From these equations, we derive

$$p(U(x_1) - U(x_0)) = (1 - p)(U(R) - U(r)) = p(U(x_2) - U(x_1)).$$

Since $p > 0$, we can conclude that

$$U(x_1) - U(x_0) = U(x_2) - U(x_1).$$

This gives an equality of utility differences that can be used to elicit different utility points by repeating this iterative process leaving the researcher with the desired number of utility points between r and R. [8]

3.3.2 P-Spline Interpolation and Error Correction for Utility Functions and Their Derivatives

How does penalized spline (P-spline) regression connect these utility points? A first approach to non-parametric estimations is linear interpolation(see e.g. Fennema and Van Assen, 1998; Abdellaoui, 2000; Etchart-Vincent, 2004; Abdellaoui, Bleichrodt, and Paraschiv, 2007). This approach is suited if distances between subsequent utility points are small, decision or measurement errors are unlikely, precision of interpolation is of lower priority, and if enough points are elicited. For computing a fourth derivative, one needs at least five utility points.

Linear interpolation is also relatively 'costly' in terms of non-classifiable utility functions , as inconsistencies in one decision carry over to the estimation of other utility points and in some cases have to be excluded from the analysis (see for example Bleichrodt and Pinto, 2000). Moreover, and important for applications, a linear interpolation generally does not establish a differentiable function.

Penalized spline estimations smooth the data in a 'global' way, thus incorporating all information available. This results in estimates for the utility and its derivatives in one single fit. Therefore, there is no need to additionally smooth the derivatives or compute them numerically. Similarly to the parametric approach, this method is very parsimonious. One point between the fixed limit points is enough to determine the sign of at least the third derivative.

Spline Regressions Spline regressions generalize the concept of linear regressions. Instead of using only the x-value of elicited utility points in a regression, they add higher powers of those values—up to degree p. This approach,

8. This procedure can be extended to elicit utility on the loss domain, to elicit probability weighting (both in the gain and the loss domain) and to elicit loss aversion, as proposed by Abdellaoui, Bleichrodt, and Paraschiv (2007).

polynomial regression and interpolation, leads to a polynomial of order p, and roots in the strong theoretical foundation of the Stone-Weierstrass theorem. The Stone-Weierstrass theorem states that every real, continuous function defined on a closed interval can be uniformly approximated arbitrarily close by a polynomial function. In practice, however, the order is limited by available data points and the data will be underfit. Further, even when the order of polynomial functions is high, interpolation quality can be very poor for some functions (e.g. *Runge's phenomenon* Runge, 1901) and at the boundaries of the interval under study, the interpolation function becomes unstable. In general, a global basis approach often lacks the flexibility to adequately adjust the degree of curvature to different, possibly asymptotically constant regions. The resulting function would thus in many cases either underfit or overfit the data.

Splines are an established solution in this case. The global basis (e.g. the polynomial basis $1, x, x^2, x^3, \ldots, x^p$) is exchanged for (or extended by) a local basis and the result is a smooth combination of piecewise polynomial functions of degree p, that is, with common implementation, $(p-1)$ times continuously differentiable.

To that end, the domain of definition $[x_{\min}, x_{\max}]$ is divided into $k-1$ subintervals, where the k boundaries are called inner knots. Local basis functions, defined depending on the knots, are placed (often equidistantly, thus independent of the data) such that they cover the domain of definition.

One such choice of basis functions, so called B-splines, have proven to be numerically stable and efficient for computation.[9] A single B-spline of degree[10] p is a combination of $p+1$ polynomial pieces of degree p that are joined smoothly (i.e. $p-1$ times continuously differentiable) at the knots. It is different from zero only on a small subinterval of the domain of definition (spanned by $p+2$ knots) and zero otherwise.

Figure 3.1 illustrates single B-splines of degree $p = 1$ and degree $p = 2$. For illustration purposes, the first $p+2$ inner knots are indicated by the gray vertical lines at $x = 0.1$, 0.2 and for the B-spline of degree 2 additionally at $x = 0.3$.

We use a B-spline basis consisting of $k+p-1$ equally spaced B-splines spanned by $k+2p$ knots; see 3.2(a) for an illustration of an exemplary B-splines basis. Also

9. De Boor (1987) gives a recursive formula for computation of B-splines from a lower degree B-spline Since a B-spline of degree 0 is just a constant between subsequent knots, this facilitates computation. However, B-splines can also be constructed as linear combinations of truncated power functions. P. H. C. Eilers and Marx (2010) show that it is numerically stable. We use the latter approach for computation of our B-spline basis.

10. Note that, in the B-spline literature, usually *order* is used instead of *degree*, where *order* = *degree* + *1*. In the P-spline literature however, degree is preferred, as order mostly is referring to the degree of differences used in the penalty.

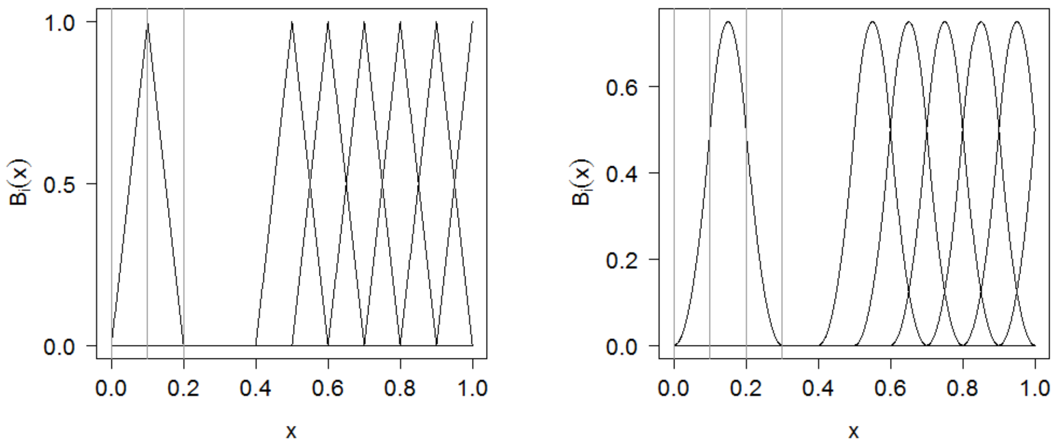

(a) B-spline of Degree $p = 1$ (b) B-spline of Degree $p = 2$

FIGURE 3.1: Illustration of a Single B-spline and the Corresponding B-spline Basis

in 3.2(a) we see that at any point x in the interval $[x_{\text{min}}, x_{\text{max}}]$, a B-spline basis decomposes 1, i.e. $\forall x \in [x_{\text{min}}, x_{\text{max}}] : \sum_{j=1}^{k+p-1} B_j(x, p) = 1$.

We denote the value of the jth B-spline (the B-spline with support interval starting at knot j) at x with $B_j(x, p)$, where p is the degree of the B-spline. For regressing the y-values of a set of N data points (x_i, y_i) on the B-spline basis $(B_j)_{j=1,\ldots,k+p-1}$, we evaluate the $k + p - 1$ B-splines at the given x-values, which yields the $(N \times (k + p - 1))$ design matrix $\boldsymbol{B}$. In matrix notation, the regression approach is to minimize

$$Q_B(\boldsymbol{\alpha}) = \|\boldsymbol{y} - \boldsymbol{B}\boldsymbol{\alpha}\|^2, \tag{3.9}$$

and the result is a fitted curve $\hat{y}(x) = \sum_{j=1}^{k+p-1} \hat{a}_j B_j(x, p)$.

A particularly useful feature of B-spline regression for our problem is established by the following formula for the m-th derivative of a B-spline function $f(x)$:

$$f^{(m)}(x) = \frac{1}{h^m} \sum_{j=m+1}^{k+p-1} \triangle^m a_j B_j(x, p - m), \tag{3.10}$$

where h is the knot distance, and $\triangle^m a_j = \triangle(\triangle^{m-1} a_j)$ with $\triangle a_j := (a_j - a_{j-1})$. For a derivation of this result, see Appendix B.3.1 or De Boor (1987, Ch. 10).

Equation (3.10) illustrates that the derivatives of a spline function can be computed conveniently by differencing its B-spline coefficients. Once a fitted curve is

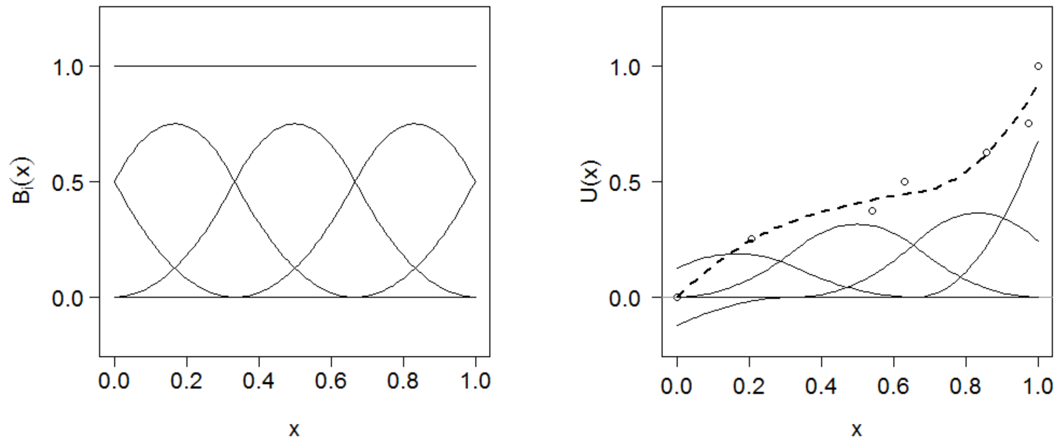

(a) B-spline Basis of Degree $p = 2$ (b) Regression on B-spline Basis

FIGURE 3.2: Illustration of Simple (Non-Penalized) Spline Regression with B-splines of Second Degree Where the Interval $[0, 1]$ is Divided Into Three Intervals

established, its derivatives are automatically obtained without the need to determine them numerically.

Although B-splines possess several advantages over other function fitting techniques, one weakness is that the fitted curve depends crucially on the knot choice. A higher amount of knots allows a higher flexibility, sometimes resulting in overfitting the data. Furthermore, knot placement also influences the fitted curve considerably.

An additional challenge is the degree of the function to be fitted: To obtain a $p-1$ times continuously differentiable curve, we need B-splines of degree p. Thus, in order to have the fourth derivative at least quadratic, B-splines of degree 6 are necessary. With barely more than 6 elicited utility points, a pure spline approach with local bases of degree 6 is impossible.

P-spline Regression P-splines solve these challenges and in addition to smoothing the utility function itself, they also smooth at least one derivative. P-splines combine the regression approach on a B-spline basis as just derived using an excessive number of (mostly equally spaced) B-splines with penalties (mostly on the curvature) to prevent the fitted curve from oscillating or fluctuating more than needed when too many knots are chosen[11]. Technically, these penalties increase

11. When introducing P-splines, P. H. Eilers and Marx state that generally the number of B-splines is moderate (10-20). In a more recent work, they note that “the size of the basis can be anywhere from 10 to over 1000”, depending on the application (P. H. C. Eilers and Marx, 2010). In

the number of conditions for the equation system to solve, and hence higher degrees can be combined with choosing a high amount of knots. Formally, the d-order differences of the B-spline coefficients is penalized by adding these differences to the objective function. We will refer to this summand as the dth-order penalty. When $\boldsymbol{D_d}$ denotes the matrix representation of the dth difference operator $\triangle^d$ (as used above), the objective function of the regression approach writes

$$Q_B(\boldsymbol{\alpha}) = \|\boldsymbol{y} - \boldsymbol{B\alpha}\|^2 + \omega\|\boldsymbol{D_d\alpha}\|^2, \tag{3.11}$$

where again, as in (3.9), $\boldsymbol{B}$ denotes the design matrix consisting of the $k+p-1$ B-splines evaluated at the x-values of the N given data points, and $\boldsymbol{y}$ denotes the vector of length N containing the respective y-values. The objective function (3.11) is minimized by $\hat{\boldsymbol{\alpha}} = (\boldsymbol{B}'\boldsymbol{B} + \omega\boldsymbol{D}_d'\boldsymbol{D_d})^{-1}\boldsymbol{B}'\boldsymbol{y}$.

The non-negative tuning parameter ω allows controlling the smoothness of the predicted function. Choosing $\omega = 0$ results in the classical linear regression of y on $\boldsymbol{B}$ as formulated in (3.9) with the well-known solution $\hat{\boldsymbol{\alpha}} = (\boldsymbol{B}'\boldsymbol{B})^{-1}\boldsymbol{B}'\boldsymbol{y}$. Therefore, as laid out above, a low value of ω will result in a fitted function that over-fits the data and possibly oscillates considerably. Depending on the order of the penalty, a linear ($d = 2$) or quadratic ($d = 3$) or more general, a polynomial function of order $d-1$ will be the result in the limit of an increasing ω (P. H. C. Eilers and Marx, 2010).

Choice of Penalty In this study and more generally in the context of the elicitation of higher order risk preferences, we are interested in smoothing the utility function itself, but also in smoothing the third and possibly higher derivatives of the utility function. Moreover, we would like to have continuous derivatives with a suitable interpolation quality at least for the derivatives of interest.

Here, we use an approach to jointly smooth the third and fourth derivative, suited for the joint elicitation of prudence and temperance, which is defined by a negative fourth derivative. This requires the balanced use of penalties of orders $d = 3$ and $d = 4$, as laid out in Appendix B.3.2. Penalization of multiple orders has been applied before in other studies with a focus on the quality of interpolation (see Marx and Eilers, 2002; Aldrin, 2006). However, both studies rely on visual inspection for determining the relation between the two penalties. In our setting, this would be too time consuming, since we consider a couple of hundreds of utility functions. Furthermore, this way of implementation could affect the classification and intensities of risk preferences and should be independent of

our application, using 15 inner knots leads to underfitting in a considerable share of cases, which is why we chose 20. This choice leads to the desired flexibility without overfitting the data.

subjective judgment. Therefore, we develop a data-driven approach to jointly use multiple orders of penalties, which we present in Appendix B.3.2 together with the details on its implementation and on the choice of the penalty orders.

Incorporating Constraints Monotonicity, or more precisely, utility as a monotone increasing function of monetary units, is a common assumption for utility functions. We follow the approach introduced by Bollaerts, Eilers, and Mechelen (2006) to incorporate a monotonicity constraint in P-splines regression. In the spirit of P-splines, this constraint is approximately enforced using a discrete, asymmetric penalty. We elucidate this approach and its implementation in Appendix B.3.3.

Normalization Due to using the certainty equivalent or the trade-off method for elicitation of utility, the points $(0, 0)$ and $(1, 1)$ are fixed and non-defective. They should therefore be exactly predicted, thus the interpolating spline function has to meet the following conditions: $f(0) = 0$ and $f(1) = 1$. We achieve this approximately by iteratively increasing weights at the points $x = 0$ and $x = 1$ until the conditions are met.[12]

Choosing the Degree of Smoothness

With perfect fidelity to the data with few data points results in a fitted P-spline function resembling a linearly interpolated function. In some cases though, it will oscillate heavily and fail to describe the true overall shape of the function. These cases are examples of overfitting: Fidelity to the data is high, but the quality of predictions is most likely poor. On the other hand, when smoothness is overweighted, the resulting function might miss important changes in curvature and also fail to describe the overall shape of the function, i.e. the function is underfitting.

In some cases it is unnecessary to determine an optimal smoothness parameter. Consider for example the case, where interest is only in the sign of a derivative of the utility function to determine whether an individual is risk averse, prudent or temperant. To that end, it might be enough to determine the sign of the respective derivative for a large number of pre-specified parameters covering a certain interval and proceed in an appropriate way. We investigate this approach further in Section 3.3.4, which is dedicated precisely to the pure classification of

12. The same effect can probably be realized with a penalty term in an iterative approach such as the one presented by Bollaerts, Eilers, and Mechelen (2006) to incorporate shape constraints. However, the existence of a solution to the minimization problem for the modified objective function has to be investigated, see Bollaerts, Eilers, and Mechelen (2006).

individuals into risk averse and risk loving, prudent and imprudent, or none of the aforementioned.

Choosing the Degree of Smoothness by Optimizing Predictive Quality When studying intensities, the precise shape of the utility curve and its derivatives are needed. Then, we wish to have a curve that perfectly balances fidelity to the data and smoothness. How should we choose a value for ω? We apply a leave-k-out cross validation (CV)—an objective datadriven decision criteria that focuses on the quality of prediction and that is robust to overfitting the data in case of correlated observations[13] (Arlot and Celisse, 2010, Chapter 8.1). Using cross validation, the model is fit with only a part of the data and the remainder, the k points left out, is used to compute prediction errors. According to CV, the smoothness parameter that minimizes the average prediction error is the preferred.

Leave-k-out CV can be seen as a mean of error correction: The more points left out when fitting the model, the more important becomes the predictive quality and the higher the smoothness parameter will be, in case some points deviate from the common trend. Since reversal rates of one third are common in choice tasks as the ones applied in this study,[14] we perform "leave-at-least-$1/3N$-out" cross validation, which results in leave-3-out CV, in case the maximum number of utility points was elicited for the individual under study.[15]

However, overfitting the data is still possible, if the distance between points is large, that is in case of sparse information per knot. The reason is that the penalized derivative of the function can change over wide intervals, thus the change needed from knot to knot to predict every point exactly may only be marginal, and is thus not sufficiently penalized. Therefore, we develop and apply a way to determine a data-driven minimum for the penalty parameter to rule out this reason of overfitting. We discuss our choice of the data-driven decision criteria and present the developed approach to rule out overfitting in case of sparse information per knot in Appendix B.3.4.

3.3.3 Intensity Measures of Risk Aversion and Prudence

Having established continuous utility functions from the elicited utility points, we can now apply well-known intensity measures of risk aversion and prudence.

13. Due to the chain structure of the experiment applied, the measurement error of single utility points might be correlated.

14. See e.g.Abdellaoui, Bleichrodt, and Paraschiv (2007), Fennema and Van Assen (1998), Abdellaoui (2000), and Etchart-Vincent (2004)

15. Note that we excluded the points $(0, 0)$ and $(1, 1)$ for computing the average prediction error.

Degree of Risk Aversion We measure the degree of risk aversion by the well-known and widely used Arrow-Pratt measure of (absolute) risk aversion, defined as $\rho(x) = -u''(x)/u'(x)$ (Pratt, 1964). For this measure, we compute the mean over the interval $[0, 1]$ based on 1000 points.

Naturally, steep increases are associated with a higher intensity than a constant, slow increase. We therefore summarize the measure of risk aversion over the interval $[0, 1]$ by taking its mean to capture such steep parts of the second derivative, even if vast parts of this derivative are actually zero. The median in such a case would be zero, which certainly is not the right measure of risk aversion in comparison with individuals exhibiting a steady, but slow increase.

Degree of Prudence We compute the measure by Kimball (1990), commonly stated as $-u'''/u''$, for (strictly) risk-averse individuals and $\pi = u'''/u'$, the measure by Crainich and Eeckhoudt (2008), for all individuals. As for the degree of risk aversion, we aggregate this information by averaging these measures over the interval $[0, 1]$.

3.3.4 Classification of Risk Aversion and Prudence

In addition to computing the intensity measures discussed above, we classify individuals as risk averse and prudent. We follow two strategies: One is based on the optimal smoothing procedure described above. Then, to account for possible overfitting, we classify a derivative as positive, if 95% of the total area between the utility curve and the x-axis are positive, and as negative, if the contrary is the case.

The second strategy resembles the popular non-linear least-squares fitting of a parametric function that 'forces' individuals to be risk averse or risk loving (and equivalently for higher order preferences), but does not leave subjects unclassified.[16] However, here, we still allow the data to 'reject' a classification, thereby following the spirit of a confidence interval.

We do so by applying the P-splines approach in (3.11) with different fixed and increasing smoothing parameters ω as opposed to a single optimal one as in the first approach. More precisely, when classifying individuals as prudent or imprudent, we apply a third-order penalty (i.e. $d = 3$) penalizing fluctuations of the third derivative and vary ω. Risk aversion is determined similarly and a

16. For the utility function $u(x) = x^b$ for example, the estimated parameter determines the sign of the second derivative and thus also determines the classification of an individual as risk averse or risk seeking (see e.g. Appendix A.2.1), independently of how good the fit actually is.

derivative is classified as positive (negative) if 95% of the total area between the utility curve and the x-axis are positive (negative).

The common pattern is that for low values of the penalty parameter, a classification is not possible due to fluctuations caused by overfitting the data. When the penalty parameter increases, classification gets more likely, until finally the data is underfit and the third derivative will vanish, as laid out above. We thus increase the penalty until for some individuals, the third derivative starts vanishing; we find that for ω of around 10,000 already a considerable share vanishes and set this as upper limit for the smoothing paramter. The lower limit is set to $.1$ and this results in virtually no smoothing.

If the classification changes with increasing penalty between risk averse and risk seeking, or prudent and imprudent, we set the respective classification to non-classifiable. Conveniently, the classification is relatively robust with respect to the smoothing parameter, so this is the case for only very few individuals.

The interpretation of the second classification approach thus is: Given enough smoothing, which classification would be the most fitting while still allowing to reject a classification? It is in some sense an alternative to confidence intervals, which we cannot compute in our setting, since standard error bands rely on asymptotic arguments. These asymptotic arguments are not likely to hold with less than 10 observations.

In some cases, we are unable to visually distinguish a derivative from a straight line. Its corresponding derivative however will never be a constant, since the numerical 'precision' of the procedure is too high. In those cases, where we are unable to distinguish the first derivative from a straight line, we decrease the precision artificially and 'snap' its second derivative, i.e. the third derivative of the utility function, to zero.[17] Formally, this means that we set a derivative to zero if its value is smaller than approximately the biggest absolute value of that derivative observed for any individual divided by 1000 (for a fixed $\omega = 50$).

3.4 Data and Definition of Variables

Our main data was collected as part of a larger study 'Savings for the Old Age'. The survey gathered financial information for a sample of 1200 subjects beneficiary of the social protection program SISBEN. The program targets the population in low social strata. We recruited the participants in a two-step procedure. First, we selected neighborhoods based on shares of the population belonging to

17. The utility curve itself is non-linear and clearly distinguishable from a straight line for all individuals in our sample.

the required strata. Neighborhoods that contained a large fraction of low income individuals and that were assessed as safe for the team to visit, were included in the study. In the second step, enumerators visited randomly selected households and verified whether they were belonging to the target group. If this requirement was not fulfilled, the enumerators visited the neighboring household. We oversampled older females to obtain a better picture of the potential vulnerabilities women are exposed to in the old age. Interviews took place in October and November 2013 and lasted on average 90 minutes.

The survey consisted of 16 sections on general demographics, wealth, general savings, pension savings, financial literacy, health behavior, expectations and hypothetical questions on psychological traits. The experiment was conducted with a subsample of 693 participants. A team of enumerators conducted the experiment on tablet computers. To meet safety demands, the experiments were conducted in a public space (e.g. communitary houses) that was easily accesible for participants. The experiment lasted around 20 minutes.

3.4.1 Net Savings

Savings In the survey we asked detailed questions on a comprehensive range of saving devices: housing, savings plans, savings and checking accounts, certificates of deposit, mutual funds, loans given out and savings in Colombian Pesos or other currencies. We use the sum of these assets to construct a savings measure denoted S_h. This variable intends to capture all liquid assets in the household.[18]

Debt Debt is defined as all financial liabilities a household has against other households, enterprises and financial institutions, including money lenders. We denote total debt as D_h.

Net Savings Following Noussair, Trautmann, and Kuilen (2014) and Fuchs-Schündeln and Schündeln (2005) we use *net savings* as our main variable of interest, which we calculate as the difference between S_h and D_h.

3.4.2 Income Uncertainty Based on Economic Activity

In order to obtain measures for the financial riskiness of the sector subjects were working in we collected measures from the Commercial Register in Bogota in 2013. We calculate for 14 sectors the empirical probability of firm closure within

18. As it is possible to withdraw money from all of these saving mechanisms in case of emergencies, we interpret all of these savings as liquid assets.

each sector. We match this then with the sector subjects are usually working in, which has the same resolution as the official statistics. This measure of working sector dynamics proxies the degree of income uncertainty and hence measures the idiosyncratic risk for which subjects like to have precautionary savings.

This measure relies on working sectors according to economic activity[19] and bases on data from the Commercial Register of Bogota from the year 2013.[20] More specifically, for all working sectors represented by individuals from our sample, we computed the ratio of cancelled to existing businesses (including microenterprises) in Bogota in 2013. Then, all individuals are assigned the corresponding ratio of the working sector they are working in.

Thus, a higher level of plant closure in the sector a subject is typically employed is associated with a higher level of income risk. As unemployment insurance is negligible in our sample, firm closure in the sector constitutes a major threat to personal income. This constitutes a substantial and unavoidable background risk.

In this sense our sample and the economic circumstances resemble the economy described in Aiyagari (1994). There are idiosyncratic earnings uncertainties, effectively no insurance markets and our subjects are borrowing constraint. Background risk, defined as uninsurable idiosyncratic shocks, is arguably substantial and not mitigated by wage insurance on the firm level as found by Fagereng, Guiso, and Pistaferri (forthcoming 2017).

Our measure is closely related for example to the measure of income uncertainty applied by Fagereng, Guiso, and Pistaferri (2017) when studying precautionary saving in Norway: They infer firm performance volatility from companies' balance sheets and use this to instrument their measure of income uncertainty. This limits their sample to employees of private firms with balance sheets available to the public and is arguably not suited in our context. We therefore use the next aggregation level, namely working sectors, since on this level, we are able to link individuals to secondary data on economic performance.

The arguments in favor of exogeneity of firm performance volatility given by Fagereng, Guiso, and Pistaferri (2017) are valid also for our measure: First, they argue, firm shocks are hard to avoid for most workers, and second, firms pass over this variation onto their workers' wages. Clearly, in our case, negative working sector dynamics will be handed over to employees due to absence

19. Working sectors as provided in ISIC Rev. IIIa A.C. by DANE, based on the economic activity.

20. Data was processed and made available in an 'Overview of Indicators' by the Knowledge Management Board of the Chamber of Commerce (in Spanish: *'Tablero de Indicadores', Dirección de Gestión de Conocimiento, Cámara de Comercio de Bogotá*).

of formal contracts as well. Thus, payment can be adjusted easily, but also employment itself is more uncertain in working sectors with a higher share of closed businesses. Regarding the first argument, we acknowledge that it is theoretically possible to switch the working sector after bad days in order to avoid income uncertainty. However, the elasticity is certainly lower than the elasticity of hiring at a different firm. Furthermore, the categorization of working sectors we use relies on the main task performed, where arguably the share of workers switching from e.g. servant to construction worker to avoid income uncertainty is considerably lower than the share of servants increasing their saving to cope with income uncertainty. Moreover, switching the working sector will most likely take more time than adapting consumption behavior.

Time Preferences We followed the experimental design by Andersen et al. (2008) to elicit time preferences: Participants decided on receiving an amount x in 30 days or an amount $x(1 + r/12)$ with $r > 0$ in 60 days. Values of r were increased gradually and subjects usually switched from choosing x in 30 days to $x(1+r/12)$ in 60 days for some r according to their time preferences. Using this switching point, we calculate a lower and an upper bound for the interest rate. This interest rate can also be interpreted as impatience, as people were deciding about the timing of receiving money.

We repeated the task with a higher delay of payment: Subjects now decided about receiving the lower amount in 180 days or receiving the higher amount in 210 days. Similarly to the case for the near future time frame, we deduce a lower and an upper bound of interest rate or impatience from the switching point.

The difference between both interest rates or impatience for the two time frames informs about consistency in interest rates. For individuals deciding consistently, the impatience to receive a monetary amount 30 days earlier should be unaffected by shifting the date of the earlier payment by 150 days. A lower impatience in the more distant future corresponds to an increasingly patient subject.

3.4.3 Further control variables

In the style of comparable, previous studies on precautionary saving (e.g. Fuchs-Schündeln and Schündeln, 2005; Noussair, Trautmann, and Kuilen, 2014), we control for other socioeconomic factors within our analysis such as age, gender, number of adult household members, number of children in a household, education and income (we use average per capita household income). To these, we add further characteristics that have been found to be important in explaining savings:[21]

Moreover, we measured financial literacy within the survey and include its result in the analysis.[22] Van Rooij, Lusardi, and Alessie (2012) find a positive correlation between financial literacy and accumulated savings. Devaney, Anong, and Whirl (2007) and Fisher and Montalto (2010, 2011) report short term planning and saving horizons (i.e. time preference for the present) having a negative effect on the likelihood of saving and net wealth. In our analysis, we use an experimental measure of impatience, time inconsistency with respect to impatience and the planning horizon with respect to financial decisions. Furthermore we calculated the BMI from weight and height of subjects. The BMI serves as a proxy for temptation and self control (Hofmann, Friese, and Roefs, 2009; Moffitt et al., 2011). It also serves as a proxy for health status that is positively associated with the likelihood of saving (Fisher and Montalto, 2010, 2011).

3.5 Results

The section is organized as follows: we first give a characterisation of our sample in terms of socioeconomic characteristics, risk and time measures. Then we use these measures in order to explain net savings and finally we model individual precautionary savings motives, i.e. income uncertainty, and relate them to net savings.

3.5.1 Descriptive Statistics

We obtain the full set of household characteristics from 680 subjects of which 72% were female. The mean age was 49 years and spanned from 24 to 87 years. The median education level is primary school. In the financial literacy test, subjects answered on average 9.29 questions out of 18 correctly. The average BMI is slightly above 25, the threshold to mild overweight. The average income per household member is 319 thousand Colombian Pesos (COP) and the average debt is 1.64 million Pesos, which leaves an average of -1.36 million Pesos in net savings.

21. Some studies focused on the likelihood of saving, others on the amount of saving. Since we use the same control variables for estimating the likelihood and the amount of saving, we include a variable (if possible) that has been found to either affect the likelihood of saving or the amount or both.

22. In total, we were asking 18 questions on financial literacy concerning interest rate, asset classes, basic math and financial math. The variable included in the regression corresponds to the number of correctly answered questions.

55 percent of our sample has neither savings nor any debt and average net-savings amount to -1,361,000 COP (710 USD). Around 85 percent of our population has no savings and average savings in the sample are 276,000 COP (140 USD), which is less than a month's average per capita household income. Around 27 percent of those reporting non-zero savings save exclusively in cash, another 20 percent save exclusively using other savings technologies. About 34 percent are saving exclusively for housing, of which roughly the half uses a special fund, whereas the other half uses any form of saving device. Average debt amounts to 1,637,000 COP (850 USD) and 38 percent of our sample hold positive debt. Table 3.1 provides an overview of the demographic characteristics.

TABLE 3.1: Summary Statistics

	Mean	Med.	s.d.	Min	Max	Obs.
Male	0.28	0.00	0.45	0.00	1.00	680
Age	48.87	49.00	13.43	24.00	87.00	680
Education	2.48	2.00	0.70	1.00	4.00	680
Financial literacy	9.29	10.00	3.39	0.00	16.00	680
BMI	25.75	25.51	4.31	12.89	42.97	680
Adult HH members	2.85	3.00	1.41	1.00	12.00	680
Children HH members	1.21	1.00	1.29	0.00	7.00	680
Income	3.19	2.77	2.25	0.01	18.00	680
Savings (100k)	2.76	0.00	15.18	0.00	200.00	680
Debt (100k)	16.37	0.00	61.85	0.00	588.04	680
Net savings (100k)	-13.61	0.00	63.94	-588.04	187.00	680
Zero net-savings	0.55	1.00	0.50	0.00	1.00	680

3.5.2 Risk Aversion, Prudence and Time Preferences

TABLE 3.2: Summary Statistics of Risk and Time Measures

	Mean	Med.	s.d.	Min	Max	Obs.
Risk Aversion (A&P)	0.03	-0.01	1.17	-2.43	2.86	588
Prudence (C&E)	7.24	7.61	6.28	-0.50	25.70	588
Prudence (Kimball)	1.59	2.13	1.03	-2.22	3.13	126
Impatience	29.60	22.00	15.41	16.00	52.00	693
Increase in patience	0.13	0.95	16.35	-38.90	36.95	693

Table 3.2 shows the Arrow-Pratt coefficient of risk aversion, measures of prudence and time preferences. The distributions of the measures of prudence is shown in Figure 3.3. The mean annual interest rate r subjects asked to receive

FIGURE 3.3: Histograms of Prudence Measures

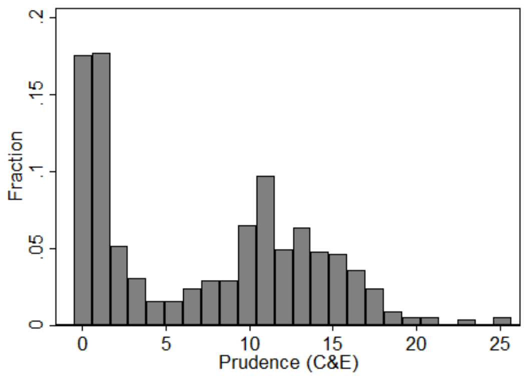

an amount $x(1 + r/12)$ in 60 days instead of an amount x in 30 days is 29.6 percent. This figure is in the range of estimates from recent experiments with the general population in Denmark (Harrison, Lau, and Williams, 2002), which seem best suited for comparisons since our data is from the general population as well. On average, this interest rate or mean impatience stays approximately constant when the timing is changed to receiving the monetary amounts in 180 or 210 days, respectively.

We obtained a full set of risk preference measures for 588 subjects. Table 3.3 gives an overview over the classification of risk and prudence as measured based on optimally smoothed P-spline regression. We observe that all combinations of risk aversion and prudence attitudes are present, confirming previous findings by Noussair, Trautmann, and Kuilen (2014) that even risk lovers can be prudent.[23] 48 percent of the subjects show risk aversion[24] and roughly 60 percent are prudent. The most unlikely combination, with below 3 percent, is being risk loving and imprudent. So utility functions that require risk-loving subjects to be imprudent are not sufficiently flexible to describe our data. When considering the strength of prudence, Table 3.4 reports a substantial correlation between the Kimball measure of prudence and risk aversion, while we do not see this correlation for the C&E measure—illustrated in Figure 3.4. This difference is important for empirical work, as the precautionary motive for saving is driven by the strength of prudence. So in order to attribute savings decisions to a precautionary motive makes it necessary to estimate these concepts separately.

23. Crainich, Eeckhoudt, and Trannoy (2013) theoretically show that prudent risk lovers devote all their income to saving.
24. When classifying individuals based on the fitted coefficient of a power utility function, around 2/3 of the sample are to be considered risk averse.

TABLE 3.3: Classification of Risk Aversion and Prudence

	Risk averse pct	Risk loving pct	Mixed pct	Total pct
Imprudent	4.76	2.55	0.68	7.99
Prudent	29.25	26.87	1.87	57.99
Mixed	14.46	19.05	0.51	34.01
Total	48.47	48.47	3.06	100.00

Notes: This table reports the share of risk-averse and prudent individuals according to the classification described in Section 3.3.4. The measures were computed using optimally smoothed spline functions, evaluated at and averaged over 1000 points in the support. Risk neutrality and prudence neutrality is a probability zero event, so none of our subjects was classified and we omit this category.

TABLE 3.4: Correlation of Risk Aversion, Different Measures of Prudence and Time Preferences

	Risk Aversion (A&P)	Prudence (C&E)	Prudence (Kimball)	Impatience
Risk Aversion (A&P)	1			
Prudence (C&E)	-0.0373	1		
Prudence (Kimball)	0.671***	0.757***	1	
Impatience	0.0275	0.0308	0.0930	1

Notes: This table reports the correlation coefficients between risk aversion and the prudence measures—calculated based on optimally smoothed spline functions—as well as time preferences. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

3.5.3 Precautionary Savings and Prudence

We now turn to the relationship between our experimental measure of prudence and wealth. As we laid out in Section 3.2, we expect to observe a positive relationship between our measure of the strength of prudence and people's accumulated wealth—given present or past income uncertainty. We run OLS regressions on net savings on our preferred Crainich and Eeckhoudt-measure of risk attitudes.[25] We use several sets of control variables motivated by previous studies (e.g. Noussair, Trautmann, and Kuilen, 2014). We report the results in Table 3.5.

In all specifications including both the risk averse and the risk loving we find a significant positive relationship between prudence and wealth. When analyzing both subgroups independently, we find a significant positive relationship

25. Robustness checks applying the measure by Kimball (1990) can be found in Appendix B.4.

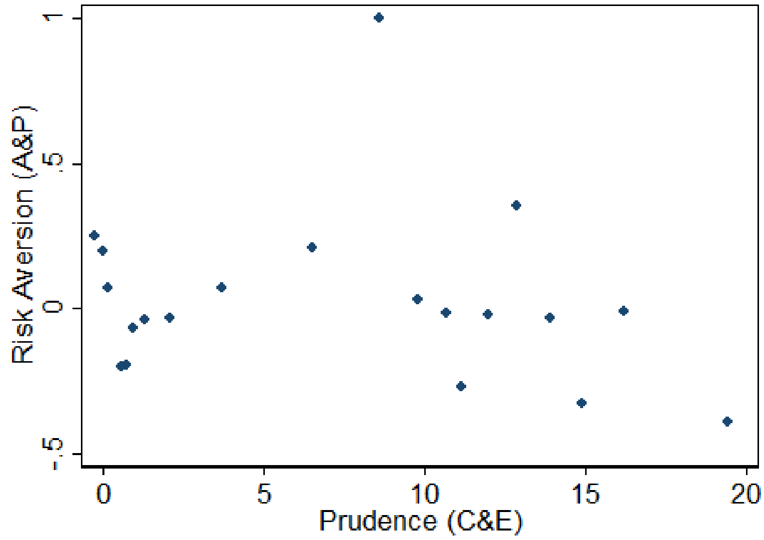

FIGURE 3.4: Bin Scatterplot of Risk Aversion and Prudence

TABLE 3.5: Net Savings and Prudence (C&E)

	Full Sample			Risk Averse		Risk Loving	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Prudence (C&E)	1.140***	1.161***	1.114**	1.707***	1.468***	0.560	0.651
	(0.422)	(0.422)	(0.431)	(0.641)	(0.559)	(0.601)	(0.724)
Risk Aversion (A&P)		3.329*	3.291		0.637		6.038
		(2.002)	(2.134)		(3.120)		(5.325)
Controls	No	No	Yes	No	Yes	No	Yes
Observations	567	567	554	270	267	279	271

Notes: This table reports the results of ordinary least squares regressions on net savings. Classification of risk aversion according to Section 3.3.4. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. Coefficients of controls can be found in Appendix B.4. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels $^{\star}\ p < 0.10$; $^{\star\star}\ p < 0.05$; $^{\star\star\star}\ p < 0.01$.

only for the risk-averse subsample. When excluding those that neither save nor are indebted, however, the relationship is significantly positive also for the risk-loving (see Table B.3 in Appendix B.4). Furthermore, as reported in Appendix B.4, we find a positive relationship between the length of the planning horizon and wealth. People who have a longer planning horizon than the next day have significantly higher wealth. Financial literacy is surprisingly negatively related to holding higher levels of wealth.

3.5.4 Income Risk

We now turn to the analysis of the impact of individual income risk, prudence and net savings. Several previous empirical studies have identified prudence parameters from consumption volatility (Guiso, Jappelli, and Terlizzese, 1992; Dynan, 1993; Fagereng, Guiso, and Pistaferri, 2017). Fagereng, Guiso, and Pistaferri (2017) instrument consumption volatility with firm specific shocks that pass through to wages. We construct a similar measure of income risk by looking at firm closures by sector, see Section 3.4. Figure 3.5 shows the distribution of our measure of income risk aggregated at the *Localidad* level.

We regress net income on the prudence measure interacted with the probability of firm closure. The empirical model can be written as follows:

$$W = \alpha + \beta_1 \text{ Prudence } + \beta_2 \text{ Shock } + \beta_3 \text{ Prudence } \times \text{ Prudence } + \boldsymbol{\beta_4 X} + \epsilon \tag{3.12}$$

This allows us to answer the question, whether prudent subjects who are confronted with a higher background risk accumulated higher levels of wealth as predicted by our theoretical framework. Results are presented in Table 3.6.

Column (1) in Table 3.6 shows the raw correlations. The main effect of prudence is positive and highly significant and so is the interaction term of income risk with prudence. Hence, more prudent people save more when facing higher income risk. This effect is robust to the inclusion of controls in column (2). When restricting the sample to only risk averse agents, the coefficients on the interaction term stay positive, however they are not significant at conventional levels for the interaction term, unless controls are included. For the risk-loving agents, however, we observe a positive and significant coefficient on the raw correlation of the interaction term. After including control variables, this coefficient is only significant if focusing on those that are saving or are indebted (see Table B.6 in Appendix B.4). When using the Kimball measure of prudence we get large, but insignificant coefficients (see Table B.4 in Appendix B.4).

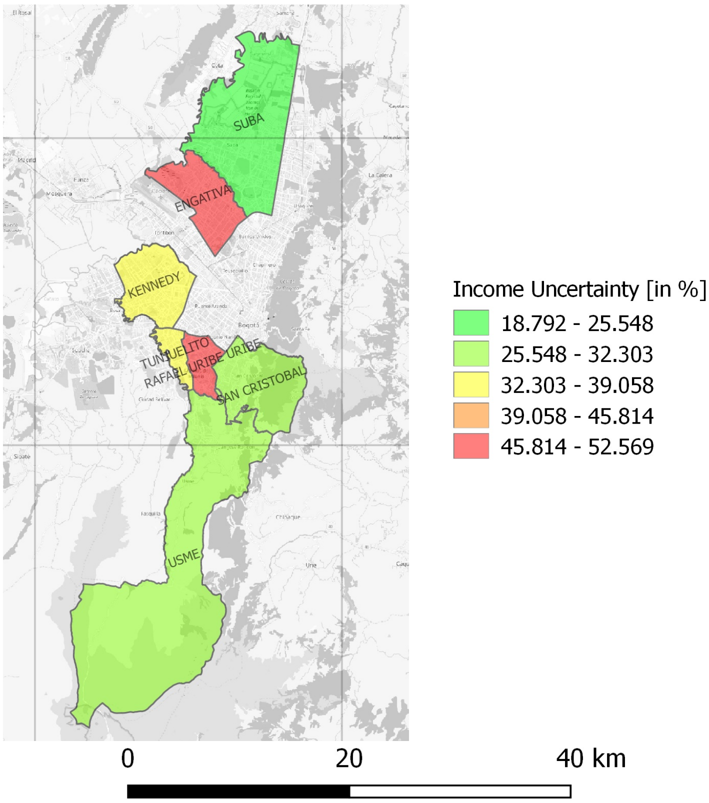

FIGURE 3.5: Income Uncertainty in Bogota: Ratio of Closed to Existing Businesses in 2013.

Notes: Individuals are categorized based on the economic activity they perform according to the ISIC Rev. IIIa A.C. categorization as used e.g. by DANE for their household survey. At this aggregation level, official data on firm closure is available e.g. from the Knowledge Management Board of the Chamber of Commerce of Bogota. Participants are assigned a level of income uncertainty corresponding to the ratio of firm closure in the working sector they are classified.

TABLE 3.6: Net Savings, Firm Closures and Prudence (C&E)

	Full Sample		Risk Averse		Risk Loving	
	(1)	(2)	(3)	(4)	(5)	(6)
Income risk	-1.746	-1.479	-1.863	-1.399	0.768	2.146
	(2.354)	(2.405)	(2.870)	(3.031)	(3.293)	(3.412)
Prudence (C&E)	1.240***	1.115**	2.224***	2.036***	0.298	0.199
	(0.461)	(0.471)	(0.769)	(0.739)	(0.613)	(0.660)
Prudence (C&E) × Income risk	1.015**	1.066**	0.856	1.037*	0.984*	0.925
	(0.424)	(0.416)	(0.600)	(0.566)	(0.588)	(0.566)
Risk Aversion (A&P)		4.644*		2.231		9.500*
		(2.408)		(4.460)		(5.666)
Controls	No	Yes	No	Yes	No	Yes
Observations	471	459	218	215	237	230

Notes: This table reports the results of ordinary least squares regressions on net savings. Classification of risk aversion according to Section 3.3.4. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. Income risk is measured as the ratio of closed to existing businesses in 2013 in the working sector an individual was usually working in at the time of the survey. Prudence and income risk are centered. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. Coefficients of controls can be found in Appendix B.4. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Fuchs-Schündeln and Schündeln (2005) show that in Germany there is sorting into risky sectors according to risk aversion and line out its implications for precautionary savings motives. We do not find a correlation with the measure of risk aversion and the sector where subjects are employed. Moreover, we are able to control for risk aversion.

3.6 Conclusion

In this study, we have investigated the connection between prudence, income riks and savings. We have demonstrated how the intensity measure for prudence by Crainich and Eeckhoudt (2008) is linked to the strength of the precautionary saving demand for risk-averse and risk-loving individuals building on the model by Leland (1968). Moreover, we have shown that especially for the latter, a saving rate strictly higher than 0 and strictly lower than 1 may be optimal, contrarily to the predictions of previous models (Crainich, Eeckhoudt, and Trannoy, 2013). To test the theoretical predictions from this model, we have introduced a new method to elicit higher order risk preferences including their intensities non-parametrically.

Applying our method to a sample of poor households in Bogota, Colombia we find comparable results with respect to prudence as for example Tarazona-Gomez (2004) in her experiment with students in Bogota. The results regarding risk aversion are also in line with previous findings (Tarazona-Gomez, 2004), although we find a higher share of risk-loving individuals.

We find strong support for the theory of precautionary saving, including our extension: According to the theoretical framework, prudent individuals react to income uncertainty by raising their saving. This should lead to higher savings, and those who are more likely to face income uncertainty should hold higher savings. Those who are more prudent than others with respect to different intensity measures should hold higher savings, irrespective of them being risk averse or risk loving, as we have shown.

This relationship can be found in the data and it is robust—even when we pool risk-averse and risk-seeking individuals. Of course, our finding relies on correlational data only and we cannot draw causal conclusions. However, given that we use experimental data for prudence, secondary data for income uncertainty and many control variables that have been found to influence savings before, we think our findings lead into the right direction.

Although our results are in line with previous findings, it would be interesting to see how our method compares to the more complicated risk apportionment tasks going back to the definition of higher order risk attitudes by Eeckhoudt and Schlesinger (2006) on the same sample, depending on the cognitive ability of participants. To that end, it would be helpful to see an implementation of these tasks that is less prone to probability weighting or the application of reference points and thus as model free as the definitions itself.

We find prudence to be strongly linked to a higher level of net savings. While preferences like risk aversion and time preferences have failed to explain a low level of wealth empirically, prudence seems to be of importance. Further studies are needed, especially in order to establish causality. Even without causality, political interventions aiming to use prudence in order to advance the population's well-being might be fruitful. Moreover, our results suggests that the sample under study has a high demand for consumption smoothing and would thus profit from a suitable solution.

Chapter 4

The min MSE Treatment Assignment Method

This chapter is joint work with Martin Schlather.

4.1 Introduction

From the early days of experiments involving random treatment allocation, researchers have thought about the methodology used to ensure the decisive fundament of every experiment—similar treatment groups in the absence of the treatment. When discussing Darwin's experiment comparing the growth of crossed and self-fertilized plants, Fisher (1935) argues that it is not enough to randomly assign plots to the treatment or the control, i.e. to crossed or self-fertilized plants. He explains that an unbiased experiment alone does not "ensure the validity of the estimates, [...] for it might well be that some unknown circumstance, such as the incidence of different illumination at different times of the day [...], might systematically favour all the plants on one [plot] over those on the other."

Had the experiment have been repeated several times, or had a large number of plots been used, any difference caused by those "unknown circumstances" would have been diminished in expectation. The same is of course true for known or observed circumstances.

Fisher noticed that an increase in the *precision* of the experiment could also be achieved differently—by making the groups more similar. His suggestion to achieve both the validity of the estimates *and* an increase in precision was to change the level of randomization by allowing plants from both groups to be planted in the two available plots, thus being exposed to the circumstances of both plots.

Hence, not only from the perspective of validity, but also with respect to efficiently estimating the size of an outcome of an experiment, it is desirable to have similar or *balanced* treatment groups. More recently, a motivation for seeking balance across treatment groups comes from the interest in subgroup analysis. Finally, a quantity often estimated in the recent impact evaluation literature—the conditional average treatment effect—can formally only be estimated if the so called *overlap condition* (Abadie and Imbens, 2006) is satisfied, which can be described as a weak criterion of balance.

In this paper, we present a new method to assign experimental units to treatment groups in a way that the resulting groups are balanced. The method falls in the category of rerandomization methods and builds on a theoretically derived statistic that aims at balancing the second moments of the covariate distributions and incorporates dependencies between covariates. Moreover, within a model where the conditional average treatment effect depends linearly on the covariates, the mean squared error of the estimator for this treatment effect is minimized.

When relying on randomization alone, balanced groups are not ensured. Following Fisher's suggestion as an example, it could—by pure chance—be that all self-fertilized and all crossed plants are still allocated to separate plots.

Therefore, it has long been recognized that group characteristics or *circumstances* should be accounted for when assigning the treatment to experimental subjects and subsequently analyzing such experiments, see e.g. Cox (1957) for an early review of the possibilities to do so. Today, several strategies are widely used, although there is no consent on how treatment assignment should be carried out, even among experts in field research (see e.g. the survey by Bruhn and McKenzie, 2009). Also from a theoretical perspective, a clear answer is missing; see e.g. Imbens (2011) for a brief discussion.

Stratification or blocking goes back to Fisher (1935). The idea is to build subgroups according to observable characteristics and to randomize within those subgroups. Although this improves 'balance' in comparison to purely random treatment assignment, it is impractical in several aspects. Using stratification, it is only possible to balance a very limited number of variables. Furthermore, continuous variables have to be arbitrarily discretized and are never really balanced with this approach. Additional problems arise in implementing this method when the number of participants is not divisible by the number of subgroups.

Pairwise matching is often seen as the limit case of stratification, when the subgroups consist of only two individuals. The subgroups, called *pairs* in the case of matching, have to be created[1] such that the two individuals are similar, where the similarity can be measured e.g. with the so-called Mahalanobis distance of the covariate vectors of the two individuals. Two types of algorithms are commonly used: the so-called greedy algorithm (Imai, King, and Nall, 2009) and an 'optimal matching' algorithm (Greevy et al., 2004; Lu et al., 2011). Matching can be realized with many possible continuous variables and thus eliminates some of the shortcomings of stratification. This, however, comes at the cost of analytical difficulties when estimating the variance of the treatment effect (e.g. Imbens, 2011; Abadie and Imbens, 2006; Klar and Donner, 1997). Additional problems arise when attrition occurs, i.e. when for some units the outcome finally is unobserved, especially in small samples or when performing randomization at the cluster level: For every unit, possibly consisting of many individuals, dropping out of the experiment, its pair should also be removed, which lowers the sample size and power and can be of major concern. Additionally, we are unaware of an existing approach to extend matching to multiple treatment arms. Furthermore,

1. Note that this is a different task to the one performed for matching in observational studies: Finding pairs when groups have already been formed is far less demanding, also from a computational aspect.

matching can only be performed when the number of units is even. Finally, the matching approach implemented by Bruhn and McKenzie (2009), needed several days to conduct treatment assignment with a sample size of 300 units, so this approach is inappropriate if time is a limiting factor.

Several so called rerandomization methods have evolved, probably because of the theoretical or practical limitations of the abovementioned approaches. The basic idea of rerandomization is to pick a random treatment assignment in some way, evaluate it with respect to a certain criteria and rerandomize until this criteria meets some condition to be specified or to rerandomize a certain number of times and choose the best assignment, according to a specified evaluation criteria. Sometimes, subjective judgment is also used (Bruhn and McKenzie, 2009). However, we are aware of only one rerandomization approach, the one by Morgan and Rubin (2012), that relies on a theoretical derivation of the statistical threshold to stop the rerandomization. This threshold, as well as the alternative ad-hoc thresholds, such as picking the maximum t-value minimizing treatment group assignment, focuses only on the mean value of one or serveral covariates, ignoring other dimensions of the distributions of the variables to be balanced. Irrespectively of this limitation, we are unaware of a software implementation of this approach or an extension to multiple treatment arms.

Kasy (2016) applies a decision theoretical, Bayesian approach to determine treatment assignment. To that end, he derives the posterior mean squared error (MSE) of an estimator for the conditional average treatment effect of interest as a function of treatment assignment. The posterior MSE, i.e. the sum of bias and variance, is then to be minimized across treatment assignments. When the estimator is modeled with a linear model, this leads to a decision criterion that balances not only the mean of the variables of interest, but also partial correlations.

Based on the introduced decision theoretical framework, Kasy (2016) argues that a deterministic assignment rule is superior to any random assignment in terms of minimizing the MSE.

However, the drawbacks of this method are the limitations to only one treatment group and the number and nature of the parameters used; in that to apply this approach, the researcher has to specify a mean vector and a covariance matrix of the regression coefficient vector in a linear model explaining the potential outcomes as a function of covariates. In addition, a guess for R^2, the coefficient of determination, of that linear regression model must be specified. Apart from that, a non-random treatment assignment rules out the possibility to perform a conventional randomization inference. To date, the only software implementation we are aware of is in Matlab (Kasy, 2016).

We develop the approach in a frequentist approach, thereby simplifying the method considerably for the case when potential outcomes are modeled as linear functions of covariates. We suggest a way to extend Kasy's (2016) method to assign multiple treatments. For the developed method, we provide a software implementation as an ado-package for Stata and thus increase its usability in the respective areas of application. Apart from that, we interpret and implement the method as a rerandomization method, which yields the possibility of randomization inference.

Bayesian modeling allows for a greater flexibility in many cases because it relies on distributions instead of parameters. In this case, however, at least when using a linear model, disproportionately many parameters have to be specified, as just explained. We think that even for experienced researchers, it is hard to come up with a reasonable guess on these parameters. Of course, one could use a *flat* prior, inducing nearly no prior information. In this case, however, one can also resign from using prior information, as it simplifies the objective function and consequently the method considerably.

Therefore, we introduce the approach in a frequentist setting. This means that we only get a point estimate instead of a distribution for any result. As the method is designed to minimize the MSE (a point estimate), this comes without limitations.

An advantageous side effect is that we can factor out variances of the decision criteria and thus, it is sufficient to specify ratios of variances relative to a base variance. The assumption of equal variances is an intuitive assumption that experienced researchers quickly can confirm or withdraw, and in the latter case, easily adjust by specifying a good guess for scaling up the variance of a treatment or an outcome.

Our result works without choosing any technical parameters while still allowing for the needed flexibility. In the treatment assignment mechanism derived here, the only parameter that must be specified by the researcher is the number of treatment groups desired; other parameters can be specified, but can be left constant unless a better guess is available. These default values have an intuitive interpretation and are not chosen by us, but follow from the theoretical derivation of the method as laid out in this paper.

Furthermore, this method can be applied irrespective of the number of treatment arms, the number of units in the experiment and its relation to the number of treatments and variables (even or uneven, divisibility by the number of treatments, ...). Another feature is its speed: Compared to the Bayesian approach, but also compared to competing methods of treatment assignment, a reasonably good balance for a sample size of 100 units and 10 variables is usually achieved in less than 5 minutes on a 2.3 Ghz dual CPU.

In a simulation study similar to the one by Bruhn and McKenzie (2009), we compare the performance of our min MSE method in various dimensions to competing methods and find that it is comparable to the matching methods and superior to stratification or pure randomization. In addition, the min MSE method is tolerant of attrition, i.e. of units whose outcome finally is unobserved, for example because treatment is never received although it was planned.

The structure of this paper is as follows: Section 4.2 introduces our approach to treatment assignment. Section 4.3 describes theoretical characteristics of the min MSE approach that are put into context of alternative methods of treatment assignment in Section 4.4. Section 4.5 explains the design of the simulation study and its implementation. Section 4.6 reports the results. Section 4.7 discusses our method and concludes.

4.2 The min MSE Treatment Assignment Mechanism

4.2.1 Estimation of the Treatment Effect

First, we define the parameter we are interested in estimating: the conditional average treatment effect. We do so by introducing the potential-outcome framework (Rubin, 1974, 1977), as this is the standard notation in the literature on program evaluation (Imbens, 2004). As we derive the minimizing MSE treatment assignment procedure for various treatment effects and various outcomes, we directly extend the framework to fit our needs.

Assume, we have N participants, randomly selected for the experiment from the population. Individual draws of a (random) variable are indicated with a subscript $i = 1, \ldots, N$ and realizations of a random variable or vector will be denoted by the corresponding lower-case letter.

In the experiment, each individual is randomly assigned to an experimental group and treated with the corresponding treatment or not treated at all if assigned to the control group.

Definition 1. Assume we have n_d treatments indicated by $1, \ldots, n_d$ and a control denoted by 0. Let $D_1, \ldots, D_N$ be random variables with values in $\{0, 1, \ldots, n_d\}$. Then, the vector $D = (D_1, \ldots, D_N)^\top$ is called a (random) *treatment group assignment*.

Irrespective of the treatment group assigned to, each participant has potential outcomes, observed outcomes and a vector of pretreatment information, which we call covariates.

Definition 2. Let $X = (X_{j,i})_{j=1,\ldots,m;i=1,\ldots,N}$ be a random matrix. Then the vector $X_i = (X_{1,i}, \ldots, X_{m,i})^\top$ is called the *vector of covariates* of individual i.

Definition 3. Let $Y_i^p = (Y_{i,t}^{p,k})_{t=1,\ldots,n_d;k=1,\ldots,n_y}$ be a random matrix for $i = 1, \ldots, N$. Then the row vector $Y_{i,t}^p = (Y_{i,t}^{p,1}, \ldots, Y_{i,t}^{p,n_y})$ is called the vector of *potential outcomes* of individual i in the case of treatment t, where n_y is the number of outcomes of interest.

These *potential outcomes* of individual i in the case of treatment t exist irrespective of whether individual i was actually treated with treatment t or not. However, for every unit and outcome of interest, we only observe the *realized outcome*.

Definition 4. Let $Y^r = (Y_i^{r,k})_{i=1,\ldots,N;k=1,\ldots,n_y}$ be a random matrix. Then the row vector $Y_i^r = (Y_i^{r,1}, \ldots, Y_i^{r,n_y})$ is called the vector of *realized outcomes* of individual i.

The realized outcomes Y_i^r of individual i can be written by means of potential outcomes and the treatment group assignment:

$$Y_i^r = \sum_{t=0}^{n_d} \mathbb{1}_{\{D_i=t\}} Y_{i,t}^p = Y_{i,0}^p + \sum_{t=0}^{n_d} \left(Y_{i,t}^p - Y_{i,0}^p\right) \mathbb{1}_{\{D_i=t\}}.$$

The right-hand side of the above formula decomposes the realized outcomes for an individual in her potential outcomes. The differences $Y_{i,t}^p - Y_{i,0}^p$, which are the causal effects of the treatment t, would be of great interest in any study, but can never be observed.

However, under certain conditions, we can estimate the population average effect of treatment t:

$$\tau_t = \mathbb{E}\left[Y_{i,t}^p - Y_{i,0}^p\right], \quad \text{for all } t = 1, \ldots, n_d,$$

which—depending on the question—is often sufficient.

If the main interest is to study a subpopulation (e.g. the poor), or when one is not sure whether or not the sample at hand is representative for the population, one should focus on the *conditional* average treatment effect (Imbens, 2004). This happens frequently in Development Economics, for instance.

Definition 5 (Conditional Average Treatment Effect)**.** Let X, $Y^p_{i,t}$ and $Y^p_{i,0}$ as in Definition 2 and 3. For every treatment $t \in \{1, \ldots, n_d\}$,

$$\tau_t(X) = (\tau_{t,1}(X), \ldots, \tau_{t,n_y}(X)) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}\left[Y^p_{i,t} - Y^p_{i,0} \,|\, X_i\right]$$

is called the *conditional average treatment effect* of treatment t. The random matrix $T = (\tau_{t,k})_{t=1,\ldots,n_d;k=1,\ldots n_y}$ contains all of the conditional treatment effects.

For identification of the conditional average treatment effect, further assumptions are needed and discussed, e.g. in Imbens (2004) or Abadie and Imbens (2006). The most important assumption, the Conditional Independence Assumption (sometimes called unconfoundedness assumption), means that potential outcomes are independent of the group and therefore treatment assignment, conditional on covariates. If the *Conditional Independence Assumption* holds, any potential selection bias vanishes and the observed difference in average outcomes conditional on the observables between the treatment and the control group can be interpreted as the causal, conditional treatment effect.

The second most important assumption is the so called *overlap assumption*, which basically says that all characteristics observed in a treatment group have to be found amongst the individuals in the control group, because otherwise, a comparison of the expected potential outcomes, given those covariates, is not possible. It is generally never guaranteed that this is possible, but a powerful treatment assignment procedure will make it more probable. Formally (Abadie and Imbens, 2006) we have

Assumption 1 (Conditional Independence Assumption and Overlap Condition)**.** For every $t = 0, 1, \ldots, n_d$, for almost every $x \in \mathbb{X}$, where $\mathbb{X}$ denotes the support of X_i and $i = 1, \ldots, N$, the following conditions hold:

$$D_i \text{ is independent of } Y^p_i \text{ conditional on } X_i = x; \qquad \text{(CIA)}$$

$$\eta < \Pr(D_i = t \,|\, X_i = x) \text{ for some } \eta > 0. \qquad \text{(Overlap)}$$

4.2.2 A Mean Squared Error Based Minimization Function

The Mean Squared Error of an estimator $\hat{\tau}$ conditional on X is defined as

$$\text{MSE}(\hat{\tau} \,|\, X) = \mathbb{E}\left[(\hat{\tau} - \tau)^2 \,|\, X\right],$$

where τ is the real-valued parameter to be estimated. The MSE can be decomposed into the variance and bias of the estimator, conditional on X, and thus results in a measure of efficiency for unbiased estimators, given a specific set of data X.

More generally, let $w^d = (w_1^d, \dots, w_{n_d}^d)$ and $w^y = (w_1^y, \dots, w_{n_y}^y)$ be non-negative vectors that weight treatments $t \geq 1$ and outcomes, respectively. Then, for the matrix of weighted estimators $\operatorname{diag}(\sqrt{w^d})(\hat{T}) \operatorname{diag}(\sqrt{w^y})$, we define the conditional weighted MSE component-wise as

$$\operatorname{MSE}(\hat{T}, w^d, w^y \mid X) = \mathbb{E}\left[\left\| \operatorname{diag}(\sqrt{w^d})(\hat{T} - T) \operatorname{diag}(\sqrt{w^y}) \right\|_F^2 \mid X \right], \qquad (4.1)$$

where $\| \cdot \|_F$ denotes the Frobenius norm. If weights shall not be considered, w^d and w^y will be vectors with entries 1 only, so $\operatorname{diag}(w^y)$ and $\operatorname{diag}(w^d)$ will be the $n_y \times n_y$ and $n_d \times n_d$ identity matrix, respectively. We assume w^d and w^y be independent of T.

The expectation of the squared Frobenius norm of the matrix $\hat{T} - T$ with its corresponding weights is—because of linearity—simply the trace of the expected 'squared' weighted error matrix:

$$\begin{aligned} &\operatorname{MSE}(\hat{T}, w^d, w^y \mid X) \\ &\quad = \mathbb{E}\left[\operatorname{tr}\left(\operatorname{diag}(\sqrt{w^y})(\hat{T} - T)^\top \operatorname{diag}(w^d)(\hat{T} - T) \operatorname{diag}(\sqrt{w^y}) \right) \mid X \right] \\ &\quad = \operatorname{tr}\left(\operatorname{diag}(\sqrt{w^y}) \mathbb{E}\left[(\hat{T} - T)^\top \operatorname{diag}(w^d)(\hat{T} - T) \mid X \right] \operatorname{diag}(\sqrt{w^y}) \right). \end{aligned}$$

Objective Function The objective is to minimize the generalized MSE (4.1). Hence, we seek an estimator $\hat{T}$ minimizing this function for the given weights w^d and w^y:

$$S_T(\hat{T}) = \operatorname{MSE}(\hat{T}, w^d, w^y \mid X).$$

As the quantity of interest, the conditional average treatment effect, is a function of the covariates, it is natural to start from this point: Suppose the estimator of the treatment effects $\hat{T}$ is a function of X, so $\hat{T} = m(X)$. As the weights do not depend on $\hat{T}$, $S_T(\hat{T})$ is given by the trace of $\mathbb{E}\left[(m(X) - T)^\top (m(X) - T) \mid X \right]$, which can be written as

$$(m(X) - \mathbb{E}[T \mid X])^\top (m(X) - \mathbb{E}[T \mid X]) + \mathbb{E}\left[(\mathbb{E}[T \mid X] - T)^\top (\mathbb{E}[T \mid X] - T) \mid X \right].$$

Since the last summand does not involve $m(X)$, $S_T(\hat{T})$ is minimized by setting $m(X) = \mathbb{E}(T \mid X)$. With that,

$$\mathbb{E}[T \mid X] \in \underset{\hat{T}}{\operatorname{argmin}}\, S_T(\hat{T}).$$

Considering the t-th row of the matrix $\mathbb{E}[T \mid X]$ and using the definition of the Conditional Average Treatment Effect (Definition 5) yields

$$\mathbb{E}[\tau_t \mid X] = \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}\left[Y_{i,t}^p - Y_{i,0}^p \mid X_i\right] \mid X\right] = \frac{1}{N}\sum_{i=1}^{N}\mathbb{E}\left[Y_{i,t}^p - Y_{i,0}^p \mid X_i\right].$$

This, however, leaves us with the challenge of estimating $\mathbb{E}\left[Y_{i,t}^p \mid X_i\right]$ for all treatment groups $t = 0, 1, \ldots, n_d$.

4.2.3 A Linear Model for Potential Outcomes

We choose a linear model for the relationship between covariates and potential outcomes.

Assumption 2 (Potential outcomes are linear functions of covariates).

$$Y_{i,t}^{p,k} = X_i^\top \beta_t^{p,k} + \varepsilon_{i,t}^{p,k}, \tag{4.2}$$

for $i = 1, \ldots, N$, $k = 1, \ldots, n_y$ and $t = 0, 1, \ldots, n_d$ with
- $Y_{i,t}^{p,k}$ a random number taking values in $\mathbb{R}$,
- X_i a random vector of length m with values in $\mathbb{R}$,
- $\beta_t^{p,k}$ the vector of deterministic parameters of dimension m and
- $\varepsilon_{i,t}^{p,k}$ a real-valued random number.

We assume (Y_i^p, X_i) independent and identically distributed for all $i = 1, \ldots, N$. For the error terms, we assume $\varepsilon_{i,t}^{p,k} \mid X_i \sim \mathcal{N}(0, \sigma_{t,k}^2)$ for all $i = 1, \ldots, N$ and all $k = 1, \ldots, n_y$, $t = 0, 1, \ldots, n_d$. Moreover, we assume independence between $\varepsilon_{i,t}^{p,k}$ and $\varepsilon_{i,0}^{p,k}$ for $i = 1, \ldots, N, k = 1, \ldots, n_y$ and $t = 1, \ldots, n_d$. The variances are expressed in relation to a 'base' variance: $\sigma_{t,k}^2 = s_{t,k}\sigma_{0,k}^2$ for all $t = 1, \ldots, n_d$, $k = 1, \ldots, n_y$ with $s_{t,k} > 0$ and $\sigma_{0,k}^2 = s_{0,k}\sigma_0^2$ with $s_{0,k} > 0$ for all $k = 1, \ldots, n_y$ and for a $\sigma_0^2 > 0$.

Let the submatrix X_t of X contain the covariate vectors of all individuals in treatment group t, that is $X_t := (X_{i_1}, X_{i_2}, \ldots, X_{i_{n_t}})$ for $\{i_1, i_2, \ldots, i_{n_t}\} = \{i : D_i = t\}$. Then, the objective function can be expressed conveniently in terms of covariates, treatment group assignment and possibly weights, as the following theorem shows.

Theorem 1. *Under Assumption 2, the minimization criterion (4.1) equals*

$$\frac{1}{N^2}\sum_i X_i^\top \left[\| \tilde{w}^y \|_1 \| w^d \|_1 \left(X_0 X_0{}^\top\right)^{-1} + \sum_k \left\{ \tilde{w}_k^y \left(\sum_{t>0} \tilde{w}_{t,k}^d \left(X_t X_t{}^\top\right)^{-1} \right) \right\} \right] \sum_i X_i,$$

where $\| \cdot \|_1$ *is the* l_1 *norm of a vector,* $\tilde{w}_k^y = w_k^y s_{0,k}$ *and* $\tilde{w}_t^d = w_{t,k}^d s_{t,k}$ *for* $k = 1, \ldots, n_y$ *and* $t = 1, \ldots, n_y$.

The proof of Theorem 1 is in Appendix C.1.

Corollary 3. *Under the condition of Theorem 1 and assuming the same variance for all outcomes and treatment groups, including the control group (i.e.* $s_{t,k} = 1$ *for all* $t = 0, 1, \ldots, n_d$ *and all* $k = 1, \ldots, n_y$*) and neglecting any weights (i.e. assuming* $w_t^d = w_k^y = (1, \ldots, 1)$*), minimizing (4.1) through choice of* D *is equivalent to minimizing*

$$\frac{1}{N}\sum_i X_i^\top \left[n_d \left(X_0 X_0{}^\top\right)^{-1} + \sum_{t>0} \left(X_t X_t{}^\top\right)^{-1} \right] \frac{1}{N}\sum_i X_i. \tag{4.3}$$

Contrary to the result by Kasy (2016), our approach is more applicable as the researcher is relieved from guessing any (absolute) values. This is because we do not assume a prior distribution for $\beta_t^{p,k}$, as such there is no need to specify its parameters: a mean or—more difficult—a covariance matrix for this parameter vector for every combination of k and t in case of an assumed normal prior distribution. Furthermore, there is no need of specifying the R^2 for the model of each potential outcome in order to express the model's variance. Instead, one can simply specify the relative scaling factors.

4.3 Characteristics of the min MSE Treatment Assignment Method

In what follows, we study the theoretical characteristics of the min MSE treatment assignment procedure.

Proposition 4. *Under the conditions of Corollary 3, (4.3) is constant under a transformation of the vector* $(X_{j,1}, \ldots, X_{j,N}) \mapsto (cX_{j,1}, \ldots, cX_{j,N})$ *for any* $c \neq 0$ *and for any* $j = 1, \ldots, m$.

The proof of Proposition 4 is in Appendix C.1.

Proposition 4 states that the min MSE Treatment Assignment Procedure is scale invariant in the sense that the corresponding minimization criterion (4.3) is unaffected by changing the scale of a covariate. This feature is desirable, since it renders rescaling of the data unnecessary, but at the same time leaves the freedem to do so.

Proposition 5. *Assume $X_{k,i}$ is orthogonal to $X_{j,i}$ for $k, j = 1, \ldots, m$, $k \neq j$ with respect to the inner product $< \cdot, \cdot >_2$ of L^2, i.e. $\mathbb{E}[X_{k,i} X_{j,i}] = 0$. Furthermore, assume all covariates have the same mean, i.e.$\mathbb{E}[X_i] = c(1, \ldots, 1)^\top$ for any $c \neq 0$. Then, for $N \to \infty$, a solution to the minimization problem according to Corollary 3 is obtained, if*

$$\sum_j \left[n_d \left(\sum_{\{i:D_i=0\}} X_{j,i}^2 \right)^{-1} + \sum_{t>0} \left(\sum_{\{i:D_i=t\}} X_{j,i}^2 \right)^{-1} \right] \tag{4.4}$$

is minimized.

The proof of Proposition 5 is in Appendix C.1.

Proposition 5 states that—provided the covariates are orthogonal and have equal mean—the MSE for $n_d = 1$ is decreased, if the sum of squared observations of the covariates are increased in both groups and among all covariates. That is, the absolute deviation from 0 is 'balanced' for all covariates across groups. In the simple case of one treatment and one control group, equally sized, with one covariate considered for treatment assignment, this is equivalent to balancing the second moment of the distribution of the covariate of interest. This makes the min MSE procedure a unique method in the sense that 'balance' incorporates not just the mean, but a higher moment of the distribution of covariates. It is exactly this property that makes the groups comparable in the sense that the different subgroups–if any–are to be found in all experimental groups.

Proposition 6. *Under the conditions of Corollary 3, now allowing for arbitrary mean values of the covariates and arbitrary relationships between covariates, the diagonal elements of the matrix $\left(X_0 {X_0}^\top\right)^{-1} + \left(X_t {X_t}^\top\right)^{-1}$ in (4.3) are given by*

$$\mathrm{Var}(\hat{\beta}_{t,j} - \hat{\beta}_{0,j} \,|\, X) \propto 1/[(1 - {R_j^t}^2) \sum_{i=1}^{n_t} (X_{j,i}^t - \bar{X}_j^t)^2] + 1/[(1 - {R_j^0}^2) \sum_{i=1}^{n_0} (X_{j,i}^0 - \bar{X}_j^0)^2] \tag{4.5}$$

for every $t \geq 1$ and every $j = 1, \ldots, m$, where $\propto$ denotes equality up to a multiplicative constant and the value of the j-th covariate of individual i in treatment group t is denoted by $X_{j,i}^t$ and $\bar{X}$ denotes the mean. ${R_j^t}^2$ is the coefficient of determination of a regression between the variable X_j^t as response and all X_p^t for $p = 1, \ldots, m$, $p \neq j$ as explanatory variables; for all $t = 0, 1, \ldots, n_d$. The number of individuals in treatment group t and the control group is denoted by n_t and n_0 respectively.

The proof of Proposition 6 is in Appendix C.1.

The sum on the right-hand side of (4.5) is decreased for every $t \geq 1$ and every covariate $j = 1, \ldots, m$, if the linear dependencies between covariates within groups are decreased. This introduces a certain orthogonality criterion, with $R_j^t = 0$ when all other covariates in the group are (partially) uncorrelated with the covariate X_j. When covariates are perfectly collinear, we would have $R_j^t = 1$ (however, in this case, the covariance matrix of the estimator of the parameter vector does not even exist, as $X_t X_t^\top$ is not invertible). Thus by having this criterion in the objective formula, we reward a grouping that avoids multicollinearity and punish a high level of similarity amongst the combination of covariates in a group. Note, however, that this grouping might not minimize off-diagonal entries of the sum of the covariance matrices.

Would the off-diagonal entries be minimized in the same way than the diagonal entries according to (4.5), then, when balancing the covariates age and household income for example, a family's twin children living in the same household should not be placed in the same group. Consider the extreme case with two twin pairs and two groups: If twin pairs are in the same group, within each group, both variables are perfectly predictable by the other. The more combinations of covariates can be found in a group, the smaller ${R_j^t}^2$. However, a reduction in ${R_j^1}^2$ might lead to an increase in ${R_j^0}^2$. Therefore, (4.5) is decreased for every $t \geq 1$ and every covariate $j = 1, \ldots, m$, if for one group, more combinations of covariates can be achieved without limiting the combinations of covariates to the same amount expressed in ${R_j^1}^2$ and ${R_j^0}^2$.

The second part influencing (4.5) is the within-group variation of variable j around its mean. The higher the variation, the lower the variance. Again, an increase in overall variance can only be achieved if an increase in a variable's variation in one group does not lower that variable's variation in the other to the same extent or more.

Especially the first part is interesting, as it shows a characteristic that is also inherent in matching: Two very similar subjects should not be allocated to the same group. This characteristic additionally distinguishes the min MSE procedure from other rerandomization methods, as it considers the complete composition of covariate values in a group instead of considering all covariates independently.

4.4 Comparison of the min MSE Treatment Assignment Method with Alternatives

4.4.1 Pair-Wise Matching

Consider treatment assignment for a treatment and a control group, where for every individual i, one covariate x_i is observed and the treatment should be assigned such that this covariate is balanced across the treatment and control groups.

Theorem 2. *Pair-wise matching before treatment assignment is a 'max min sum of variances' approach.*

The proof of Theorem 2 is in Appendix C.1.

In essence, this theorem shows that also matching aims at balancing a higher moment of the covariate distribution than the mean, as does the min MSE approach. Basically, it ensures that the most similar observations are assigned to different groups.

4.4.2 An (Alternative) Linear Model for Potential Outcomes

The criteria considered by Greevy et al. (2004) to compare the efficiency of treatment assignment could also be used as an optimization criterion for treatment assignment.

For every treatment t, a linear, additive model is specified as follows:

$$Y_t = \begin{bmatrix} Z_t & X_t \end{bmatrix} \begin{bmatrix} \tau_t \\ \beta^t \end{bmatrix} + \varepsilon,$$

where the subscript t of Y_t, Z_t and X_t indicates that row entries are from individuals of $\{i : D_i = t \vee D_i = 0\}$. Y_t contains the potential outcomes for the control group or treatment group t, depending on Z_t, which is the treatment status, with $Z_{i,t} = \mathbb{1}_{\{D_i=t\}}$ for those in treatment group t and $Z_{i,t} = -\mathbb{1}_{\{D_i=0\}}$ for the control group. X_t contains the covariate vectors $X_i^\top$ of individuals in treatment group t and in the control group. In this model, $2\tau_t$ is the estimate for the conditional average treatment effect.

Under the Gauss-Markov assumptions (additive errors, that are uncorrelated conditional on X_t with constant variance σ^2), the MSE of the estimated treatment effect is proportional to $1/Z_t^\top P_t Z_t$ with $P_t = I - X_t \left(X_t^\top X_t\right)^{-1} X_t^\top$ and is minimized for $X_t^\top Z_t = 0$ (Greevy et al., 2004). Thus, with the assumption of constant variances across treatments and without imposing weights, an alternative objective function for minimization would be

$$S^*_T(\hat{T}) \propto \sum_t 1/Z_t^\top P_t Z_t. \tag{4.6}$$

In this model, the treatment effect is assumed to be constant across individuals, so potential outcomes of the control group and the treatment group of interest are assumed to differ only by a constant. Contrarily to a simple difference in means estimator for the average treatment effect, here covariates are controlled for, which induces a criterion of balance for covariates. As this criterion is minimized by $X_t^\top Z_t = 0$, it is enough to have equal mean values of a covariate to minimize this criterion (given equal group sizes), independent of the distribution in the respective groups.

Thus, comparing this result with the results derived in Section 4.3 shows that if there is reason to assume that any of the treatment effects might differ across individuals and be a function of the covariates, it is necessary to focus on more distributional characteristics of the covariates than their means.

4.4.3 Morgan and Rubin (2012)

The approach by Morgan and Rubin (2012) considers the Mahalanobis distance between the vector of covariate means of the control group and the vector of covariate means of the treatment group. When group averages are equal, the distance is minimal. For the derived statistic, a threshold to stop re-randomization is derived. This approach is closely related to the omnibus test for multivariate covariate balance by Hansen and Bowers (2008); in fact, for the case without strata or matching pairs, the statistic is the same. It is related to Hotelling's T-test, but it treats treatment assignment instead of covariates as random and thus follows a χ^2-distribution, rather than an F-distribution under the null of no difference between groups (Hansen and Bowers, 2008).

With respect to the notion of balance, the statistic shares its properties with the objective function (4.6) of the just discussed alternative linear model as considered in Greevy et al. (2004). Balance in these approaches equals balancing group means,[2] thus balance is limited to the first moment of the distribution of covariates.

2. For (4.6) this holds when the treatment and control groups are of equal size.

4.5 Simulation Study

In order to investigate the performance of the treatment assignment procedure described in Section 4.2, we perform a simulation study similar to the one by Bruhn and McKenzie (2009), henceforth cited as BK09, to compare our new mechanism to established ones.

BK09 compare five treatment assignment methods (purely random assignment, pairwise greedy matching, stratification and two rerandomization schemes) in terms of 'balance' of relevant observable and "unobservable" variables when creating one treatment and one control group. To rule out the possibility that results depend on the characteristics of a specific dataset or sample size, they consider several data sets.

We extend this study by adding the scenario of multiple treatment arms and a scenario where attrition occurs randomly. In terms of treatment assignment mechanisms, we also include an 'optimal matching' approach (Lu et al. (2011) as introduced by Greevy et al. (2004)) and our new min MSE procedure, as introduced in Section 4.2 of this paper.

4.5.1 Study Design

Data

We use the same data as BK09 for reasons of comparability. It consists of four panel datasets, with different data from different contexts.

The first dataset contains data on microenterprises in Sri Lanka and is from an actual randomized experiment by De Mel, McKenzie, and Woodruff (2008). The outcome variable of interest is firms' profits, and data on firm and owner characteristics at the time of the baseline study is available. It is either used for treatment assignment or treated as "unobservable" and studied after treatment assignment to assess the effect of the different methods on "unobservables".

The second dataset consists of a subsample of the Mexican employment survey (ENE), where we used the same subsamples as BK09. In this dataset, the outcome of interest is the income of household heads that were employed and between age of 20 and 65 when the baseline survey was conducted in 2002. In addition to this, the dataset includes additional characteristics on the household and its head, which again are used either for treatment assignment or as "unobservables".

The third dataset is comprised of subsamples with two waves (IFLS 2 and IFLS 3) from the Indonesian Family Live Survey (IFLS):[3] The year 1997 (IFLS 2) is used as the baseline and the data from 2000 (IFLS 3) is treated as the follow-up. We only use data on household expenditure from this survey.[4]

The fourth dataset is from the Learning and Educational Achievement project in Pakistan, which is also used by Andrabi, Das, and Khwaja (2015). It contains child and household data, and the outcome variables of interest are math test scores and z-scores of children between the ages of 8 to 12 at the baseline.

It is noteworthy that the subsamples of 30 and 100 observations sometimes differed considerably: The share of the variation in the follow-up variable explained by the group used for treatment assignment in the dataset on firms' profit for the small subsample is around 6 percent, whereas it amounts to 18 percent for the subsample of 100 observations. A larger difference is observed in the dataset from Mexico[5] and in the data on height z-scores, in the smaller subsample, an even higher share of variation in z-scores could be explained than in the larger sample (64% and 51%, respectively). For the remaining datasets, however, no meaningful differences are found.

Nevertheless, this observation gives rise to a method for drawing comparative samples of a 'universe'. The treatment assignment procedure derived in Section 4.2 can also be used in this setting.[6]

Treatment Assignment Mechanisms

A number of treatment assignment mechanisms is common (BK09).

In this study, we additionally consider the minimal MSE procedure as introduced in this paper, and the matching method called 'optimal matching' as implemented in the *R* package *nbpMatching* (Lu et al., 2011), going back to the work of Greevy et al. (2004). We give a short overview over the assignment mechanisms applied in this study. In Section 4.7, we discuss the weaknesses and strengths of the different treatment assignment mechanisms in comparison with our min MSE method.

3. See http://www.rand.org/labor/FLS/IFLS.html.

4. BK09 also use data on the schooling of children from this dataset. Given that they do not report results for this dataset in all graphs and tables, we limited ourselves to the inclusion of household expenditure.

5. Roughly 7 and 32 percent (for the subsamples of 30 and 100 observations, respectively) of the variation in household expenditure is explained by "observable" variables.

6. A Stata software package for this purpose can be obtained from the author.

Pure randomization refers to the realization of a single random draw, performed by using anything from a coin to a random number generator on a computer. Stratification, also known as blocking, is attributed to Fisher (1935). The idea is to build subgroups according to observable characteristics (covariates) and randomize within those subgroups.

Pairwise matching is in a certain sense the limit case of stratification, with only two units per strata, which are then randomly assigned to the treatment or the control group. BK09 apply a 'greedy algorithm' laid out in Imai, King, and Nall (2009), an approach popular in the literature on matching observational data at least since Rubin (1973). This implementation of the greedy algorithm computes pairwise Mahalanobis distances[7] between two units for the whole sample and pairs the two with the smallest distance; those then are taken out of the sample of units to be matched and the procedure is repeated. Overall distance is not necessarily minimized by this approach, because it is not 'forward-looking' (Rosenbaum, 1989; Greevy et al., 2004). The approach called 'optimal matching', as introduced by Rosenbaum (1989) for observational and Greevy et al. (2004) for experimental studies aims at achieving this goal.[8] We use both implementations in our study.

Finally, different rerandomization methods have evolved and are widely used (BK09). A rerandomization approach is basically any method that performs a somehow random treatment assignment, and repeats randomization until a certain condition is reached. This condition might either be a certain number of iterations or a statistical threshold or even subjective judgment. In the first case, the "best" assignment is chosen; in the second, usually the first to reach the statistical threshold is kept. In this sense, the "best" assignment can be determined in various ways. A representative of the first group is, for example, the min max t-stat method, in which 1000 random assignments are made, and the one chosen is the one in which the maximal t-statistic on any variable to consider is the smallest. A variant of the second group is the 'big stick' method in BK09, in which a new treatment assignment is drawn if any difference in means between treatment and control group is significantly different from zero. A more sophisticated approach is the one by Morgan and Rubin (2012, 2015), where the Mahalanobis distance between the group means of the covariates is considered and a theoretical threshold

7. The use of Mahalanobis distance in matching for observational data has been discussed e.g. in Cochran and Rubin (1973).

8. Forming pairs is referred to as non-bipartite matching in the optimization but also the matching literature. This procedure (non-bipartite matching) is different than finding matches in already existing groups (bipartite matching) and is considerably more difficult (Lu et al., 2011). Therefore, not all results for matching, and more importantly, software implementations, can be applied in this setting.

is derived. Since the ad-hoc approaches lack a theoretical foundation and as for the last method, a software implementation is missing, we do not consider these approaches in this study.

Min MSE Method The Min MSE Method can be considered as a rerandomization method, in which a certain number of iterations is drawn. Unlike the rerandomization approaches considered in BK09, our implementation of the min MSE approach improves on previous draws, as we use the stochastic simulated annealing optimization algorithm (Kirkpatrick, Gelatt, and Vecchi, 1983): Given a treatment assignment, a new one is obtained by randomly exchanging the treatment status for a certain number of units. The new assignment is then evaluated according to the formula derived in Section 4.2 and either kept or withdrawn. Thus, the min MSE method maximizes the balance in a more efficient way with respect to time than the other discussed approaches of rerandomization.

Furthermore, apart from the approach by Morgan and Rubin (2012), we are not aware of a rerandomization criterion that has a theoretical foundation and is not an ad-hoc measure. For their approach—to the best of our knowledge—there is no extension to multiple treatment groups[9] and the criterion is only based on the mean differences of the treatment groups. Additionally, we were unable to find any software implementation of their approach.

Variables for Balancing

We used the same variables for treatment assignment as BK09. They include the baseline outcome of an outcome of interest, and add six other variables that may affect the outcome of interest, with the exception of stratification, where only subsamples are used. This means, however, that the results of the study regarding stratification can only be conditionally compared to the other results, since stratification is tested with a lower number of variables and thus has a higher likelihood to achieve balance on those, and in particular, the baseline outcome variable. For reasons of comparability, we stick to this approach despite its shortcomings. For the exact reasoning for the choice of variables to be balanced, we refer to Bruhn and McKenzie (2009).

We use the same variables with the newly added treatment assignment mechanisms as we did for greedy matching: the baseline outcome and six additional variables.

9. Although they name a possible way of extending their criteria in this sense.

Attrition

Researchers might be concerned about the consequences of attrition when a sophisticated method of treatment assignment has been applied. In case obtaining the outcome from a unit fails when having used a matching approach, it is common to also exclude its pair from the analysis (Imai, King, and Nall, 2009). While mostly perceived as a disadvantage for the diminished sample size—especially when performing cluster randomization—Imai, King, and Nall (2009) consider this practice an advantage, as they argue the remaining sample is still balanced.

We investigate this claim by randomly removing $1, 3, 5$ and 7 units after the treatment groups have been assigned with a sample size of $N = 30$. While we exclude the pair of a randomly removed unit from the treatment groups assigned by the matching approaches, for the other treatment assignment methods, such a possibility is missing and we leave the groups unaltered after simulated attrition. Subsequently, we investigate how attrition has affected balance.

4.5.2 Comparing Treatment Assignment Mechanisms

Pre-Treatment Balance

We investigate balance using the measures of pre-treatment balance on baseline variables as BK09 for the cases of one treatment arm. In the main text, we will report suitably aggregated results over all variables used for treatment assignment, as our interest lies rather in overall performance than in performance on an arbitrarily selected variable. Results for the latter case are printed in Appendix C.2

For reasons of comparability, we also assess balance in follow-up outcomes. However, we think that balance on follow-up outcomes is rather important when assessing the general value of covariate based treatment assignment mechanisms in panel studies, which is beyond the scope of this study. We therefore report those results in Appendix C.2.

For the cases of multiple treatment arms, we extend the measures used by BK09 in a suitable way.

Balance in a single variable, one treatment arm To assess balance in a single variable for the case of one treatment and one control group, BK09 compare the difference in means for one draw, expressed in the variable's standard deviation. Of all draws, they then graphically compare the distribution of the differences and report the average, and the 95% quantile of the distribution of (absolute) differences in the group means. Additionally, they perform a t-test to assess whether or not estimates for differences are "significantly" different from zero, and report the share of draws in which this was the case. We assess balance using these measures and report results in Appendix C.2 for comparison reasons.

Balance in a group of variables, one treatment arm First, standardized differences in means are calculated for every single variable of the group of variables for one draw. Then, for the average (absolute) difference and for the share of estimates significantly different from zero as calculated for a single variable, overall averages are built. For assessing the 95% quantile of differences, first the 95% quantile of every single variable in the group is determined. Then, the maximum among the variables in the group is reported as the 95% quantile. To detect extreme imbalances, we also compute the maximum difference of group means of a group of variables and evaluate the distribution of these maximal differences across 10,000 iterations graphically.

Balance in a group of variables, several treatment arms When aggregating the balance of several treatment arms, taking the average difference of the means between the several treatment groups and the control groups before taking the average over all variables of interest in all performed draws is one option. Another possibility is to take the overall average over the largest difference in the means between the treatment and control groups in one variable and in one draw. We think both are relevant and perform both.

4.6 Results

All results are based on 10,000 simulations, unless otherwise stated. The sample size, which was used for the tables and graphs, is indicated in the respective caption. For sample sizes, where results are not reported in the text, we provide the respective tables and graphs in Appendix C.2.

4.6.1 Scenario 1: One Treatment Group

We first present the results for the scenario considered in BK09: Units have to be assigned to either one treatment or the control group. In the main text, we focus on aggregate measures over all variables considered for treatment assignment, since no single variable has received a higher focus or a higher weight. In Appendix C.2, we present results for one single variable, as BK09 report these results.

Table 4.1 shows the average differences between the group means, average absolute differences between group means and the 95% quantile of the differences in group means among the 10,000 iterations performed. A lower measure indicates a better balance in group means. The main message of Table 4.1(a) is that, on average, all differences equalize and we observe balance (note that results are reported in 1000 standard deviations and thus are zero to the third or fourth digit in the upper panel of Table 4.1(a)).

A measure that is informative with respect to imbalance in a single random draw is the absolute difference in group means. The averaged absolute difference is reported in Table 4.1(b) for every treatment assignment mechanism and every dataset. On average and for the variables of every single dataset, the min MSE method performs better than all competing methods. Compared to the second best method, 'optimal matching', it reduces the average absolute difference in the balance of group means by nearly 30%.

Finally, Table 4.1(c) shows the average over all 95% quantiles of the absolute differences in a single variable considered for treatment assignment. Averaging all of the datasets, the min MSE method again performs superior to competing methods.

In summary, the min MSE method not only performs better than to competing methods on average, but also reduces extreme differences in the mean values between the treatment and control groups the most as compared to a single random draw. This last finding is supported by the results in Figure 4.1, which shows the distribution of the largest differences between the group means of any variable for a single draw. In all graphs, we see that the mass of differences close to zero is largest for the min MSE approach, which also always yields a favourable mass in the tails as compared to competing methods.

4.6.2 Scenario 2: Multiple Treatment Arms

The second scenario we consider is an experiment, in which multiple (variants of) interventions are tested. Units shall be assigned to the control or one of the treatment groups while keeping all groups comparable.

For this scenario, we were unable to find a software implementation of a competing method, so we compare the min MSE procedure to a single random draw.

The findings are graphically presented in Figure 4.2. We first compute the maximum and the mean difference between the treatment group means and the control group mean of one variable for a single draw. We aggregate the measure over the variables of one draw and over all iterations by averaging over the mean or maximal difference. The first aggregate measure is shown by the dashed lines, the latter aggregate measure by the solid lines. Both lines, the solid and the dashed one, start at the same point, as for only one group, the maximum and the

TABLE 4.1: Comparison of Treatment Assignment Methods Regarding Balance in a Group of Baseline Variables (N=30)

(a) Average difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Indonesia	0.025	0.444	1.134	-1.010	1.729	3.101
Pakistan (height scores)	-0.875	0.785	-0.872	1.035	-0.489	-0.062
Pakistan (test scores)	1.944	-1.579	0.949	-0.653	-1.903	-0.138
Mexico	2.191	1.996	-0.477	0.481	-0.021	-0.310
Sri Lanka	0.624	2.111	-0.748	-0.328	-1.594	-1.667

(b) Average abs. difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Indonesia	295.9	198.1	174.0	126.2	233.5	243.8
Pakistan (height scores)	291.0	180.1	172.3	134.6	235.4	230.5
Pakistan (test scores)	293.1	287.9	180.3	103.0	244.7	257.3
Mexico	299.1	223.8	174.1	147.8	258.5	262.1
Sri Lanka	292.4	253.9	170.9	93.4	248.8	267.0
Total	294.3	228.8	174.3	121.0	244.2	252.1

(c) 95% quantile of the difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Indonesia	788.3	655.5	579.9	643.4	742.6	702.3
Pakistan (height scores)	802.8	655.5	535.2	393.8	783.3	722.9
Pakistan (test scores)	715.2	905.4	628.4	533.8	729.3	715.2
Mexico	701.3	677.1	445.0	348.2	727.7	702.3
Sri Lanka	802.8	863.6	535.2	311.8	792.3	744.7
Total	762.1	751.4	544.7	446.2	755.1	717.5

Note: Statistics based on 10,000 iterations. Details on the study and the computation of each measures are explained in Section 4.5.2. For every dataset, several variables were considered for treatment assignment. The results in this table report aggregate measures of differences in treatment group means for the group of considered variables. Differences are weighted by standard deviation. Lower values indicate better balance with respect to equality of group means.

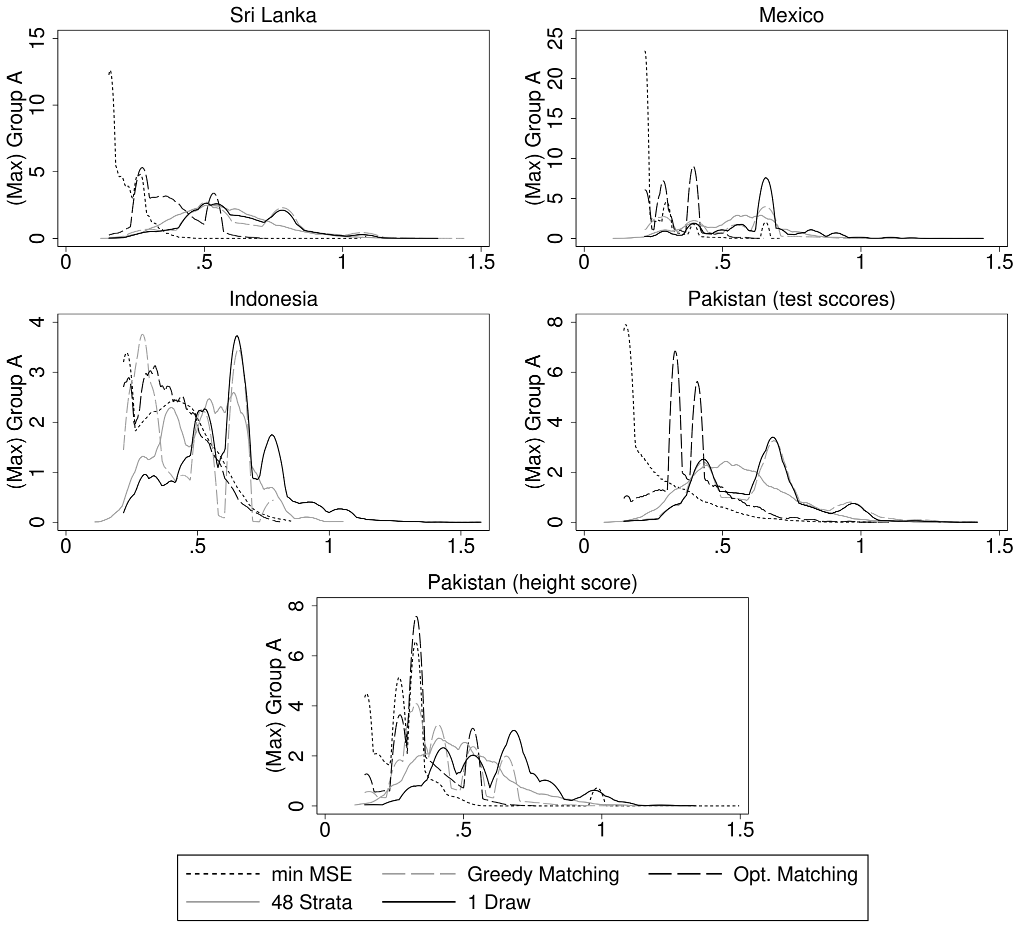

Note: Distributions of the (maximal) differences in treatment group means among the group of variables to consider for treatment assignment are based on 10,000 treatment assignments. Differences in group means are expressed in standard deviations. A high mass around a difference of 0 indicates a good balance with respect to equality of group means.

FIGURE 4.1: Distributions of the Maximal Differences in Group Means (N=30)

mean difference in group means is the same, as there is only one group difference to consider.

When applying the min MSE procedure, the maximum difference—typically the one a researcher is worried about—is always increasing at a lower rate with an additional treatment group to assign than when drawing completely random. For 9 treatment groups, which means 10 groups of 10 units, the average maximal difference (in SD) is—for all datasets—between .4 and .6, whereas when randomly drawing, the average maximal group difference is mostly around .75 or .8 SD. In one case (household expenditure in Indonesia), this average *maximum* difference for 6 treatments (thus 7 groups) when using the min MSE procedure was as high as the average *mean* difference when relying on a single random draw.

In all datasets, the min MSE procedure was able to lower the average maximum difference across group means compared to drawing randomly by between .1 SD (height z-score in Pakistan and labor income in Mexico) and up to .3-.4 SD (math test score in Pakistan and household expenditure in Indonesia).

It is also worth noting that with the min MSE procedure, we can assign between 2 to 5 more treatments compared to randomly drawing with the same maximum difference in group means to be expected.

4.6.3 Scenario 3: Attrition

The third considered scenario corresponds to the first scenario, where units have to be assigned to either the treatment or the control group. After treatment assignment, however, some units fail to provide the outcome of interest, i.e. the study suffers from attrition. We randomly remove $1, 3, 5$ and 7 units after the treatment groups have been assigned with a sample size of $N = 30$. We "correct" for attrition in case matching approaches were used for treatment assignment, see Section 4.5.1 for details.

Table 4.2 reports the results of the attrition scenario. Table 4.2(a) shows the absolute difference between group means, averaged over variables considered for treatment allocation, datasets and iterations. We report the average absolute difference instead of the average difference, since average differences are close to zero for all mechanisms, and imbalances obtained for the different attrition levels could average out.

The first observation is that attrition always strictly worsens balance for all treatment assignment mechanisms. Except for the case when 7 units (roughly 25% of the sample) drop out, balance achieved with the min MSE mechanism is best for all levels of attrition, as indicated by the lowest average absolute difference. Only the 'optimal matching' algorithm achieves a better balance in that

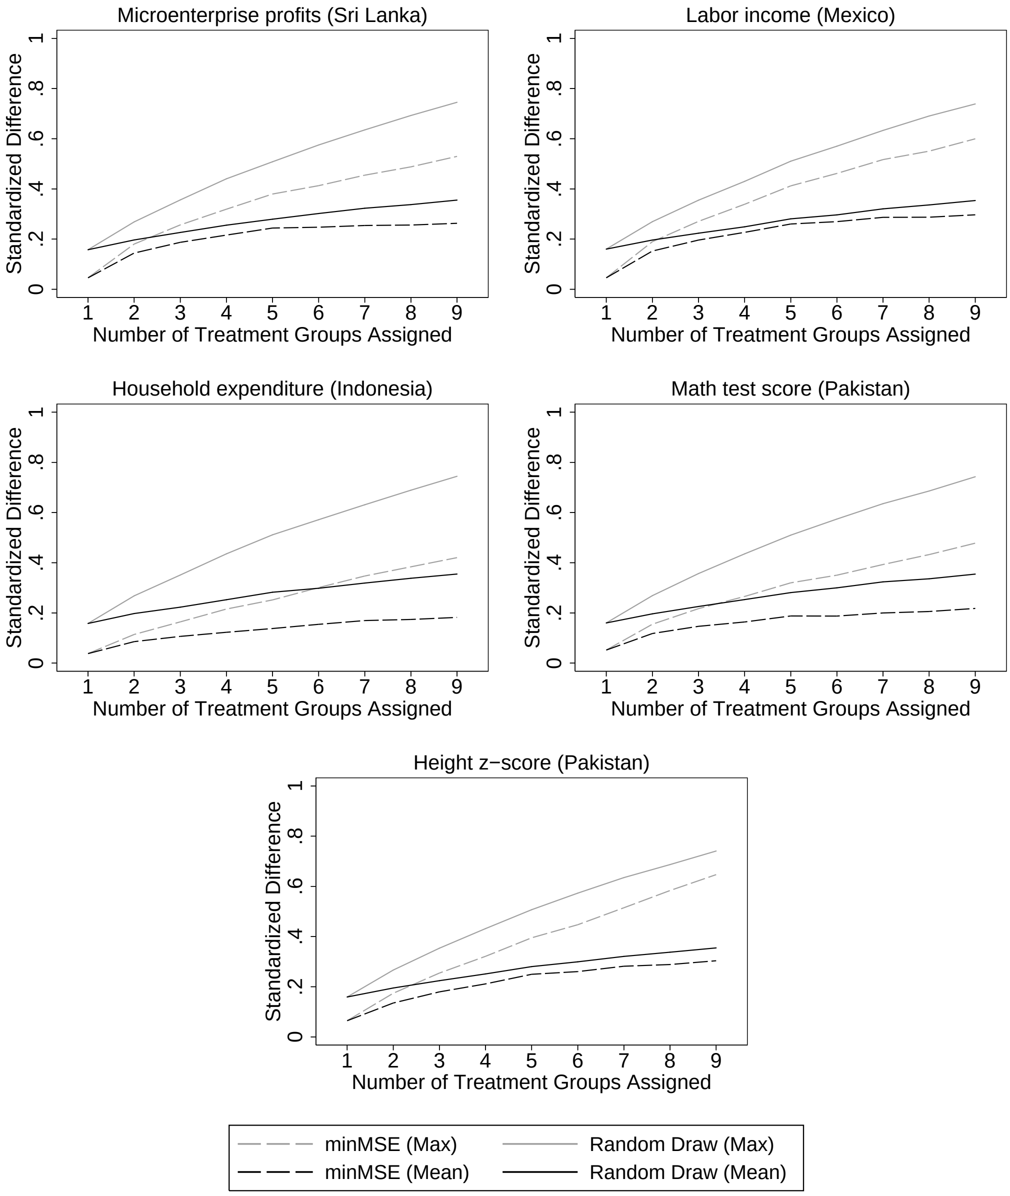

Note: Multiple Treatment Assignment. Evolution of the (mean and maximal) difference in group means between the treatment groups and the control group among the group of variables to consider for treatment assignment. X-axis: Number of treatment groups to assign. For each dataset, the difference in group means in all variables used for treatment assignment are computed. The line labeled 'max' shows the average over all iterations and variables of the maximum of these differences amongst the treatment groups. The line labeled 'mean' shows the differences amongst the group differences by building the average. Distributions are based on 10,000 treatment assignments. Differences in group means are expressed in standard deviations. Lower mean and maximal differences in group means between the treatment groups and the control group indicate a better balance with respect to equality of group means.

FIGURE 4.2: Multiple Treatment Assignment: Evolution of the Differences Between Treatment Groups and Control Group in the Group of the Baseline Variables for an Increasing Number of Treatments to Assign (N = 100)

case, and the difference between both mechanisms in that case is only marginal, compared to the balance of the other mechanisms. However, the sample assigned with the min MSE mechanism then still consists of 23 units, whereas for the matching approaches, it diminishes to 16 (see Section 4.5.1). On average, however, it is the min MSE mechanism, that performs best when attrition happens with respect to the absolute difference between group means. Compared to the min MSE mechanism and 'optimal matching', stratification and greedy matching cannot considerably improve balance compared to a single random draw.

The second panel of Table 4.2 shows the worst case scenario across 10,000 iterations. For every method, every level of attrition and every variable in every dataset, we compute the 95% quantile of the absolute difference in group means of all iterations. This measure is then averaged across all variables and datasets and reported for every method and level of attrition. When assessing balance according to this measure, again the balance strictly worsens when attrition occurs. On average, it is again the min MSE approach yielding the most favorable results, and with exception of the 'optimal matching' approach, other mechanisms only provide a limited improvement on a single random draw.

In Appendix C.2, we report results for individual datasets and when assessing balance using t-tests.[10]

Summing up, we cannot confirm the claim by Imai, King, and Nall (2009) that the practice of removing the matched pair of a unit dropping out an experiment is actually beneficial. Only when using the 'optimal matching' approach, is balance in cases of attrition comparable to the min MSE approach. On average, however, the min MSE approach outperforms all competing mechanisms in our study in cases of attrition, while mantaining the maximal possible sample size in contrast to the matching approaches.

10. Testing for equality of group means in a single variable of a random draw is often mistaken as a test for successful randomization. BK09 report the share of treatment assignments for every mechanism, where a test for equal group means yields a p-value below .1. In Table C.4(a) in Appendix C.2, we report corresponding results as these tests are frequently conducted. However, we think that it is misleading to assess the probability that a statistic excesses a certain value by pure chance when actually *knowing* that it is pure chance that drives the differences and thus the statistic. Therefore, we forgo presenting these results in the main text. Figure C.6 in Appendix C.2 shows the results of Panel (b) of Table 4.2 for the individual datasets.

TABLE 4.2: Comparison of Treatment Assignment Methods Regarding Balance in a Group of Baseline Variables in Cases of Attrition (N=30)

(a) Average abs. differences in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
No attrition	294.3	228.8	174.3	121.0	244.1	252.1
1 (2) unit removed	300.7	235.9	179.9	139.1	252.0	258.5
3 (6) units removed	312.3	263.8	187.2	163.3	265.6	271.3
5 (10) units removed	328.3	277.0	198.4	186.9	280.5	285.7
7 (14) units removed	352.7	294.3	205.3	211.7	297.8	302.9
Total	317.6	259.9	189.0	164.4	268.0	274.1

(b) 95% quantile of the differences in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
No attrition	705.8	520.7	356.3	231.2	575.4	600.2
1 (2) unit removed	723.6	544.4	372.0	289.1	593.3	614.5
3 (6) units removed	765.6	601.7	381.8	377.1	634.9	650.8
5 (10) units removed	781.1	636.2	395.6	437.5	674.7	687.6
7 (14) units removed	844.8	676.1	408.9	503.6	714.9	730.0
Total	764.2	595.8	382.9	367.7	638.6	656.6

Note: Statistics based on 10,000 iterations. Details on the study and the computation of each measures are explained in Section 4.5.2. For every dataset, several variables were considered for treatment assignment. After treatment assignment, units were randomly removed from the study to simulate attrition. The results in this table report aggregate measures of differences in treatment group means for the group of considered variables. Differences are weighted by standard deviation. Lower values indicate better balance with respect to equality of group means.

4.7 Discussion

4.7.1 Treatment Assignment Mechanism

In what follows, we discuss the strengthes and weaknesses of the min MSE approach as compared to alternative mechanisms.

Pure Randomization Depending on the transparency of the actual implementation, randomization can be considered to be the fairest method for treatment allocation and it certainly is the fastest.

When comparing the means of randomly allocated groups across 20 variables with a conventional t-test and a significance level of 5%, we have to expect that for one variable, the hypothesis of no difference will be rejected. Pure randomization does not yield any device for controlling undesired imbalances that may happen by chance.

Furthermore, it is not guaranteed, especially when the sample size is small, that all characteristics of a variable appear in all experimental groups at all and additionally with the same frequencies; this is a problem when subgroup analysis is desired to study heterogeneous treatment effects.

Stratification The main advantage of stratification is to ensure the possibility of subgroup analysis while ideally increasing the efficiency of the analysis. The idea is to build subgroups according to observable characteristics (covariates) and randomize within those subgroups. This design is probably still considered relatively fair.

One problem of this approach is that continuous variables have to be discretized arbitrarily, and that stratification is only possible for a limited number of variables: Consider a sample of 50 units, where subgroup analysis for age, income and gender is desired. If three categories for age and income are desired, 18 strata have to be created, where at most 3 persons are in one strata. Stratification on another variable is thus not feasible with a comparable sample size. This example points to another drawback: Difficulties arise in implementation if sample size is not divisible by the number of strata. Although solutions to this have been suggested, a simple implementation is no longer possible. Moreover, building the strata requires expertise on both the data and the question under investigation.

The time needed to conduct treatment assignment using stratification depends on the actual implementation, but in simple cases, e.g. with two dichotomous variables, it takes only slightly longer than pure randomization.

Matching Pairwise matching resolves the problems of stratification. Theoretically, pairwise matching may be performed with an unlimited number of variables considered for treatment allocation. Moreover, the possibility to balance continuous variables is an advantage (e.g. Greevy et al., 2004). It is arguably considered to be fair and the design is relatively clear and easy to explain.

Subgroup analysis, however, is not ensured in cases, where balance on a certain variable could not be achieved, which, however, should not be the case in moderate sized samples and a moderate amount of variables to balance.

Implementation of pairwise matching may take considerable time[11] when relying on the 'greedy algorithm' used by BK09. Yet, the software implementation of the 'optimal matching' algorithm in the *R* package *nbpMatching* (Lu et al., 2011) is considerably faster.

However, the biggest disadvantage is probably attrition, but also with perfect compliance, analysis is a major concern. Regarding the analysis, Abadie and Imbens (p. 236, 2006) note that matching estimators for the average treatment effect "include a conditional bias term whose stochastic order increases with the number of continuous matching variables". They show that the simple matching estimator is not $N^{1/2}$ efficient and propose an alternative. Imai, King, and Nall (2009) claim that the variance can be consistently estimated, but they refer to the variance not conditional on covariates (Imbens, 2011). Imbens (p. 17, 2011) writes that "this variance is larger than the conditional one if treatment effects vary by covariates. In stratified randomized experiments we typically estimate the variance conditional on the strata shares, so the natural extension of that to paired randomized experiments is to also condition on covariates." In contrast to this, BK09 estimate the variance conditional on pair dummies, but not conditional on covariates. It thus seems that even among experts, it is unclear how to correctly assess the variance of estimates.

With respect to attrition, Imai, King, and Nall (2009) note that an advantage of matching is that if a unit drops out, its pair can also be taken out of the experiment while the remaining sample still remains balanced. In the simulation study, we have seen that even with attrition, the matching techniques perform, on average, worse than the min MSE approach, which by design is unable to "correct" for attrition in this way. It thus might be perceived an overall disadvantage that for every unit dropping out of the experiment, its pair also has to be discarded, as this leads to a lower sample size, and consequently, lower power—irrespective of the exact nature of the treatment effect.

11. BK09 note that in the 300 observation sample, the algorithm takes several days to run and that ample time is needed to perform matching techniques.

Rerandomization Methods Bruhn and McKenzie (p. 210, 2009) name a key advantage of ad-hoc rerandomization methods: "[They] may offer a way of obtaining approximate balance on a set of relevant variables in a situation of multiple treatment groups of unequal sizes." As pointed out before, for the approach by Morgan and Rubin (2012), an extension allowing the assignment of multiple treatment groups is theoretically possible, but missing to date.

Furthermore, implementation time may differ considerably—depending on the approach. Drawing a thousand treatment assignments may take some time, and in small samples with many variables, the 'big stick' method that aims at finding a treatment assignment where no difference in group means exceeds a certain t-value might need even more time.

Yet, all of those rerandomization methods aim at balancing group means, and with the exception of Morgan and Rubin (2012), fail to consider dependencies of the different variables included in treatment assignment. However, in their approach, the dependency between variables is constant accross treatment assignments.

Nevertheless, all of the rerandomization methods discussed here are able to consider continuous, categorial and binary variables in a theoretically unlimited number.

The approach by Kasy (2016) In a Bayesian setting, Kasy (2016) analyzes the task of treatment assignment from a decision theoretical perspective, where the mean squared error of an estimator is to be minimized. Kasy (2016) argues that randomization never increases precision. In the technical appendix of his paper, he discusses several modelling aspects of conditional expectations of potential outcomes, which are the basis of his analysis. One of the discussed models for potential outcomes is the bayesian linear model, which gives rise to a treatment assignment mechanism using the framework of his paper.

While Kasy (2016) provides software implementation in Matlab, an extension to treatment assignment for multiple groups is neither discussed nor implemented. Moreover, the Bayesian setting requires the choice of parameter values that are hard to guess without analyzing the pre-treatment version of the outcome of interest such as the covariance matrix of the estimator of the parameter vector in the linear model or the coefficient of determination of this model. Yet, in practice, a pre-treatment version of the outcome of interest may be unavailable. The choices, however, are consequential, as they distort the balance of treatment groups, which might be desired if the interest is limited to the precision of estimation of the specified treatment effect. In case the researcher is interested in comparable treatment groups for the sake of credible research, or in treatment effects for a specific subgroup, the balance in treatment groups might be an equally important goal.

However, the approach is—as the other rerandomization methods—able to consider any type of variable and any number of covariates desired. Moreover, dependencies between variables are taken into consideration and dispersion of variables within groups is encouraged—if not impeded by the researcher through an unwise choice of parameter values.

The min MSE Treatment Assignment Mechanism The min MSE Treatment Assignment Mechanism retains the advantages of the linear model as discussed by Kasy (2016). In particular, it is able to perform treatment assigment using possibly various continuous, categorical and binary variables. In that aspect, it is as powerful as all other rerandomization schemes and the matching approach.

Contrarily to the Bayesian model applied in Kasy (2016), we rely on a standard frequentist model. This relieves the researcher from the obligation to choose abstract parameters, for which arguably a good guess is sometimes impossible. Moreover, this setup ensures that dependencies amongst variables are equally distributed accross treatment groups and that dispersion of variables is maximized within treatment groups—without any possibly unwanted distortion. Thus, if possible, subgroup analysis is ensured as the distribution of variables in treatment groups is taken into consideration.

Still, our min MSE approach allows for considerable flexibility if one has reasons to believe that an outcome will have a higher variance than the other in general; and more specifically, that treatment t, which mainly affects outcome k, might have compliance problems, resulting in an expected higher variance than for the control group. By this feature, of all treatment assignment mechanisms we are aware of, the min MSE approach comes closest to the ability of the 'optimal matching' approach to achieve approximate equality of the distributions of the covariates in the treatment and control groups—while being able to assign multiple treatment groups. Moreover, the min MSE treatment assignment mechanism is attrition tolerant in the sense that the balance stays favorable and the sample size is reduced only by the units dropping out.

Finally, software implementation for Stata is available. Thus, implementation takes only a few minutes and may easily be performed in the field.

Chapter 5

The min MSE Stata ado-Package

5.1 Introduction

The min MSE ado-package is the software implementation of the treatment assignment mechanism introduced in the paper for the use with the statistical software Stata.

When implementing an experiment, researchers can use this software to assign subjects to different treatment groups before treatment is executed. Treatment assignment is conducted based on observed pre-treatment characteristics, e.g. age, income or education. The created treatment groups are *balanced* and, within the model laid out in the paper, the mean squared error of the estimator for the conditional average treatment effect is minimized. Note that the analysis of the experiment can be based on a different model than the one used the paper; in particular, additional variables to the ones used for treatment assignment can be included in the analysis.

The min MSE ado-package is suited for two use cases: First, for treatment assignment before any treatment is executed and second, in case treatment is not implemented at the same time for all units and attrition is happening, the `assignment` option can be used for re-balancing the treatment groups, given the units and their treatment group for which treatment has already been implemented.

The package provides a convenient single command line interface for the user. Being implemented with the use of Stata's fast Mata language, this software makes treatment assignment possible in less than five minutes and is thus ideally suited for implementation in the field.

In this article, I provide an overview of the functionality of the package and briefly explain its usage.

5.1.1 Software Requirements

The min MSE ado-package provides the `assignMinMSETreatment` command for Stata.The ado-file has to be read in after starting Stata—either automatically[1] or manually. It is available at http://minMSE.sebastianschneider.info. The min MSE ado-package was written for Stata 13.1 and has been tested with Stata 14.

1. For an ado-file to be automatically loaded, it has to be placed in a folder listed by Stata's `sysdir` command. It is recommended to create a folder for user installed ado-packages (e.g. Personal) if not already existing.

5.2 Overview and Reference Manual

The command `assignMinMSETreatment` takes one or several variable name(s) as argument(s) and returns treatment group number(s) in a variable to be created.[2]

The number of participants might be uneven and the software deals with missing values automatically.

Minimization of the statistic derived in the paper is performed by a combination of random draws to find suitable starting points and a variant of the simulated annealing (Kirkpatrick, Gelatt, and Vecchi, 1983) optimization algorithm with default values as implemented in the ***R*** command **`optim`**[3]; see details below.

General Syntax

The general syntax for treatment assignment (in the conventional Stata notation) is

`assignMinMSETreatment` *`varlist`*`, generate(`*`newvar`*`) [`*`options`*`]`

where

`varlist` is `var1[ var2[...]]` and used for covariate input and

`newvar` is the name of the variable to be created containing treatment group number(s).

Options

Any of the following options might be specified.

`treatments(#)` specifies the number of treatment groups desired (in addition to the control group); minimum and default value is `treatments(1)`.

`iterations(#)` specifies the number of iterations the algorithm performs; the default value is `iterations(50)`. With small samples and few covariates, depending on the desired method of inference, a relatively small value is recommended, see Section 5.3 for details. Depending on the number of units and the number of covariates to consider for group assignment, a high value could result in a long run-time.

2. Therefore, the data has to be organized such that rows consist of an individual's covariates and columns contain individual values of a single covariate, consistent with the Stata-typical workflow.

3. As Stata lacks a built-in routine to perform optimization using the simulated annealing algorithm, we relied on default values of an established implementation in ***R***.

`assignment(varname)` takes a numerical vector of partial treatment assignment as argument, and assigns the missing units (where *varname* == .) to a treatment group while minimizing the objective function. Non-missing values are copied to the new vector, i.e. treatment group assignment of these observations is unaffected.

Advanced Options

The following options control the minimization process performed by the simulated annealing algorithm. The `plot` option might be useful for adjusting these parameters. Default values (except for the `plot` option) correspond to the default values applied in the ***R*** implementation of simulated annealing (in the command **`optim`**).

`change(#)` sets the number of units to exchange treatment in each iteration; the default value is `change(3)`. In case of big datasets (e.g. with more than 100 units), one might consider increasing the default value.

`cooling(#)` specifies the cooling scheme for the simulated annealing algorithm to use. `cooling(1)`, which is the default scheme, sets the temperature to

$$t_0/\log(\text{floor}((k-1)/t_{\text{max}})t_{\text{max}} + \exp(1)),$$

whereas `cooling(2)` sets the temperature to the faster decreasing sequence

$$t_0/(\text{floor}((k-1)/t_{\text{max}})t_{\text{max}} + 1).$$

In praxis, cooling schemes are mostly of one of these forms. One might want to change the cooling scheme if the plot indicates a too slow deacrease of objective values. For a theoretical discussion of cooling schemes Belisle (see 1992, p. 890).

`t0(#)` sets the starting temperature for the simulated annealing algorithm, see Belisle (1992) for theoretical convergence considerations. In praxis, a lower starting temperature t_0 decreases the acceptance rate of a worse solution more rapidly. Specifying a negative number allows values proportional to the objective function, i.e. $t0(-5)$ sets the starting temperature to $1/5$ of the objective function for the starting point, and thus—for the first t_{max} iterations of the algorithm—the difference of the old and the proposed solution is scaled by $1/5$. When changing the default value, it should be considered that also worse solutions have to be accepted in order for the algorithm to escape a local minimum, so it should be chosen high enough. The default value is `t0(10)`.

`tmax(#)` specifies the number of function evaluations at each temperature: For instance, `tmax(10)` makes the algorithm evaluate 10 treatment assignments that are found based on the current solution, before the temperature is decreased and thus the probability of accepting a worse solution is decreased. The default value is `tmax(10)`.

`plot(#)` can be used to suppress drawing a plot showing the value of the objective function for the last iterations by setting `plot(0)`. The default setting is `plot(1)`, which shows a plot. While the convergence plot is a helpful tool for setting the control parameters of the simulated annealing algorithm and for detecting convergence, it might be less interesting when generating a big number of alternative treatment assignments e.g. for performing Fisher's exact test.

5.3 Remarks and Examples

In case Fisher's exact test should be applied to evaluate significance of estimates, a certain number of alternative treatment group assignments are needed. Those alternative treatment group assignments should be established under the same conditions as the true (and implemented) treatment group assignment. While it is in many cases possible to gather the alternative treatment assignments after having conducted the experiment, it should be noted that especially in small samples and with few (possibly categorical) variables this could be problematic. In those cases, it could happen that the algorithm always proposes the same treatment group assignment. Then, computing Fisher's exact test is impossible and bootstrap mechanisms have to be applied (see for example Bertsimas, Johnson, and Kallus, 2015).

Therefore, if computation of Fisher's exact p-values is desired, it is advised to have the algorithm run several times (e.g. 10 times in small samples with up to 30 individuals per treatment groups) and compare the value of the objective function and compare treatment assignments. If among those treatment assignments duplicates are present, the number of iterations for the algorithm to run is too high and should be decreased. Alternatively, the number of covariates could be increased.

In the near future, a function for automatically establishing alternative group assignments to compute exact p-values will be implemented. This function will automatically adjust the number of iterations. For now, this task has to be performed by the researcher—given exact p-values should be computed.

Example 1

The following command creates a treatment and a control group, balancing age and gender:

```
assignMinMSETreatment age gender, gen(Treatment)
```

Treatment numbers 0 and 1 are then stored in the variable *Treatment*.

Example 2

Partial rebalancing of treatment assignment can be performed using the following command, where overall balance of age and gender should be achieved (across units that already have been treated and those still to be treated):

```
assignMinMSETreatment age gender, gen(Treatment) ///
    assignment(treatmentDayOne)
```

As in Example 1, treatment numbers are stored in the variable *Treatment*, and *treatmentDayOne* contains the treatment group numbers of individuals already treated (missing for non-treated).

Example 3

Assigning 5 treatment groups in addition to the treatment group (which will have the number 0), can be achieved with the following command:

```
assignMinMSETreatment age gender income hasChildren, ///
    gen(Treatment) treatments(5) iterations(100) ///
    change(5) cooling(2) t0(-3) tmax(50) plot(0)
```

This will perform 100 iterations of the algorithm, where in each step of the algorithm 5 randomly selected subjects exchange their treatment status. The cooling scheme of the form $1/n$ is selected, with starting temperature set to $1/3$ of the value of the first treatment assignment evaluated. In this example, the temperature decreases every 50 iterations. Plotting the convergence curve is supressed.

5.4 Conclusion and Future Work

The min MSE ado-package for Stata implements the treatment assignment procedure as derived in the paper. This method has theoretically appealing features, such as balance in the dependence of covariates within a group as well as balance in the dispersion of the covariates. As implemented in the min MSE ado-package, this method is particularly suited for treatment assignment in the field, as it is very fast compared to several competing methods.

Furthermore, it features the possibility to 'rebalance': For example in interventions that are implemented over several days and in several locations, attrition might happen and might leave the sample unbalanced. The min MSE package provides an option to adjust treatment assignment in those cases. Treatment group assignment of subjects already treated is taken into account while the yet untreated are possibly assigned a different treatment group to keep overall balance.

While the implemented features already make the min MSE package a practical tool for applied researchers, its applicability will be additionally increased. First, the possibility to weight treatments, outcomes or variances will be implemented. In a second step, the possibility to conveniently prepare randomization inference using e.g. Fisher's exact test will be added by gathering a number of prespecified alternative treatment group assignments under the same conditions as the true treatment group assignment.

Chapter A

Appendix: Income Risk, Precautionary Saving, and Loss Aversion: An Empirical Test

A.1 Theoretical Framework: Details

A.1.1 General Version of the Two-period Model by Kőszegi and Rabin (2009)

As in Section 2.2, 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.

$$U = u_1(c_1) + \mathbb{E}[u_2(c_2)]. \tag{A.1}$$

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.[1] Favourite means that there are possibly several such credible plans, and the decision-maker

1. Details about how these plans are formed are given in Appendix A.1.2 or in Kőszegi and Rabin (2009).

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 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 $\lambda > 0$.[2]

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 from deviating from the plan, instantaneous utility in period 1 is given by

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

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 it 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 from 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.1.2 Rational Beliefs

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

2. 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, Minehart, and Rabin (1999).

'Beliefs' are the result of a plan: They "must be rationally based on credible plans for state-contingent behavior"[3]. 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 because they 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.

A.1.3 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)

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

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. □

A.2 Data: Details

A.2.1 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.[4] Thus, when allowing for gains and losses, the *power family* reduces to

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

Figure A.1(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

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

4. 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 exptected to occur."

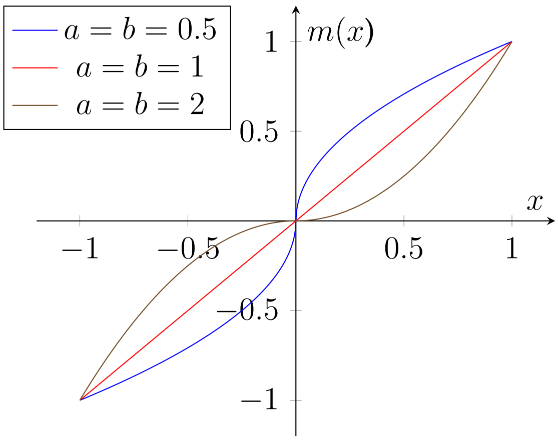

(a) Curvature of the power family for different values of a and b.

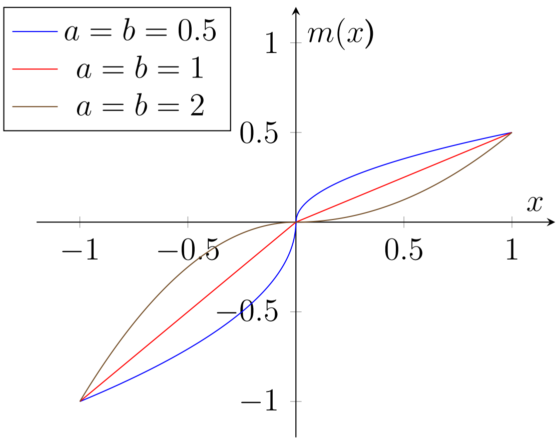

(b) Estimated power utility functions plotted for different values of a and b.

FIGURE A.1: Illustration of the Power Family Utility Function with Different Values of a and b

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:

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

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

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

This equation is illustrated in Figure A.1(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 A.1 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.

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

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

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

A.3 Results: Details

A.3.1 Discussion: Model Choice

In this part, we shortly discuss alternatives to the model chosen and assess their appropriateness in the setting of Chapter 2

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 2.4 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 is a requisite 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 is 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.

A.4 Further Results and Robustness Checks

Figures A.2 and A.3 show conditional marginal effects for the likelihood of saving and conditional savings with confidence intervals.

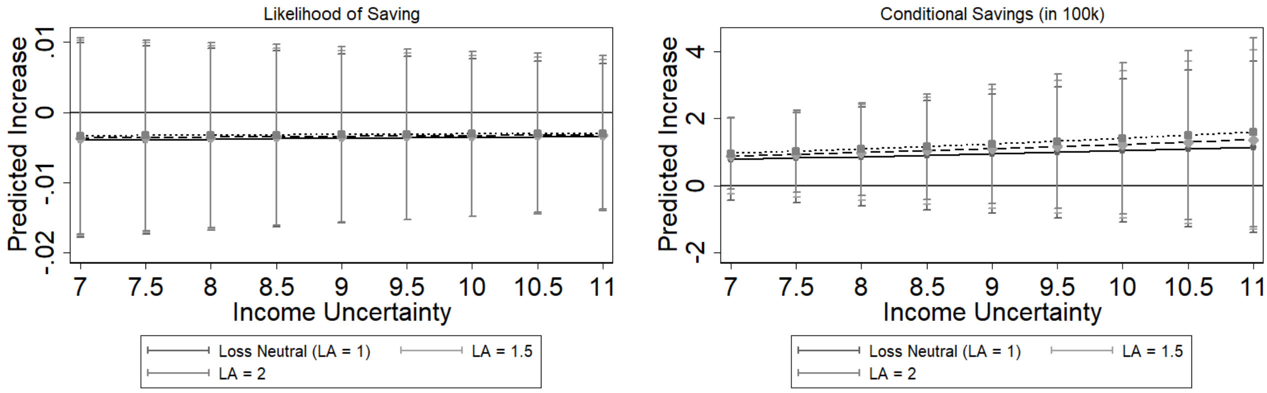

(a) Measure of Income Uncertainty: Regional Unemployment Rate (DANE)

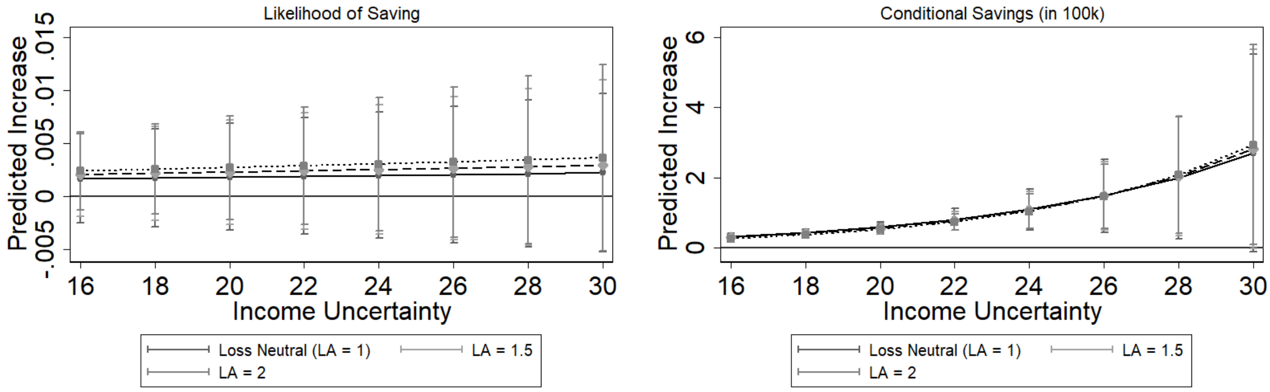

(b) Measure of Income Uncertainty: Regional Unemployment Risk (Survey)

FIGURE A.2: Conditional Marginal Effect of Income Uncertainty on the Predicted Likelihood to Save (left) and on Predicted Conditional Savings (right) for Different Degrees of Loss Aversion

Note: Mean values of the covariates used for prediction. Estimates for calculating the marginal effects result from fitting Model 3; see Table 2.4 for the corresponding coefficients. Income Uncertainty is expressed in percent and the predicted increase in conditional savings in 100,000 COP.

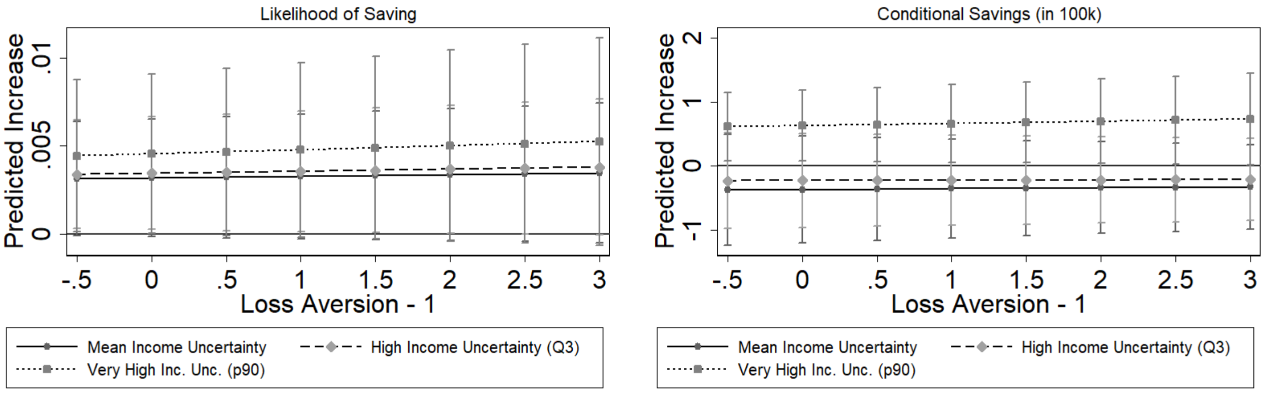

(a) Measure of Income Uncertainty: Regional Unemployment Rate (DANE)

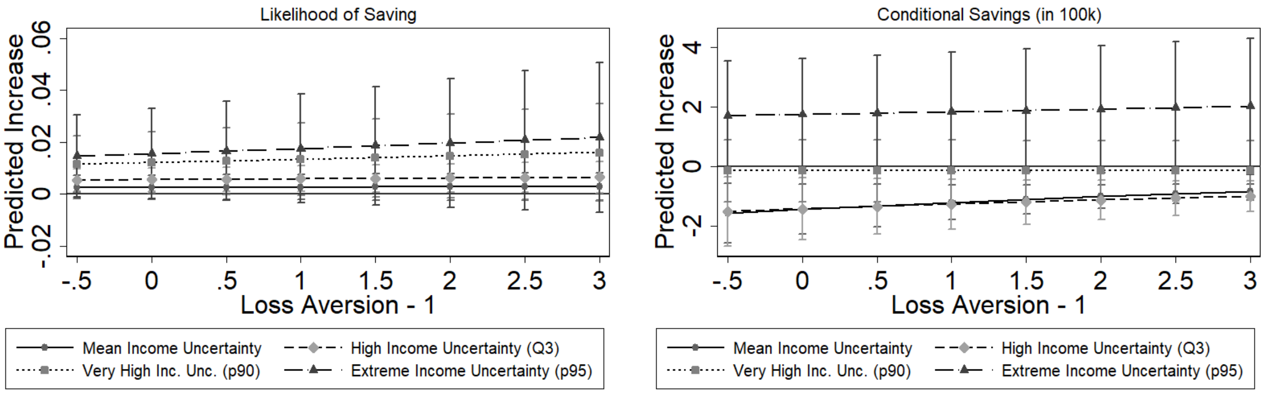

(b) Measure of Income Uncertainty: Regional Unemployment Risk (Survey)

FIGURE A.3: Conditional Marginal Effect of Loss Aversion on the Predicted Likelihood to Save (left) and on Predicted Conditional Savings (right) for Different Levels of Income Uncertainty

Note: Mean values of the covariates used for prediction. Estimates for calculating the marginal effects result from fitting Model 3; see Table 2.4 for the corresponding coefficients. The predicted increase in conditional savings is expressed in 100,000 COP.

Chapter B

Appendix: Higher Order Risk: An Application To Savings of the Poor in Bogota

B.1 Dynan (1993)

Dynan (1993) assumes a concave utility, which is additive over time and establishes the following link between expected consumption growth, prudence and the variation in consumption growth:

$$\mathbb{E}_t\left[\frac{c_{i,t+1}-c_{i,t}}{c_{i,t}}\right]=\frac{1}{\rho_r}\left(\frac{r_i-\delta}{1+r_i}\right)+\frac{\xi}{2}\mathbb{E}_t\left[\left(\frac{c_{i,t+1}-c_{i,t}}{c_{i,t}}\right)^2\right], \tag{B.1}$$

where $\mathbb{E}_t$ is the expectation conditional on information available at time t, $c_{i,t}$ represents consumption, δ is the constant time preference rate, and r_i is the household's interest rate. $\rho_r = -c_{i,t}(u''/u')$ is the coefficient of relative risk aversion and $\xi = -c_{i,t}(u'''/u'')$ is the coefficient of relative prudence as defined by Kimball (1990). This equation has been derived using a Taylor series approximation for $u'(c_{i+1,t})$ around $c_{i,t}$:

$$u'(c_{i,t+1}) \approx u'(c_{i,t}) + u''(c_{i,t})(c_{i,t+1}-c_{i,t}) + \frac{u'''(c_{i,t})}{2}(c_{i,t+1}-c_{i,t})^2. \tag{B.2}$$

Note that this approximation does not impose any structural constraints on any of the derivatives. This result is combined with the first order condition Dynan (1993, Equation 3) resulting from solving the maximization problem as stated in

Dynan (Equation 1 1993):

$$\left(\frac{1+r_i}{1+\delta}\right)\mathbb{E}_t\left[u'(c_{i,t+1})\right]=u'(c_{i,t}). \tag{B.3}$$

What is problematic about that approach is that ξ as used in equation (B.1) does only correspond to Kimball's definition, when $u'''(c_{i,t})u''(c_{i,t}) < 0$. Only then, it can be interpreted as a measure of prudence with respect to a precautionary savings demand. From the description of the empirical approach applied by Dynan (1993), it does not seem like this is taken into account. In addition, the estimations of $1/\rho_r$ (the coefficients of $(r_i-\delta)/(1+r_i)$ and its large standard errors suggest that the sign of ρ_r is not necessarily positive.

B.2 Shortcomings of Parametric Utility Functions for the Study of Higher Order Risk Preferences

Widely used parametric functions are not suited for (empirically) studying higher order risk preferences, because they are too stylized to possess the flexibility to combine any shape of the utility function with any shape of the second, third and higher order derivatives. This is in particular true for utility functions belonging to the family of switching sign utilities (i.e. functions where $\operatorname{sgn}(u^{(n)}) = -\operatorname{sgn}(u^{(n-1)})$). According to these functions, a risk-averse individual (negative second derivative) would always be classified as prudent (positive third derivative) and temperant (negative fourth derivative), whereas risk-seeking individuals (positive second derivative) would always be classified as imprudent and intemperant (negative third and positive fourth derivative, respectively).

We illustrate this shortcoming for the power (CRRA) utility family.

Power utility family For $x > 0$, the family is defined as

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

Note that, if u is an interval scale (meaning u is unique up to unit and level), it can be multiplied by any positive factor and any constant can be added without affecting any relevant empirical aspect (Wakker, 2008). This is, in mathematical terms, a monotonic transformation, that does not affect the maximization process. In particular, u as defined above is then equivalent to alternative formulations (with restricted domain, i.e. where $b < 1$) such as

$$u(x) = \begin{cases} \frac{x^{1-\eta}-1}{1-\eta} & \text{if } \eta > 0, \eta \neq 1 \\ \ln(x) & \text{if } \eta = 1, \end{cases}$$

where $b = \eta - 1$.

Thus, we analyze the shape of utility and its derivatives according to the power family as defined in its general definition. The second derivative is given by

$$u''(x) = \begin{cases} \frac{\partial}{\partial x} bx^{b-1} = b(b-1)x^{b-2} & \text{if } b > 0 \\ \frac{\partial}{\partial x}\frac{1}{x} = -\frac{1}{x^2} & \text{if } b = 0 \\ \frac{\partial}{\partial x} - bx^{b-1} = -b(b-1)x^{b-2} & \text{if } b < 0. \end{cases}$$

Note that as $x > 0$, the first derivative of u is positive for all $b \in \mathbb{R}$. The second derivative is negative and the utility itself concave for $b < 1$ and thus, in the expected utility framework would correspond to a risk averse individual.

The third derivative is given by

$$u'''(x) = \begin{cases} \frac{\partial}{\partial x} b(b-1)x^{b-2} = b(b-1)(b-2)x^{b-3} & \text{if } b > 0 \\ \frac{\partial}{\partial x} - \frac{1}{x^2} = \frac{1}{x^3} & \text{if } b = 0 \\ \frac{\partial}{\partial x} - b(b-1)x^{b-2} = -b(b-1)(b-2)x^{b-3} & \text{if } b < 0. \end{cases}$$

For $b \leq 0$, we see immediately that the third derivative is positive. For $0 < b < 1$ and $b > 2$, we have a positive third derivative. For $1 < b < 2$, we have a negative third derivative.

Thus, only for an b in the interval $(1, 2)$, could the third derivative be negative. On that interval however, the second derivative is always positive, and thus—using this utility family—an imprudent individual can never be risk averse. Similarly, a risk averse individual can never be imprudent.

Exponential utility family We now turn to the exponential (CARA) utility family. Assuming u is unique up to unit and level, the formulation of this family again does not depend on multiplicative factors or the addition of constants (Wakker, 2008).

We define the exponential family for $x > 0$ as:

$$u(x) = \begin{cases} (1 - e^{-bx})/b & \text{if } b \neq 0 \\ x & \text{if } b = 0. \end{cases}$$

Its derivatives are given by

$$u'''(x) = \begin{cases} \frac{\partial^2}{\partial x^2} e^{-bx} = \frac{\partial}{\partial x} - be^{-bx} = (-b)^2 e^{-bx} & \text{if } b \neq 0 \\ \frac{\partial^2}{\partial x^2} 1 = \frac{\partial}{\partial x} 0 = 0 & \text{if } b = 0. \end{cases}$$

We see that the third derivative can never be negative and thus, assuming this utility family, we would never classify an individual as imprudent. Further, the fourth derivative will always have the same sign as the second derivative, and so we can never classify an individual as intemperant and risk averse relying on the utility functions of this family.

Expo-Power Family The expo-power family has been proposed by Abdellaoui, Barrios, and Wakker (2007) and was used e.g. by Holt and Laury (2002). It exhibits an increasing measure of relative risk aversion in x and a decreasing measure of absolute risk aversion in x for $0 < b < 1$ and x in the interval $[0, 1]$ (resulting from normalizations of x, see Abdellaoui, Barrios, and Wakker, 2007). On the interval $(0, 1]$, it is defined by

$$u(x) = \begin{cases} -e^{-x^b/b} & \text{if } b \neq 0 \\ -1/x & \text{if } b = 0. \end{cases}$$

We take a look at its derivatives:

$$u''(x) = \begin{cases} \frac{\partial}{\partial x} e^{-x^b/b} x^{b-1} = e^{-x^b/b} x^{b-2} \left(-x^b + b - 1\right) & \text{if } b \neq 0 \\ \frac{\partial}{\partial x} 1/x^2 = -2/x^3 & \text{if } b = 0. \end{cases}$$

For $b = 0$, the second derivative is always negative and the function itself is concave. For $b \neq 0$, the sign of the second derivative depends on the sign of $-x^b + b - 1$. It is negative for $b - 1 < x^b$. As $x \in (0, 1]$, this is the case for $b < 1$. Contrarily, for $b > 2$, the above term is always positive and so is the second derivative, implying a convex utility function.

Let's now turn to the third derivative:

$$u'''(x) = \begin{cases} \frac{\partial}{\partial x} u''(x) = e^{-x^b/b} x^{b-3} \left(b^2 + x^{2b} + 3x^b - 3b(x^b + 1) + 2\right) & \text{if } b \neq 0 \\ \frac{\partial}{\partial x} u''(x) = 6/x^4 & \text{if } b = 0. \end{cases}$$

For $b \neq 0$, the sign of the third derivative depends on the term $b^2 + x^{2b} + 3x^b - 3b(x^b + 1) + 2$. Numerically, one finds that this term is uniformly negative in $x \in (0, 1]$ for values of b between roughly 1.27 and 2. That is, for a risk-loving or a risk-averse individual (i.e. individuals to which a paramter value of $r < 1$ and $r > 2$ correspond, respectively) there is at least one point $x \in (0, 1]$, for which the third derivative is positive. Therefore, neither a risk averse nor a risk loving individual will ever be classified as imprudent according to this utility family.

B.3 Methodology: Details

B.3.1 Derivatives of a B-spline Function

As laid out in De Boor (1987, Ch. 10, Eqs. (12) & (16)), the derivative of a B-spline function $f(x)$ is given by

$$\begin{aligned} f^{(1)}(x) = \frac{\partial f(x)}{\partial x} &= \frac{\partial}{\partial x} \sum_{j=1}^{k+p-1} a_j B_j(x,p) = \sum_{j=1}^{k+p-1} a_j \frac{\partial}{\partial x} B_j(x,p) \\ &= \sum_{j=1}^{k+p} (a_j - a_{j-1}) \frac{1}{h} B_j(x, p-1) = \frac{1}{h} \sum_{j=1}^{k+p} \triangle a_j B_j(x, p-1), \end{aligned} \tag{B.4}$$

where h is the knot distance, and $\triangle a_j := (a_j - a_{j-1})$, $a_0 := 0$ and $a_{k+p} := 0$. Note that, sticking to the indices, the B-splines $B_j(x, p-1)$ vanish on the interval $[x_{\min}, x_{\max}]$ for $j < 2$ and $j > k + p - 1$. Accounting for this fact and iteratively applying (B.4) yields

$$\begin{aligned} f^{(m)}(x) &= \frac{1}{h^m} \sum_{j=1}^{k+p+m-1} \triangle^m a_j B_j(x, p-m) \\ &= \frac{1}{h^m} \sum_{j=m+1}^{k+p-1} \triangle^m a_j B_j(x, p-m), \end{aligned} \tag{B.5}$$

where $\triangle^m a_j = \triangle(\triangle^{m-1} a_j)$ is the m-order difference of the sequence (a_j).

B.3.2 Choice of Penalty: Jointly Smoothing Multiple Derivatives

For their exposition of P-splines, P. H. Eilers and Marx (1996) use an unspecified order d of penalization. A penalty based on the second derivative was introduced in the smoothing context by Reinsch (1967), mainly "because it leads to a very simple algorithm". Probably, a penalty of order 2 is still the most common penalty used. However, P. H. Eilers and Marx (1996) note that, besides convenient computation, there is no specific reason for this choice.

Here, we are interested in smoothing the utility function itself, and additionally the third and the fourth derivative of the utility function. In this regard, P. H. Eilers and Marx (1996) note that the "kth difference penalty is a good discrete approximation to the integrated square of the kth derivative"[1]. Further, for a penalty of order d, the fitted curve approaches a polynomial of degree $d-1$, as the penalty increases (P. H. Eilers and Marx, 1996). Lastly, interpolation is affected by the order of the penalty: With a penalty of order d, interpolation of utility is of degree $2d-1$. This means that the third derivative of the interpolating curve has degree $2d-4$ and the fourth derivative has degree $2d-5$ (P. H. C. Eilers and Marx, 2010).

Considering these aspects suggests using a penalty of order 3 or 4, where the latter is to be preferred for the limiting behavior of the spline function when the penalty increases and for the degree of the fourth derivative of the interpolation curve. However, this choice leads to fluctuations in the third derivative not caused by the data, so additionally introducing a penalty on the third derivative is necessary.

The objective function then writes

$$Q_B(\boldsymbol{\alpha}) = \|\boldsymbol{y} - \boldsymbol{B\alpha}\|^2 + \omega_3\|\boldsymbol{D_3\alpha}\|^2 + \omega_4\|\boldsymbol{D_4\alpha}\|^2. \tag{B.6}$$

Penalization of multiple orders has been applied in other studies, remarkably with a focus on the quality of prediction (i.e. interpolation): Marx and Eilers (2002) introduced the use of multiple orders independently, including a ridge penalty (corresponding to order $d=0$) in addition to any penalty order $d=1,2,3$. Aldrin (2006) shows in a simulation experiment that the prediction performance when penalizing both slope ($d=0$) and curvature ($d=2$) is always at least as good as penalizing curvature only.

Our goal is to smooth the third and fourth derivative jointly corresponding to using both penalties of order $d=3$ and $d=4$. When choosing an optimal parameter with respect to prediction quality, the third order penalty and the fourth order

1. They illustrate this point with a penalty on second order differences.

penalty are to a high degree exchangeable. As the third order differences of the B-spline coefficients generally will be much higher than the fourth order differences, the third order penalty will in general dominate the fourth order penalty, causing the effect of the latter to vanish and resulting in unnecessary fluctuations of the fourth derivative. Using a penalty of order $d = 4$ alone, however, results in unnecessary fluctuations of the third derivative.

This issue has been dealt with by Eilers and Goeman (2004), who present the first approach we are aware of to jointly smooth multiple orders $d > 0$. They study how signals consisting of largely flat areas combined with a sharp pulse can be smoothed. This phenomenon resembles to a certain degree the pattern of a considerable share of our utility functions, where parts of nearly zero marginal utility follow parts of sharp increases or vice versa. Eilers and Goeman (2004) use a combination of a first and second order penalty by setting $\omega_2 = \nu\omega_1^2$, corresponding to penalty terms of order $d = 2$ and $d = 1$, where they found by trial and error that $\nu = 1/4$ works well.

We consider a couple of hundreds of utility functions, and in this case, visual inspection is clearly too time consuming. Furthermore, the choice of ν should be independent of subjective judgment as it could affect the classification and intensities of risk preferences. Therefore, we develop a data-driven approach.

We propose a solution in which the fourth order penalty 'drives the shape' of the utility function while the third order penalty is limited to avoiding unnecessary fluctuations in the third derivative. This is achieved by using one penalty parameter, appropriately scaled for both orders.

Specifically, we choose the scaling parameter ν such that

$$\omega_4 \|\boldsymbol{D_4 \alpha_0}\| \approx \nu\omega_4 \|\boldsymbol{D_3 \alpha_0}\|,$$

i.e. such that the penalty terms have approximately equal weight, where we used $\nu = 0.001$ for a computation of $\|\boldsymbol{D_3 \alpha_0}\|$. Then, we set $\omega_3 = \nu\omega_4/5$, to ensure $\omega_4\|\boldsymbol{D_4 \alpha}\| > \omega_3\|\boldsymbol{D_3 \alpha}\|$.

The choice of ν for a first computation of $\|\boldsymbol{D_3 \alpha_0}\|$ practically neglects any third order penalty and a good fit using solely the fourth order penalty is achieved. Then, ν is set as a fifth of the ratio of the sum of absolute differences of the fourth order differences of the B-spline coefficients over the third order differences. The ratio measures how much stronger a third order penalty will affect the smoothing process as compared to a fourth order penalty and will in most cases be the maximum factor needed to weight both penalties equally. To avoid the third order penalty dominating the fourth order penalty, we divide this ratio by five, which proved to yield the desired behavior.

Thus, we elaborate an approach which is to be preferred when visual inspection is inappropriate—be it for objectivity reasons or for time constraints impeding the researcher to visually investigate a large amount of individual curves.

B.3.3 Incorporating a Monotonicity Contraint In P-Spline Regression

The rationale of the approach by Bollaerts, Eilers, and Mechelen (2006) to approximately incorporate a monotonicity constraint in P-Spline regression is simple: Although in P-Spline regression, a penalization term is added to the objective function, the predicted function is still a B-spline function, and thus, its derivatives are given by equations (B.4) and (3.10). From these formulae, a sufficient condition for the first derivative to be positive and the utility to be a monotone increasing function can be easily deduced: Since h, the knot distance, is positive and B non-negative for all x, p and j, all $\triangle a_j$ have to be positive.

Thus, differences of coefficients of the B-splines that are negative have to be penalized, whereas positive differences do not, which makes the penalty asymmetric. This penalization is achieved with the following penalty:

$$\sum_{j=2}^{k+p-1} w(\alpha)_j (\triangle a_j)^2,$$

where

$$w(\alpha)_j = \begin{cases} 0, & \text{if } \triangle a_j \geq 0 \\ 1, & \text{otherwise.} \end{cases}$$

The objective function (B.6) now writes

$$Q_B(\boldsymbol{a}) = \|\boldsymbol{y} - \boldsymbol{B\alpha}\|^2 + \nu\omega_4\|\boldsymbol{D_3\alpha}\|^2 + \omega_4\|\boldsymbol{D_4\alpha}\|^2 + \kappa\|\boldsymbol{W^{1/2}D_1\alpha}\|^2 \tag{B.7}$$

which—in case Q_B is convex—is minimized if

$$(\boldsymbol{B'B} + \nu\omega_4\boldsymbol{D_3'D_3} + \omega_4\boldsymbol{D_4'D_4} + \kappa\boldsymbol{D_1'WD_1})\hat{\boldsymbol{\alpha}} = \boldsymbol{B'y},$$

where $\boldsymbol{y}$, $\boldsymbol{B}$, $\boldsymbol{D_d}$ and ω_4 are defined as in (B.6) and ν as determined in Section B.3.2. $\boldsymbol{W}$ and $\boldsymbol{W^{1/2}}$ are diagonal matrices with elements $w(\alpha)_j$ and $\sqrt{w(\alpha)}$, respectively, and the impact of the constraint penalty on the solution is tuned by a sufficiently high[2] (positive) constraint parameter κ.

Bollaerts, Eilers, and Mechelen (2006) show that Q_B is convex in $\boldsymbol{\alpha}$ and propose using a Newton-Raphson procedure to find an optimal solution. We follow this suggestion, stopping the algorithm after 10 iterations, which led to a monotone increasing function in 99.5% of all cases.

B.3.4 Choosing the Degree of Smoothness

P. H. Eilers and Marx (1996) suggest two classical objective and data-driven selection criteria, AIC and (leave-one-out) cross validation, for the choice of the degree of smoothness of the P-spline function to be established as tuned by ω in order to balance between fidelity to the data and smoothness.

Cross validation (CV) is independent of distributional or asymptotic assumptions and "should be preferred to any model selection procedure relying on assumptions which are likely to be wrong" (Arlot and Celisse, 2010). In our case, for the computation of AIC, at least the asymptotic argument that the approximation of standard errors relies on is unlikely to hold for our moderate number of data points. We thus apply cross validation as selection criteria for choosing ω_4 as in (B.7). The principle of CV is simple: The model is fit with only a part of the data and the remainder is used to compute prediction errors. With leave-one-out CV, for N data-points, the model is fit N times, and each time, one data point is left out and used for computation of the prediction error. Then, the model—in our setting the penalty parameter—is chosen that minimizes the average prediction error over N predictions. For this case, exact formulae for convenient computation exist without the need to actually fit the model N times (P. H. Eilers and Marx, 1996; Eilers, Marx, and Durbán, 2015).

In literature, however, it has been noted that leave-one-out CV is not the ideal choice regarding model choice (see e.g. Kohavi, 1995, and the references therein). One argument against leave-one-out CV is that the probability of choosing the model with the best predictive quality does not converge to 1 as the number of observation increases (Shao, 1993). Further, Eilers, Marx, and Durbán (2015) warn that leave-one-out CV severly overfits the data in case of correlated observations.

We address this issue with the following strategy: As proposed in literature for the purpose of model identification (e.g. Shao, 1993; Arlot and Celisse, 2010), we increase the number of points left out for prediction, which—fortunately,

2. We chose $\kappa = 10^8$, Bollaerts, Eilers, and Mechelen (2006) chose $\kappa = 10^6$ in their application.

with a maximum of 9 elicited points—is computationally still feasible to do exhaustively. More specifically, we perform permuted leave-k-out cross validation (Aldrin, 2006): We choose the degree of smoothness, that minimizes

$$\frac{1}{V}\sum_{v=1}^{V}\sum_{i\in I^{(v)}}\{y_i-\hat{y}_{(-v)}(x_i)\}^2, \tag{B.8}$$

where $V=\binom{N}{k}$ is the number of possibilities to chose k out of N points for validation, $I^{(v)}$ is the set containing the vth choice of k points for validation and $\hat{y}_{-v}(x_i)$ the prediction of y_i, obtained by estimating the model using all points but those in $I^{(v)}$.

Increasing the number of points left out for prediction results in less 'weight' for a single point; it can therefore be seen as a mean of error correction, as long as the remaining points are sufficient to establish a meaningful utility curve. As pointed out in the main text, reversal rates of one third are common in choice tasks.[3] For this reason, we perform 'leave-at-least-$1/3N$-out' cross validation. In case the maximum number of utility points was elicited for the individual under study,[4] this choice results in leave-3-out CV. This strategy also accounts for correlated observations (Arlot and Celisse, 2010, Chapter 8.1) possibly resulting from the chain structure of the experiment.

In addition, we develop and apply a data-driven minimum for the penalty parameter to rule out overfitting resulting from large distances between utility points. We compute the number of balls with radius equal to the knot distance needed to cover the elicited points. For a minimum of two balls, overfitting the data is impossible and the method has to compromise between data-fidelity and smoothness. For a maximum of nine balls, however, and for low values of ω_4, the fitted function usually perfectly predicts every data point used for estimation, and minimizes the prediction error for those points left out for validation. Thus, according to CV, the minimal ω_4 is chosen in those cases—but the fitted function is considerably overfitting the data. We impede this by setting a higher minimal penalty parameter in those cases.

Specifically, for individual i, the minimal smoothness parameter is calculated using the following formula:

$$\omega_{4,i}^{\min}=(b_i*(n_{\max}/n_i)-1)^{2.5} \tag{B.9}$$

3. See e.g.Abdellaoui, Bleichrodt, and Paraschiv (2007), Fennema and Van Assen (1998), Abdellaoui (2000), and Etchart-Vincent (2004)

4. Note that we excluded the points $(0,0)$ and $(1,1)$ for computing the average prediction error.

where b_i is the number of balls with radius equal to the knot-distance needed to cover all elicited utility points of individual i, $n_{\max}$ is the maximum number of elicited points possible for all individuals (we have $n_{\max} = 9$) and n_i is the number of elicited points for individual i.

If the maximum of 9 points are elicited, and all points have a pair-wise distance above the knot-distance, then the number of balls to cover all points will be 9, and the minimal value for $\omega_{4,i}$ will be roughly 180. This value is still low enough to allow the fitted function to be a polynomial of degree $p > 3 = d - 1$, but in most cases, it is high enough to prevent overfitting. In some cases, the data is overfit, indicating that the minimal smoothness parameter is chosen conservatively. If all elicited points in $(0, 1)$ lie close together, the minimal smoothness parameter $\omega_{4,i}$ would be 1, i.e. a minimal smoothness parameter that results in hardly any penalization.

For some individuals, less than 9 utility points were elicited, since due to the implementation of the protocol followed (Abdellaoui, Bleichrodt, and Paraschiv, 2007), in some choice tasks, one option is stochastically dominated and the resulting utility point has to be erased following, (e.g. Bleichrodt and Pinto, 2000). If the number of elicited utility points is less or equal to the order of penalty, we have to perform leave-one-out CV according to the formula by P. H. Eilers and Marx (1996). In those cases, overfitting is only prevented by the increased minimal ω_4 in case of sparse information per data knot as expressed by formula (B.9).

B.4 Full Tables and Robustness Tests

TABLE B.1: Net Savings and Prudence (C&E) Showing Coefficients on Covariates

	Full Sample			Risk Averse		Risk Loving	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Prudence (C&E)	1.140***	1.161***	1.114**	1.707***	1.468***	0.560	0.651
	(0.422)	(0.422)	(0.431)	(0.641)	(0.559)	(0.601)	(0.724)
Risk Aversion (A&P)		3.329*	3.291		0.637		6.038
		(2.002)	(2.134)		(3.120)		(5.325)
Male			-7.647		-9.299		-2.863
			(6.847)		(8.889)		(10.34)
Age			0.101		0.0610		0.426*
			(0.152)		(0.186)		(0.218)
Financial literacy			-1.601*		0.390		-3.841**
			(0.933)		(1.109)		(1.702)
BMI			-0.161		-0.846		0.783
			(0.618)		(0.809)		(1.000)
Adult HH members			-2.588		-5.937*		-0.357
			(1.868)		(3.196)		(2.720)
Children HH members			0.192		1.434		0.0408
			(2.243)		(2.934)		(3.564)
Income			-0.695		2.053		-2.382
			(1.392)		(1.909)		(2.085)
Impatience			0.0535		0.0442		0.170
			(0.169)		(0.240)		(0.255)
Increase in patience			0.246		0.338		0.196
			(0.243)		(0.288)		(0.396)
Planning horizon			ref.		ref.		ref.
– Next months			9.439*		3.429		16.65*
			(5.138)		(6.698)		(8.949)
– Next year			-1.060		2.360		-4.154
			(9.930)		(15.58)		(13.28)
– Next two to five years			13.02*		5.171		31.24**
			(6.971)		(8.873)		(14.90)
– 5 or more years			11.20*		7.431		15.19
			(6.667)		(11.00)		(10.89)
Constant	-14.69***	-14.64***	-0.453	-11.78***	17.46	-17.04***	-44.73
	(2.722)	(2.705)	(22.32)	(3.443)	(30.90)	(4.321)	(29.45)
Education	No	No	Yes	No	Yes	No	Yes
Observations	567	567	554	270	267	279	271

Notes: This table reports the results of ordinary least squares regressions on net savings. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. $Income$ in 100k Colombian pesos. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

TABLE B.2: Net Savings and Prudence (C&E & Kimball) Showing Coefficients on Covariates

	(1)	(2)	(3)	(4)	(5)	(6)
Prudence (Kimball)	12.73⋆		7.236		7.245	
	(7.428)		(4.774)		(6.190)	
Prudence (C&E)		2.191⋆		1.096⋆		1.293
		(1.206)		(0.619)		(0.795)
Risk Aversion (A&P)			12.55	13.47	10.24	10.18
			(9.618)	(10.36)	(9.797)	(9.678)
Male					6.227	5.018
					(10.17)	(9.790)
Age					0.164	0.190
					(0.333)	(0.333)
Financial literacy					0.915	0.798
					(2.237)	(2.213)
BMI					-1.093	-1.244
					(1.081)	(1.156)
Adult HH members					-4.062	-4.327
					(4.260)	(4.258)
Children HH members					2.146	2.291
					(4.330)	(4.326)
Income					4.647⋆	4.850⋆
					(2.756)	(2.808)
Impatience					0.742⋆⋆	0.745⋆⋆
					(0.363)	(0.360)
Increase in patience					-0.0836	-0.109
					(0.219)	(0.216)
Planning horizon					ref.	ref.
– Next months					-2.176	-0.177
					(15.66)	(15.13)
– Next year					32.58	33.40
					(20.36)	(21.08)
– Next two to five years					10.77	9.571
					(17.66)	(15.76)
Constant	-14.31⋆⋆⋆	-15.79⋆⋆⋆	-33.37⋆	-35.49⋆	-37.68	-33.41
	(5.178)	(5.937)	(19.03)	(20.54)	(42.81)	(42.04)
Education	No	No	No	No	Yes	Yes
Observations	120	120	120	120	119	119

Notes: This table reports the results of ordinary least squares regressions on net savings. The prudence measure by Kimball is only defined for subjects that are risk averse at all points of evaluation, which reduces the number of observations to 120. For comparison, columns (2), (4) and (6) show the results for the C&E measure for the same sample. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. $Income$ in 100k Colombian pesos. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels ⋆ $p < 0.10$; ⋆⋆ $p < 0.05$; ⋆⋆⋆ $p < 0.01$.

TABLE B.3: Net Savings and Prudence (C&E) for Non-zero Net Savings Showing Coefficients on Covariates

	Full Sample			Risk Averse		Risk Loving	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Prudence (C&E)	2.297***	2.435***	2.749***	2.992**	2.629***	1.561	3.565**
	(0.861)	(0.863)	(0.804)	(1.168)	(0.974)	(1.358)	(1.712)
Risk Aversion (A&P)		7.030*	10.25**		-0.207		17.00
		(4.033)	(4.608)		(6.468)		(11.98)
Male			-17.84		-13.73		-13.11
			(14.30)		(17.98)		(23.27)
Age			-0.179		-0.660		0.858
			(0.380)		(0.472)		(0.575)
Financial literacy			-3.645*		2.987		-7.631**
			(2.024)		(3.019)		(2.963)
BMI			0.379		-0.615		2.338
			(1.314)		(1.700)		(2.451)
Adult HH members			-9.435**		-9.395		-9.444
			(4.334)		(5.929)		(7.890)
Children HH members			1.379		4.760		5.386
			(4.531)		(5.790)		(7.426)
Income			0.264		3.970		-1.639
			(2.272)		(3.577)		(3.549)
Impatience			-0.0205		-0.360		0.547
			(0.358)		(0.541)		(0.561)
Increase in patience			0.319		0.608		0.329
			(0.453)		(0.610)		(0.748)
Planning horizon			ref.		ref.		ref.
– Next months			14.86		0.637		31.79*
			(10.57)		(13.85)		(18.92)
– Next year			-2.107		4.974		-17.91
			(17.67)		(26.44)		(19.41)
– Next two to five years			40.30***		89.71**		71.81**
			(15.06)		(38.82)		(34.55)
– 5 or more years			27.02*				35.82
			(15.54)				(22.27)
Constant	-31.94***	-31.54***	-38.58	-23.41***	51.30	-37.80***	-202.0*
	(5.712)	(5.586)	(67.88)	(6.700)	(75.05)	(9.305)	(115.8)
Education	No	No	Yes	No	Yes	No	Yes
Observations	257	257	249	124	124	127	120

Notes: This table reports the results of ordinary least squares regressions on net savings. The sample is restricted to subjects who have non-zero net-savings. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. $Income$ in 100k Colombian pesos. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

TABLE B.4: Net Savings, Firm Closures and Prudence (C&E & Kimball) Showing Coefficients on Covariates

	Full Sample C&E		Risk Averse C&E		Strictly Risk Averse Kimball	
	(1)	(2)	(3)	(4)	(5)	(6)
Income risk	-1.746	-1.479	-1.863	-1.399	-4.985	-4.154
	(2.354)	(2.405)	(2.870)	(3.031)	(6.150)	(7.472)
Prudence	1.240***	1.115**	2.224***	2.036***	18.37*	9.263
	(0.461)	(0.471)	(0.769)	(0.739)	(9.852)	(9.574)
Prudence × Income risk	1.015**	1.066**	0.856	1.037*	9.654	12.99
	(0.424)	(0.416)	(0.600)	(0.566)	(9.694)	(9.948)
Risk Aversion (A&P)		4.644*		2.231		18.98
		(2.408)		(4.460)		(18.73)
Male		-10.50		-16.27		10.19
		(8.132)		(10.95)		(13.49)
Age		-0.00373		-0.155		-0.269
		(0.194)		(0.275)		(0.468)
Financial literacy		-1.469		1.966		3.523
		(1.208)		(1.465)		(3.197)
BMI		0.0508		-1.284		-1.391
		(0.721)		(1.146)		(1.473)
Adult HH members		-1.627		-3.763		-1.443
		(2.026)		(3.145)		(3.969)
Children HH members		-2.266		-1.202		-2.708
		(2.506)		(3.452)		(5.328)
Income		-0.889		3.040		5.866
		(1.637)		(2.577)		(4.058)
Impatience		-0.103		-0.127		0.558
		(0.181)		(0.325)		(0.406)
Increase in patience		0.381		0.431		-0.259
		(0.297)		(0.400)		(0.329)
– Next months		9.617		7.929		12.88
		(6.426)		(9.350)		(24.90)
– Next year		1.775		14.00		53.13
		(10.76)		(17.33)		(32.59)
– Next two to five years		14.13		27.74		55.94*
		(9.795)		(25.44)		(29.22)
– 5 or more years		13.87		-6.131		
		(9.759)		(11.93)		
Constant	-15.33***	-6.033	-12.58***	20.66	-16.45**	-61.69
	(2.995)	(29.66)	(4.158)	(39.63)	(6.542)	(71.05)
Controls	No	Yes	No	Yes	No	Yes
Observations	471	459	218	215	93	92

Notes: This table reports the results of ordinary least squares regressions on net savings. Prudence is the Crainich-Eckhoud measure of prudence in columns (1) to (4) and the Kimball measure in columns (5) and (6). Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. Income risk is measured as the ratio of closed to existing businesses in 2013 in the working sector an individual was usually working in at the time of the survey. Prudence and income risk are centered. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. $Income$ in 100k Colombian pesos. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

TABLE B.5: Net Savings, Firm Closures and Prudence (C&E) Showing Coefficients on Covariates

	Full Sample		Risk Averse		Risk Loving	
	(1)	(2)	(3)	(4)	(5)	(6)
Income risk	-1.746	-1.479	-1.863	-1.399	0.768	2.146
	(2.354)	(2.405)	(2.870)	(3.031)	(3.293)	(3.412)
Prudence (C&E)	1.240***	1.115**	2.224***	2.036***	0.298	0.199
	(0.461)	(0.471)	(0.769)	(0.739)	(0.613)	(0.660)
Prudence (C&E) × Income risk	1.015**	1.066**	0.856	1.037*	0.984*	0.925
	(0.424)	(0.416)	(0.600)	(0.566)	(0.588)	(0.566)
Risk Aversion (A&P)		4.644*		2.231		9.500*
		(2.408)		(4.460)		(5.666)
Male		-10.50		-16.27		-4.157
		(8.132)		(10.95)		(12.26)
Age		-0.00373		-0.155		0.391
		(0.194)		(0.275)		(0.244)
Financial literacy		-1.469		1.966		-4.361**
		(1.208)		(1.465)		(2.085)
BMI		0.0508		-1.284		1.928*
		(0.721)		(1.146)		(0.994)
Adult HH members		-1.627		-3.763		0.386
		(2.026)		(3.145)		(3.102)
Children HH members		-2.266		-1.202		-3.604
		(2.506)		(3.452)		(3.174)
Income		-0.889		3.040		-3.419
		(1.637)		(2.577)		(2.248)
Impatience		-0.103		-0.127		0.0365
		(0.181)		(0.325)		(0.237)
Increase in patience		0.381		0.431		0.270
		(0.297)		(0.400)		(0.437)
Planning horizon		ref.		ref.		ref.
– Next months		9.617		7.929		12.01
		(6.426)		(9.350)		(9.271)
– Next year		1.775		14.00		-12.33
		(10.76)		(17.33)		(14.24)
– Next two to five years		14.13		27.74		33.20*
		(9.795)		(25.44)		(18.35)
– 5 or more years		13.87		-6.131		19.77
		(9.759)		(11.93)		(15.42)
Constant	-15.33***	-6.033	-12.58***	20.66	-17.37***	-53.10
	(2.995)	(29.66)	(4.158)	(39.63)	(4.519)	(33.94)
Controls	No	Yes	No	Yes	No	Yes
Observations	471	459	218	215	237	230

Notes: This table reports the results of ordinary least squares regressions on net savings. The sample is restricted to subjects who have non-zero net-savings. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. Income risk is measured as the ratio of closed to existing businesses in 2013 in the working sector an individual was usually working in at the time of the survey. Prudence and income risk are centered. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. $Income$ in 100k Colombian pesos. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

TABLE B.6: Net Savings, Firm Closures and Prudence (C&E) for Non-zero Net Savings Showing Coefficients on Covariates

	Full Sample		Risk Averse		Risk Loving	
	(1)	(2)	(3)	(4)	(5)	(6)
Income risk	-0.959	-0.760	-0.562	0.374	3.103	-0.0879
	(3.959)	(4.329)	(4.299)	(6.526)	(5.844)	(6.291)
Prudence (C&E)	2.247**	2.484***	3.628***	3.478***	0.856	2.309*
	(0.871)	(0.785)	(1.289)	(1.302)	(1.302)	(1.327)
Prudence (C&E) × Income risk	1.551**	1.841***	1.433	1.702	1.574	2.398**
	(0.683)	(0.684)	(0.935)	(1.080)	(0.988)	(1.154)
Risk Aversion (A&P)		11.43**		1.979		11.00
		(4.615)		(7.762)		(10.39)
Male		-25.27		-28.75		-19.87
		(15.93)		(21.98)		(26.13)
Age		-0.116		-0.631		0.913
		(0.436)		(0.653)		(0.698)
Financial literacy		-2.727		5.485		-7.977**
		(2.272)		(3.483)		(3.374)
BMI		0.908		-0.860		4.629**
		(1.383)		(1.784)		(2.314)
Adult HH members		-9.051*		-5.144		-7.627
		(5.125)		(7.019)		(10.06)
Children HH members		-1.777		-0.121		2.373
		(4.964)		(6.405)		(6.734)
Income		-0.724		2.981		-2.307
		(2.578)		(4.009)		(3.928)
Impatience		-0.335		-1.074		0.229
		(0.373)		(0.687)		(0.509)
Increase in patience		0.697		0.880		0.801
		(0.523)		(0.700)		(0.810)
Planning horizon		ref.		ref.		ref.
– Next months		9.929		2.249		27.81
		(13.33)		(17.55)		(21.05)
– Next year		0.588		12.86		-27.17
		(17.11)		(28.80)		(24.58)
– Next two to five years		46.50**		126.6**		83.38*
		(21.63)		(51.44)		(45.08)
– 5 or more years		33.90*				40.41
		(18.62)				(25.39)
Constant	-30.33***	-56.16	-20.82***	54.77	-36.02***	-252.6**
	(5.829)	(80.54)	(7.246)	(96.24)	(9.216)	(117.6)
Controls	No	Yes	No	Yes	No	Yes
Observations	231	223	111	111	114	107

Notes: This table reports the results of ordinary least squares regressions on net savings. The sample is restricted to subjects who have non-zero net-savings. Prudence (C&E) is the Crainich-Eckhoud measure of prudence. Risk Aversion (A&P) is the Arrow-Pratt measure of risk aversion. Income risk is measured as the ratio of closed to existing businesses in 2013 in the working sector an individual was usually working in at the time of the survey. Prudence and income risk are centered. The controls are time preferences, gender, age, financial literacy, body mass index (BMI), household members (adults and children), income as measured as the average income per household member, planning horizon and education. $Income$ in 100k Colombian pesos. We account for potential heteroskedasticity by robust standard errors. Results of t-tests indicated at following significance levels * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Chapter C

Appendix: The Min MSE Treatment Assignment Method

C.1 Proofs

Proof of Theorem 1. Similar to the definition of X_t, define the subvector of the k-th potential outcome $Y^{p,k}_{\{i:D_i=t\},t} := (Y^{p,k}_{i_1,t}, Y^{p,k}_{i_2,t}, \dots, Y^{p,k}_{i_{n_t},t})^\top$ and the respective subvector of error terms $\varepsilon^{p,k}_{\{i:D_i=t\},t} := (\varepsilon^{p,k}_{i_1,t}, \varepsilon^{p,k}_{i_2,t}, \dots, \varepsilon^{p,k}_{i_{n_t},t})^\top$, where again $\{i_1, i_2, \dots, i_{n_t}\} = \{i : D_i = t\}$. That is, using observed information, (4.2) for all t and k in matrix notation writes $Y^{p,k}_{\{i:D_i=t\},t} = X_t^\top \beta^{p,k}_t + \varepsilon^{p,k}_{\{i:D_i=t\},t}$. For this linear model, it is well-known that for all $t = 0, 1, \dots, n_d$ and $k = 1, \dots, n_y$,

$$\hat{\beta}^{p,k}_t - \beta^{p,k}_t \sim \mathcal{N}(0, \sigma^2_{t,k}(X_t {X_t}^\top)^{-1}).$$

Using this result, the squared error of the estimator of the treatment effect for treatment t and outcome k becomes

$$\begin{aligned}
&\mathbb{E}\left[(\hat{\tau}_{t,k} - \tau_{t,k})^2 \mid X\right] \\
&= \mathbb{E}\left[\left(\frac{1}{N}\sum_i \left(\hat{Y}^{p,k}_{i,t} - \hat{Y}^{p,k}_{i,0}\right) - \frac{1}{N}\sum_i \left(E\left[Y^{p,k}_{i,t} \mid X_i\right] - E\left[Y^{p,k}_{i,0} \mid X_i\right]\right)\right)^2 \mid X\right] \\
&= \frac{1}{N^2}\mathbb{E}\left[\left(\sum_i X_i^\top \left(\left(\hat{\beta}^{p,k}_t - \beta^{p,k}_t\right) - \left(\hat{\beta}^{p,k}_0 - \beta^{p,k}_0\right)\right)\right)^2 \mid X\right] \\
&= \frac{1}{N^2}\sum_i X_i^\top \left(\mathrm{Cov}\left(\hat{\beta}^{p,k}_t - \beta^{p,k}_t \mid X\right) + \mathrm{Cov}\left(\hat{\beta}^{p,k}_0 - \beta^{p,k}_0 \mid X\right)\right)\sum_i X_i \\
&= \frac{1}{N^2}\sum_i X_i^\top \left(\sigma^2_{t,k}(X_t {X_t}^\top)^{-1} + \sigma^2_{0,k}(X_0 {X_0}^\top)^{-1}\right)\sum_i X_i,
\end{aligned}$$

where we used independence of the error terms $\varepsilon_{i,t}^{p,k}$ and $\varepsilon_{i,0}^{p,k}$.

Now denote the l_1 norm of a vector with $\| \cdot \|_1 = \sum |\cdot|$ and summarize weights and scaling factors for the variance as $\tilde{w}_k^y = w_k^y s_{0,k}$ and $\tilde{w}_t^d = w_t^d s_{t,k}$. Then, applying the just derived result to the objective function, the generalized MSE (4.1), completes the proof:

$$
\begin{aligned}
S_T(\hat{T}) = & \frac{\sigma_0^2}{N^2} \sum_i X_i^\top \\
& \left[\sum_k \left\{ w_k^y s_{0,k} \left(\sum_t w_t^d s_{t,k} \left(X_t X_t^\top \right)^{-1} + \| w^d \|_1 \left(X_0 X_0^\top \right)^{-1} \right) \right\} \right] \\
& \sum_i X_i \\
\propto & \frac{1}{N^2} \sum_i X_i^\top \\
& \left[\| \tilde{w}^y \|_1 \| w^d \|_1 \left(X_0 X_0^\top \right)^{-1} + \sum_k \left\{ \tilde{w}_k^y \left(\sum_t \tilde{w}_t^d \left(X_t X_t^\top \right)^{-1} \right) \right\} \right] \\
& \sum_i X_i,
\end{aligned}
$$

where $\propto$ denotes equality up to multiplicative constants. $\square$

Proof of Proposition 4. According to Cramer's rule, the inverse of the $p \times p$ matrix $\boldsymbol{A}$ is given by

$$
(A^{-1})_{ij} = \frac{(-1)^{i-j} \det(\boldsymbol{A_{ji}})}{\det(\boldsymbol{A})}, \tag{C.1}
$$

where $\boldsymbol{A_{ji}}$ is the $(p-1) \times (p-1)$ matrix resulting from deleting row j and column i.

Assume $n_d = 1$ and consider the first summand of (4.3) with a realization of the sample, i.e., $\bar{\boldsymbol{x}}^\top (\boldsymbol{X}\boldsymbol{X}^\top)^{-1} \bar{\boldsymbol{x}}$, where, for notational simplicity, we omit group indicators and $\bar{\boldsymbol{x}}$ is the vector containing the mean values of the covariates. Suppose the pth covariate vector $\boldsymbol{x_p}^\top$ is multiplied with some scalar $c \neq 0$ and denote the product with $\boldsymbol{w_p}^\top$. The covariate matrix changes accordingly from $\boldsymbol{X}$ to $\boldsymbol{W} = \mathrm{diag}(1, 1, \ldots, c)\boldsymbol{X}$. Thus, $\det(\boldsymbol{W}\boldsymbol{W}^\top) = c^2 \det(\boldsymbol{X}\boldsymbol{X}^\top)$.

Now consider the denominator of (C.1) for the possible combinations of i and j and denote $\boldsymbol{M} = \boldsymbol{X}\boldsymbol{X}^\top$, $\boldsymbol{N} = \boldsymbol{W}\boldsymbol{W}^\top$, and $\boldsymbol{M_{ij}}$, $\boldsymbol{N_{ij}}$ the matrices resulting from deleting row i and column j of the matrices $\boldsymbol{M}$ and $\boldsymbol{N}$, respectively. Since $\det(\boldsymbol{N_{ji}}) = c^2 \det(\boldsymbol{M_{ji}})$ for $i \neq p$ and $j \neq p$, $(M^{-1})_{ij} = (N^{-1})_{ij}$ in those cases. If either $j = p$ or $i = p$, we have $\det(\boldsymbol{N_{ji}}) = c\det(\boldsymbol{M_{ji}})$, thus $(N^{-1})_{ij} = 1/c(M^{-1})_{ij}$. Finally, for $i = p$ and $j = p$, as $\boldsymbol{N_{pp}} = \boldsymbol{M_{pp}}$, we have $(N^{-1})_{ij} = 1/c^2(M^{-1})_{ij}$.

Then, $\bar{\boldsymbol{w}}^\top \boldsymbol{N}^{-1}\bar{\boldsymbol{w}} = (1/n)^2 \sum\sum \bar{x}_i\bar{x}_j(M^{-1})_{ij} = \bar{\boldsymbol{x}}^\top \boldsymbol{M}^{-1}\bar{\boldsymbol{x}}$ applies to both summands of (4.3), also for $n_d > 1$, which completes the proof. □

Proof of Proposition 5. For $N \to \infty$, $\sigma^2_{t,k}\left(X_t {X_t}^\top\right)^{-1}$, the sample covariance matrix of $\hat{\beta}^{p,k}_t$, converges to the population covariance matrix for any $t = 0, 1, \ldots, n_d$ and any $k = 1, \ldots, n_y$. The elements of the inverse of the population covariance matrix are given by $\sigma^{-2}_{t,k}\mathbb{E}[X_{g,i}X_{j,i}]$ for all $g, j = 1, \ldots, m$. As $\mathbb{E}[X_{g,i}X_{j,i}] = 0$ for $g \neq j$, this is a diagonal matrix and so is its inverse, the population covariance matrix. Hence, $\left(X_t {X_t}^\top\right)^{-1}$ for $t = 0, 1, \ldots, n_t$ in (4.3) converges to a diagonal matrix with elements $\mathbb{E}[X^2_{j,i}]^{-1}$, $j = 1, \ldots, m$ for $N \to \infty$. Since $\frac{1}{N}\sum_i X_i$ is independent of treatment assignment and equals $c(1, \ldots, 1)^\top$ for some $c \neq 0$, for $N \to \infty$, (4.3) is minimized if

$$\sum_j \left[n_d \mathbb{E}[X^2_{j,i}]^{-1} + \sum_{t>0} \mathbb{E}[X^2_{j,i}]^{-1} \right]$$

is minimized. Noting that (4.4) converges to this sum as $N \to \infty$ completes the proof. □

Proof of Proposition 6. It is known that the diagonal elements of the covariance matrix of the estimator for the parameter vector in linear regression models such as (4.2) are given by

$$\mathrm{Var}(\hat{\beta}_j \,|\, X) = \frac{\sigma^2}{(1 - R^2_j)\sum_{i=1}^n (X_{j,i} - \bar{X}_j)^2},$$

for $j = 1, \ldots, m$, with notations as in Proposition 6 (see e.g. Wooldridge, 2013).

As in the proof of Theorem 1, the covariance matrix of $\hat{\beta}^{p,k}_t - \hat{\beta}^{p,k}_0$ for any $t = 1, \ldots, n_d$ and any $k = 1, \ldots, n_y$ is given by

$$\left(X_t {X_t}^\top\right)^{-1} + \left(X_0 {X_0}^\top\right)^{-1},$$

the claim follows, noting that we assume equal variances in Corollary 3. □

Proof of Theorem 2. We use $m_{i,j} = m_{j,i} = 1$ to indicate individual i is matched to individual j; otherwise $m_{i,j} = 0$. Every individual is matched exactly once, so $\sum_i \sum_j m_{i,j} = N$, assuming the sample consists of N individuals. Usually, in that case, the goal is to minimize $\sum_i \sum_j m_{i,j}(y_i - y_j)^2$ through the choice of $m_{i,j}$, although sometimes the absolute difference is also used (Rubin, 1973). For being a special case of the squared Mahalanobis distance, we prefer the squared euclidean distance. The set of solutions to this optimization problem is given by

$$\text{argmin}_{(m_{i,j})} \sum_i \sum_j m_{i,j}(y_i^2 + y_j^2 - 2y_iy_j) = \text{argmax}_{(m_{i,j})} \sum_i \sum_j m_{i,j}(y_iy_j).$$

We now show that elements of this set maximize the minimal sum of the variances of the groups to be created. This sum of variances is given by

$$(N/2)^{-1} \sum_{\{i:D_i=0\}} y_i^2 - \bar{y}_0^2 + (N/2)^{-1} \sum_{\{j:D_j=1\}} y_j^2 - \bar{y}_1^2 = (N/2)^{-1} = \sum_i y_i^2 - \bar{y}_0^2 - \bar{y}_1^2. \tag{C.2}$$

Since $\bar{y}_t^2 = (N/2)^{-2}(\sum_{\{i:D_i=t\}} y_i^2 + \sum_{\{i:D_i=t\}} \sum_{\{j:D_j=t,j\neq i\}} y_iy_j)$, for $t = 0, 1$, (C.2) can be rewritten as

$$((N/2)^{-1} - (N/2)^{-2}) \sum_i y_i^2 - (N/2)^{-2}(\sum_{\{i:D_i=0\}} \sum_{\{j:D_j=0,j\neq i\}} y_iy_j + \sum_{\{i:D_i=1\}} \sum_{\{j:D_j=1,j\neq i\}} y_iy_j).$$

The first summand is independent of group or treatment assignment. We rewrite the elements of the subtrahend as

$$\sum_{\{i:D_i=0\}} \sum_{\{j:D_j=0,j\neq i\}} y_iy_j + \sum_{\{i:D_i=1\}} \sum_{\{j:D_j=1,j\neq i\}} y_iy_j \tag{C.3}$$

$$= \sum_i \sum_j y_iy_j - \sum_i y_i^2 - 2 \sum_{\{i:D_i=0\}} \sum_{\{j:D_j=1\}} y_iy_j \tag{C.4}$$

$$= \sum_i \sum_j y_iy_j - \sum_i y_i^2 - 2 \sum_{\{i:D_i=0\}} \sum_{\{j:D_j=1\}} m_{i,j}y_iy_j - 2 \sum_{\{i:D_i=0\}} \sum_{\{j:D_j=1\}} (1 - m_{i,j})y_iy_j, \tag{C.5}$$

where we have split the cross product between group observations into those that are matched and those that are unmatched. The first two parts are again independent of group or treatment assignment, and so is the third for a fixed m. Thus, by matching we have maximized the sum of group variances across feasible treatment assignments. In other words, the sum of group variances resulting from the worst treatment assignment in this aspect from the set of possible treatment assignments after matching $\{D : D_i = |D_j - 1| \text{ for } \hat{m}_{i,j} = 1, \hat{m} \in \text{argmax}_{(m_{i,j})} \sum_i \sum_j m_{i,j}x_ix_j\}$ is still maximized over m. □

C.2 Additional Results

Table C.1: Comparison of Treatment Assignment Methods Regarding Balance in a Group of Baseline Variables (N=100)

(a) Average difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Indonesia	-0.056	-0.634	0.044	-0.151	0.111	0.892
Pakistan (height scores)	1.508	-0.667	-0.098	-0.105	0.578	0.263
Pakistan (test scores)	0.383	2.001	-0.330	-0.334	0.315	-0.112
Mexico	-0.383	-0.492	-0.097	0.074	0.565	-1.546
Sri Lanka	-0.633	0.632	0.708	-0.294	0.219	0.106

(b) Average abs. difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Indonesia	159.2	95.6	62.6	37.8	121.5	121.8
Pakistan (height scores)	159.7	85.7	52.3	65.4	116.8	95.2
Pakistan (test scores)	160.1	137.4	67.2	51.4	123.5	122.0
Mexico	159.2	118.0	56.5	44.3	134.6	124.4
Sri Lanka	159.1	128.5	75.7	46.1	126.2	120.6
Total	159.4	113.0	62.9	49.0	124.5	116.8

(c) 95% quantile of the difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Indonesia	409.8	263.9	199.7	260.9	390.3	397.5
Pakistan (height scores)	426.6	350.3	193.6	219.8	403.4	398.6
Pakistan (test scores)	394.7	438.6	218.4	276.4	394.7	407.8
Mexico	445.2	372.8	196.0	140.2	399.9	372.8
Sri Lanka	399.3	453.7	319.4	131.3	401.2	380.0
Total	415.1	375.8	225.4	205.7	397.9	391.3

Note: Statistics based on 10,000 iterations. Details on the study and the computation of each measures are explained in section 4.5.2. For every dataset, several variables were considered for treatment assignment. The results in this table report aggregate measures of differences in treatment group means for the group of considered variables. Lower values indicate a better balance with respect to equality of group means.

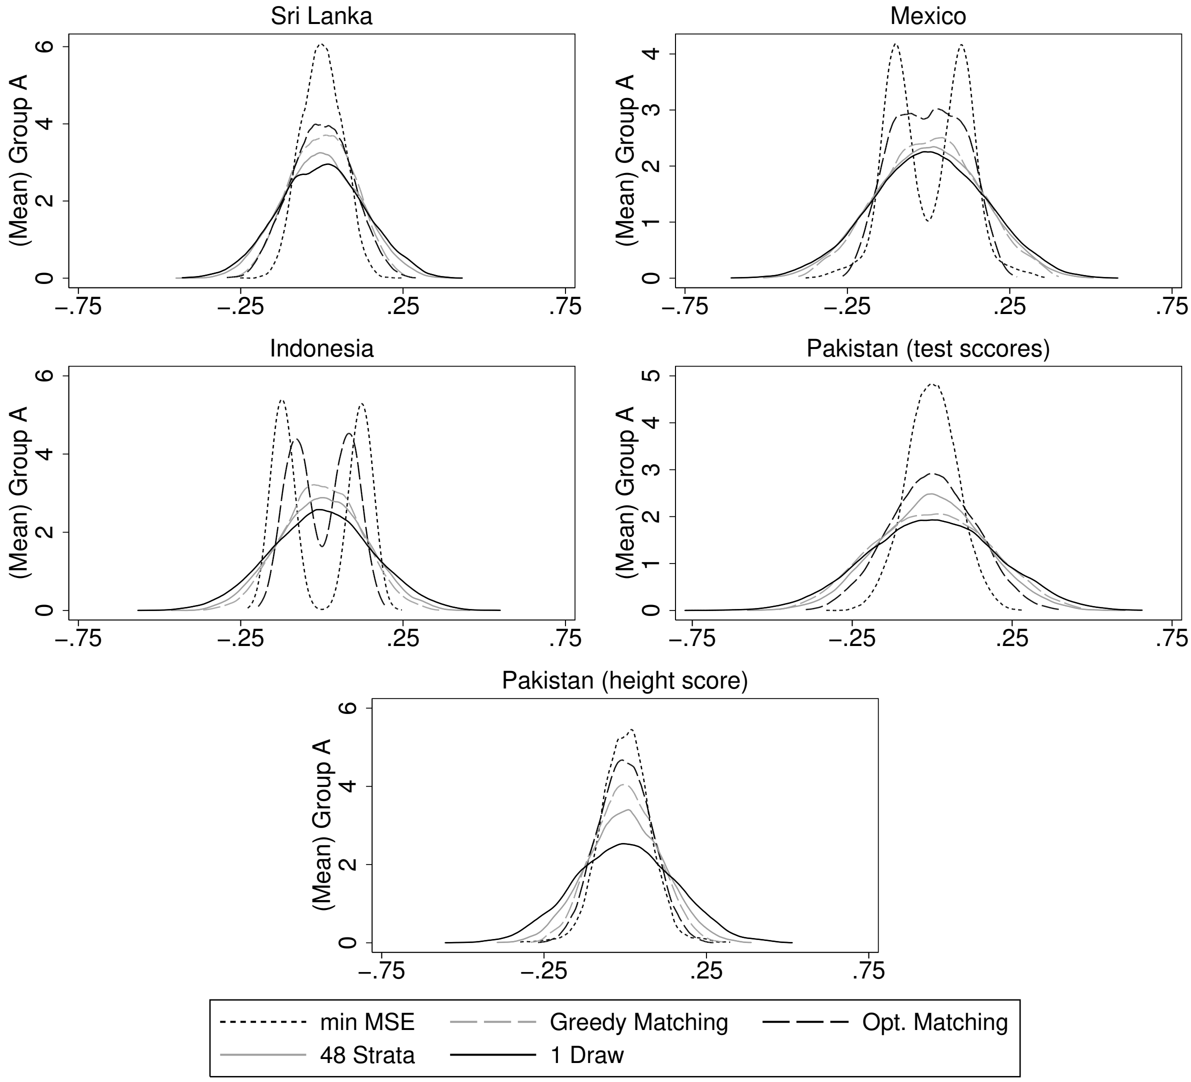

Note: Distributions of the (mean) differences in the group means of the respective group of variables are based on 10,000 treatment assignments, where the difference in one draw is the average of the distances in group means of all variables included in the assignment of treatments. Group means are expressed in standard deviations. A high mass around a difference of 0 indicates a good balance with respect to equality of the group means.

FIGURE C.1: Distributions of the Mean Differences in Group Means (N=30)

We present the results for the scenario considered in BK09, when focusing on one of the up to seven considered variables for treatment assignment. For two datasets, groups of variables—referred to as “unobservables”—are available, but not considered during treatment assignment. Balance on these variables is reported ‘group-wise’ (see Section 4.5.2).

Figures C.2 and C.3 show the distribution of the differences in group means for the indicated variables, which are the baseline and follow-up variable of interest, for the five datasets considered with sample sizes 30 and 100. Tables C.2 and C.3 consist of three panels: The upper panel shows the average difference in group means, the middle panel shows the 95% quantile of this difference and the lower panel shows the proportion of draws, in which the p-value of a t-test of the differences in group means was lower than 0.1.

For the single random draw method as well as the stratification methods, results are identical to those in BK09. Differences in the pairwise greedy matching approach are probably due to the order in which we run the scripts. However, the essential part of the do-file for performing the greedy matching is the same as the one provided by BK09.

The newly introduced methods—the optimal matching approach and the min MSE procedure—perform comparable to the others and the conclusion here is the same as in BK09, namely that “on average all methods lead to balance”.

In terms of average balance in baseline variables, we conclude that the min MSE procedure outperforms the other methods: For four of the five baseline variables and all unobservables, an average difference of zero to the third digit was achieved.

With respect to the whole distribution, as shown in Figures C.2 and C.3, the min MSE procedure shows the most favorable distribution with the highest mass at 0 and thinnest tails in half of the cases considered. Stratification seems to be superior in one dataset, where household expenditure is studied, whereas pairwise greedy matching dominates the competing mechanisms in achieving balance with the height z-score data. These findings are numerically underlined not only by the group means as discussed above, but also by the 95% quantile of the differences in group means as shown in the middle panel of Table C.3, although in this panel, no mechanism clearly shows more favorable figures than another.

Consistent with the findings of BK09, we also note that with increasing sample size, balance improves. This can be seen in Figures C.2 and C.3, where the distributions of group means for the bigger sample sizes are mostly nearly half as wide as the distributions on the left.

With respect to the balance observed in follow-up variables, we find, and consistent with BK09, no major differences; especially for the cases in which baseline variables explained little of the variation in follow-up outcomes (Microenterprise profits in Sri Lanka and Indonesian expenditure figures). For the bigger sample size, there is hardly any difference between the covariate based treatment assignment mechanisms.

Summing up, in most comparisons, either one of the matching methods or the min MSE procedure dominates the competing mechanisms. All methods achieve balance, on average, and all decrease extreme imbalances considerably in comparison with a single random draw. However, we think results that consider the overall balance of all variables considered in treatment assignment are more informative. These results are discussed in the main text, see Section 4.6.1, Tables 4.1 and C.1 and Figures 4.1 and C.1.

TABLE C.2: Comparison of Treatment Assignment Methods Regarding Balance in the Baseline Outcome (N=30)

(a) Average difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Microenterprise profits (Sri Lanka)	-4.214	-0.636	4.004	-1.302	-5.657	0.239
Household expenditure (Indonesia)	-2.094	-2.908	2.218	-2.794	0.644	2.739
Labor income (Mexico)	2.627	1.375	-1.346	-1.062	-0.774	-0.195
Height z-score(Pakistan)	-2.502	0.670	-3.486	0.034	-0.506	-0.228
Math test score (Pakistan)	-1.741	-0.632	-2.306	-0.301	-1.464	-1.571
Baseline unobservables (Sri Lanka)	-0.641	-0.331	0.837	0.896	-1.192	1.144
Baseline unobservables (Mexico)	-0.112	-0.130	-0.066	-0.390	-0.114	-0.673

(b) 95% quantile of the difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Microenterprise profits (Sri Lanka)	705.6	598.0	415.9	227.6	415.9	538.0
Household expenditure (Indonesia)	716.2	458.1	478.0	643.4	346.9	500.9
Labor income (Mexico)	690.8	176.5	223.7	228.1	409.1	581.8
Height z-score(Pakistan)	710.1	257.9	467.3	393.8	444.8	445.6
Math test score (Pakistan)	712.8	256.8	362.6	161.2	408.8	586.3
Baseline unobservables (Sri Lanka)	802.8	879.4	879.4	889.3	824.4	804.6
Baseline unobservables (Mexico)	834.3	834.3	879.4	834.3	771.3	774.9

(c) Proportion p-values <0.1 when testing the difference in baseline group means

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Microenterprise profits (Sri Lanka)	0.097	0.049	0.001	0.000	0.000	0.021
Household expenditure (Indonesia)	0.100	0.005	0.011	0.089	0.000	0.012
Labor income (Mexico)	0.099	0.000	0.000	0.000	0.000	0.037
Height z-score(Pakistan)	0.102	0.000	0.007	0.000	0.005	0.005
Math test score (Pakistan)	0.103	0.000	0.000	0.000	0.001	0.041
Baseline unobservables (Sri Lanka)	0.090	0.088	0.068	0.079	0.094	0.094
Baseline unobservables (Mexico)	0.088	0.077	0.083	0.079	0.077	0.079

Note: Statistics based on 10,000 iterations. Details on the study and the computation of each measures are explained in section 4.5.2. For every dataset, several variables were considered for treatment assignment. The results in this table consider the differences in treatment group means for a single variable of the group of considered variables. Lower values indicate a better balance with respect to equality of group means.

TABLE C.3: Comparison of Treatment Assignment Methods Regarding Balance in the Baseline Outcome (N=100)

(a) Average difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Microenterprise profits (Sri Lanka)	1.388	0.590	0.655	-0.638	-0.053	-0.777
Household expenditure (Indonesia)	-2.223	-1.665	-0.153	0.409	0.810	-0.679
Labor income (Mexico)	-0.428	-0.493	-1.051	-0.002	0.024	-0.295
Height z-score(Pakistan)	1.336	0.025	0.413	0.287	0.832	0.117
Math test score (Pakistan)	2.946	-0.260	-1.419	-0.288	-0.116	-0.555
Baseline unobservables (Sri Lanka)	-0.205	-0.593	1.102	-0.447	0.305	0.031
Baseline unobservables (Mexico)	0.135	-0.087	-0.113	-0.062	0.250	-0.218

(b) 95% quantile of the difference in baseline group means between the treatment and the control group in 1000 SD

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Microenterprise profits (Sri Lanka)	386.3	314.6	183.0	119.3	195.5	240.7
Household expenditure (Indonesia)	390.2	263.9	198.5	260.9	145.0	191.0
Labor income (Mexico)	383.9	99.5	153.8	100.1	280.2	304.0
Height z-score(Pakistan)	394.9	102.7	189.9	185.2	160.1	206.0
Math test score (Pakistan)	392.2	74.5	184.5	106.5	163.6	237.3
Baseline unobservables (Sri Lanka)	434.2	434.2	434.2	434.2	417.0	414.2
Baseline unobservables (Mexico)	456.5	456.5	456.5	456.5	447.6	439.0

(c) Proportion p-values <0.1 when testing the difference in baseline group means

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
Microenterprise profits (Sri Lanka)	0.097	0.039	0.000	0.000	0.000	0.005
Household expenditure (Indonesia)	0.102	0.010	0.001	0.013	0.000	0.000
Labor income (Mexico)	0.100	0.000	0.000	0.000	0.015	0.029
Height z-score(Pakistan)	0.100	0.000	0.000	0.000	0.000	0.001
Math test score (Pakistan)	0.100	0.000	0.000	0.000	0.000	0.006
Baseline unobservables (Sri Lanka)	0.101	0.083	0.097	0.091	0.096	0.095
Baseline unobservables (Mexico)	0.108	0.104	0.091	0.089	0.095	0.093

Note: Statistics based on 10,000 iterations. Details on the study and the computation of each measures are explained in section 4.5.2. For every dataset, several variables were considered for treatment assignment. The results in this table consider the differences in treatment group means for a single variable of the group of considered variables. Lower values indicate a better balance with respect to equality of group means.

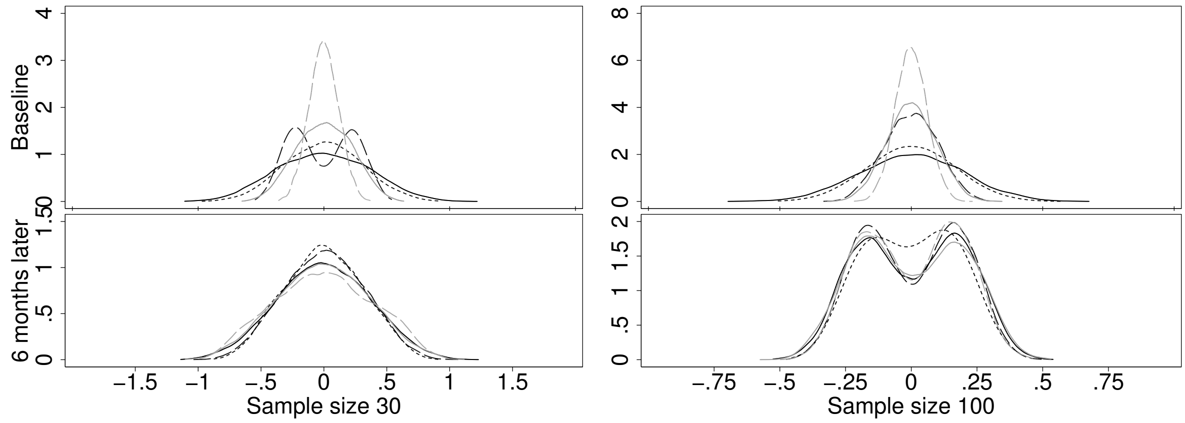

(a) Panel A. Sri Lanka. Differences in average profits (weighted by standard deviation)

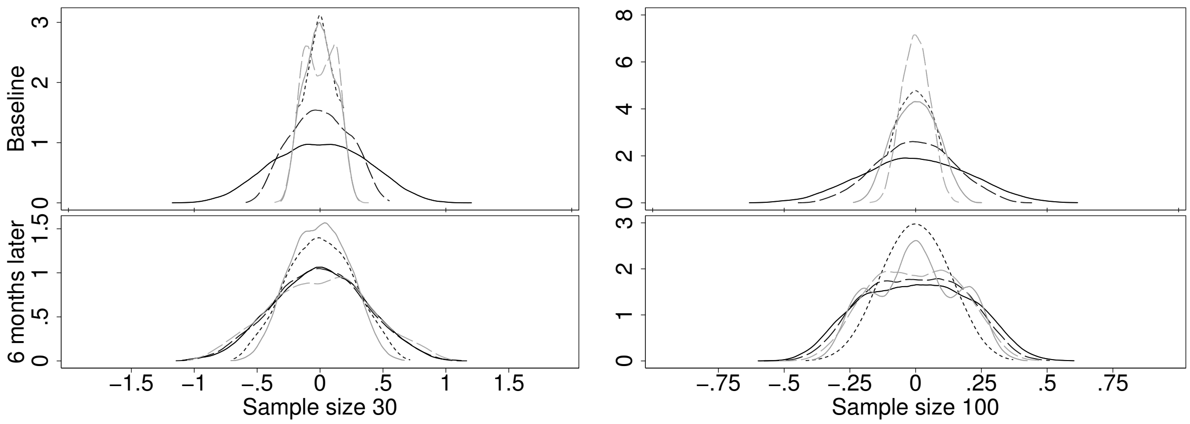

(b) Panel B. Mexico ENE. Differences in average income (weighted by standard deviation)

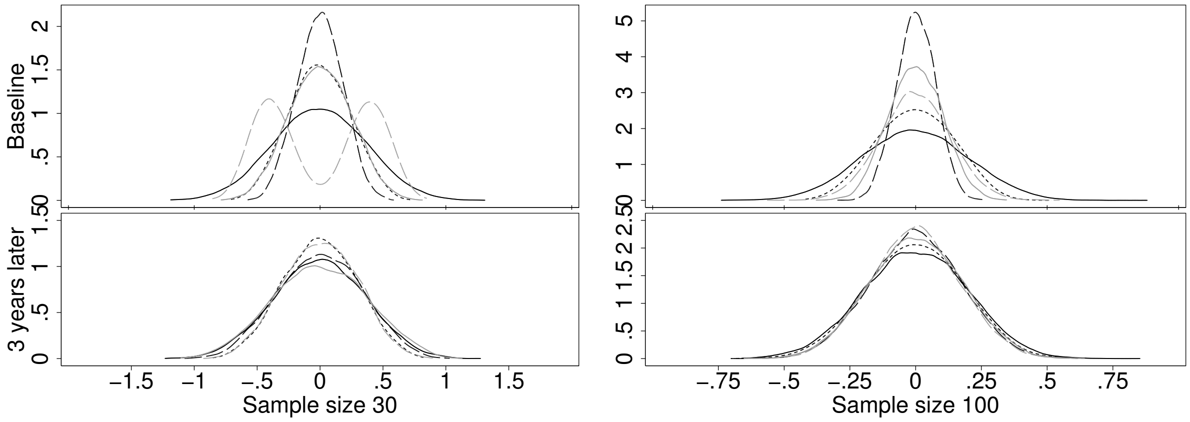

(c) Panel C. IFLS. Differences in average *hh* expenditure *p cap* (weighted by standard deviation)

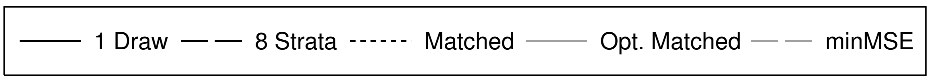

Note: Panel A, B and C show the distributions of the differences in treatment group means of the indicated variable for 10,000 treatment assignments, expressed in standard deviations. The higher the mass around a difference of 0, the better the treatment assignment with respect to balancing group means. Baseline variable is the outcome of interest; one of up to six variables used for treatment assignment. Follow-up variable is the same variable, measured six months after the baseline variable was collected; note that it is not included in treatment assignment.

FIGURE C.2: Distributions of the Differences in Group Means Between the Treatment and the Control Group in the Baseline Variable and the Follow-up Variable (N=30)

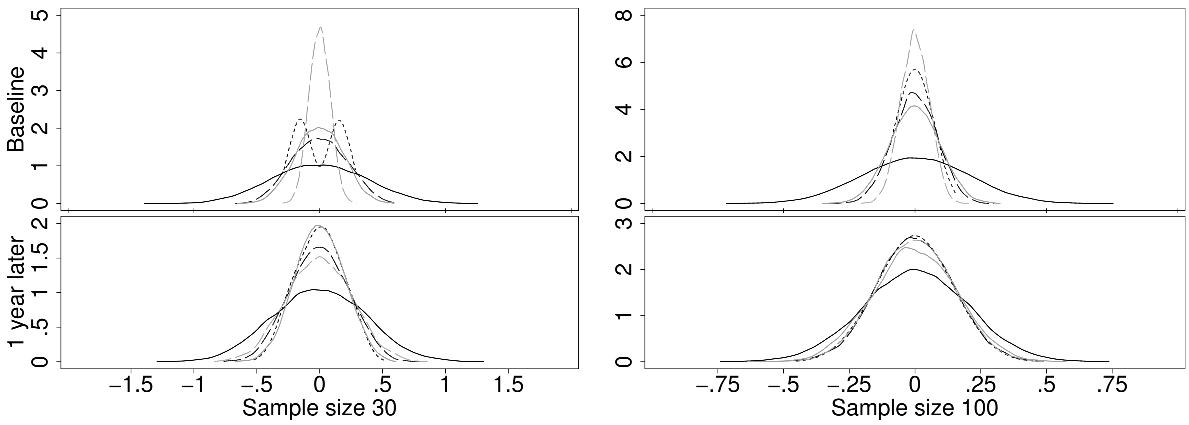

(a) Panel D. LEAPS. Differences in average in math test score (weighted by standard deviation)

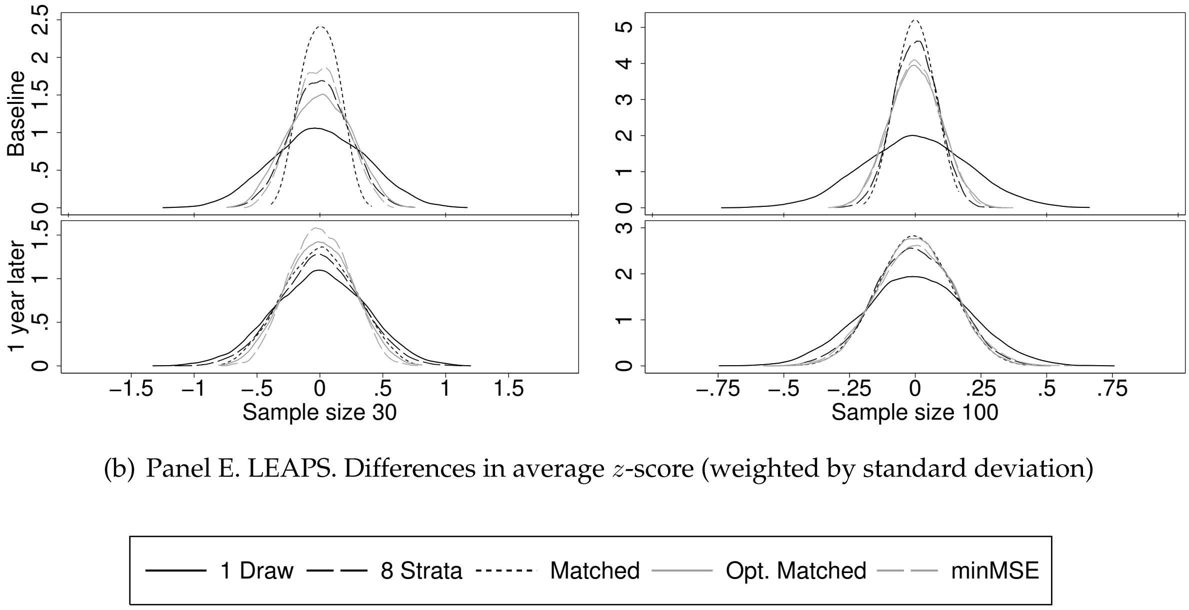

(b) Panel E. LEAPS. Differences in average z-score (weighted by standard deviation)

Note: Panel D and E show the distributions of the differences in treatment group means in standard deviations of the indicated variable for 10,000 treatment assignments. The higher the mass around a difference of 0, the better the treatment assignment with respect to balancing group means. Baseline variable is the outcome of interest; one of up to six variables used for treatment assignment. Follow-up variable is the same variable, measured six months after the baseline variable was collected; note that it is not included in treatment assignment.

FIGURE C.3: Distributions of the Differences in Group Means Between the Treatment and the Control Group in the Baseline Variable and the Follow-up Variable (N=30)

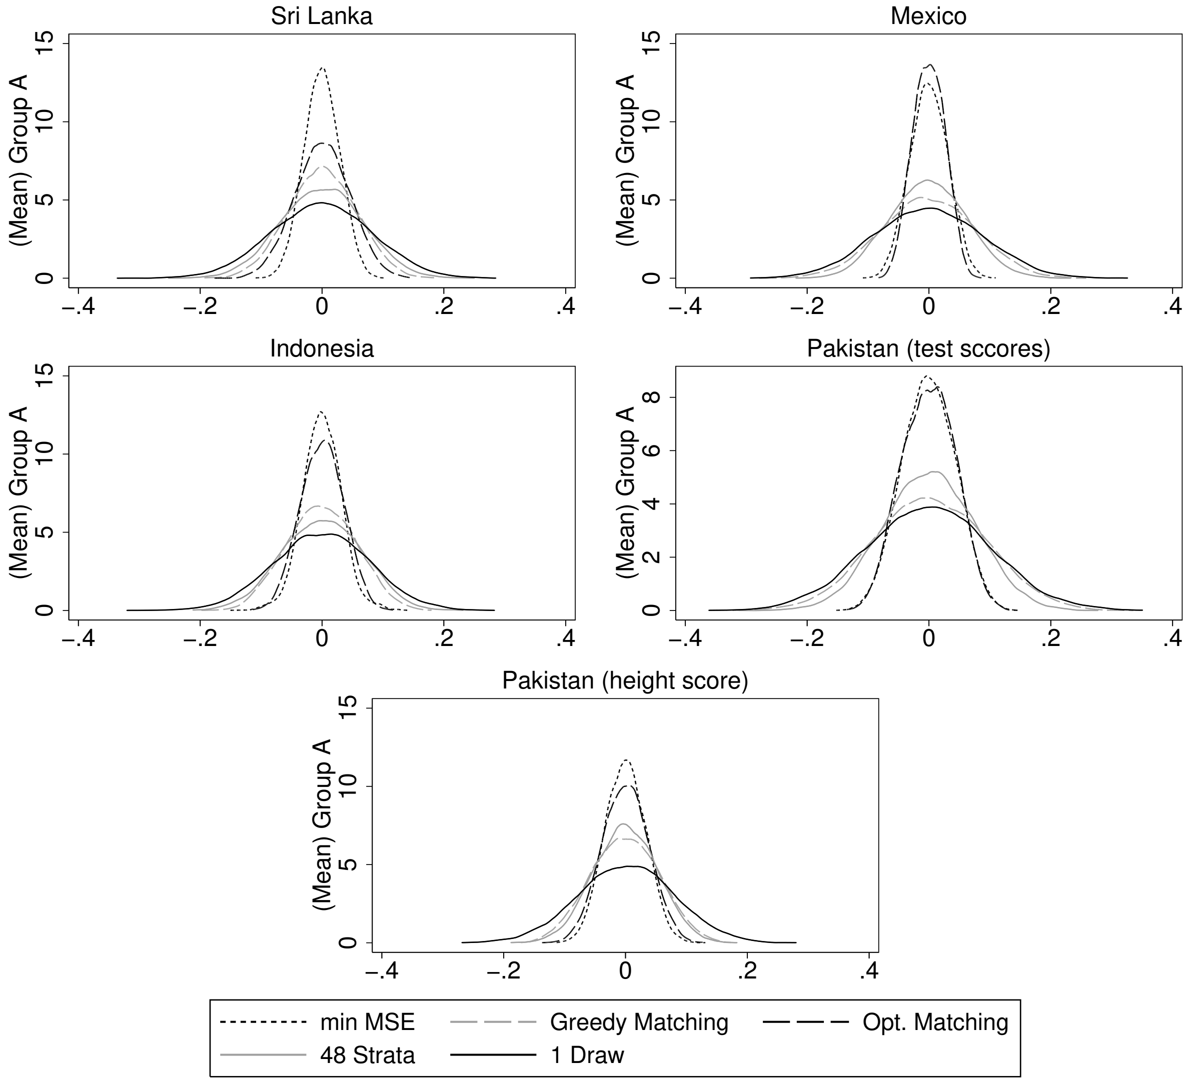

Note: Distributions of the (mean) differences in treatment group means among the group of variables to consider for treatment assignment are based on 10,000 treatment assignments. Differences in group means are expressed in standard deviations. A high mass around a difference of 0 indicates a good balance with respect to equality of group means.

FIGURE C.4: Distributions of the Mean Differences in Group Means (N=100)

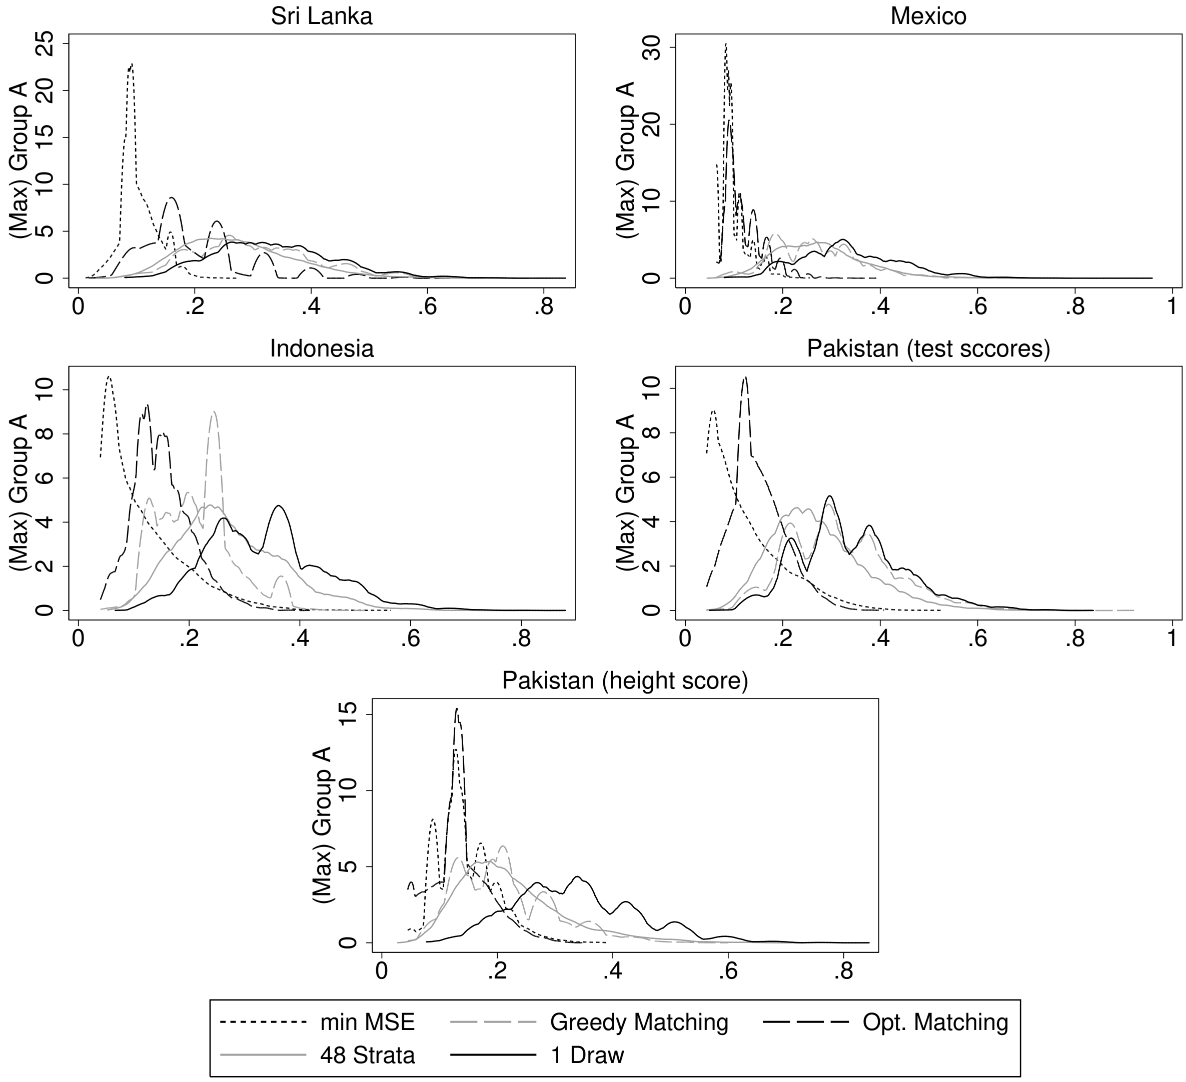

Note: Distributions of the (maximal) differences in treatment group means among the group of variables to consider for treatment assignment are based on 10,000 treatment assignments. Differences in group means are expressed in standard deviations. A high mass around a difference of 0 indicates a good balance with respect to equality of group means.

FIGURE C.5: Distributions of the Maximal Difference in Group Means (N=100)

Table C.4: Comparison of Treatment Assignment Methods Regarding Balance in Baseline Variables in Case of Attrition (Assessing Balance Using T-tests, N=30)

(a) Proportion p-values <0.1 when testing the difference in baseline group means

	Single random draw	Pairwise greedy matching	'Optimal matching'	Min MSE procedure	Stratified on two variables	Stratified on eight variables
No attrition	0.1131	0.0652	0.0045	0.0058	0.0678	0.0562
1 (2) unit removed	0.1123	0.0597	0.0047	0.0061	0.0665	0.0545
3 (6) units removed	0.0909	0.0642	0.0035	0.0069	0.0640	0.0571
5 (10) units removed	0.0926	0.0671	0.0048	0.0089	0.0635	0.0580
7 (14) units removed	0.0918	0.0686	0.0051	0.0131	0.0644	0.0606
Total	0.1001	0.0650	0.0045	0.0082	0.0652	0.0573

Note: Statistics based on 10,000 iterations. The smaller the proportion of p-values $< .1$, the better the balance with respect to similar group means. Details on the study and the computation of each measure are explained in section 4.5.2.

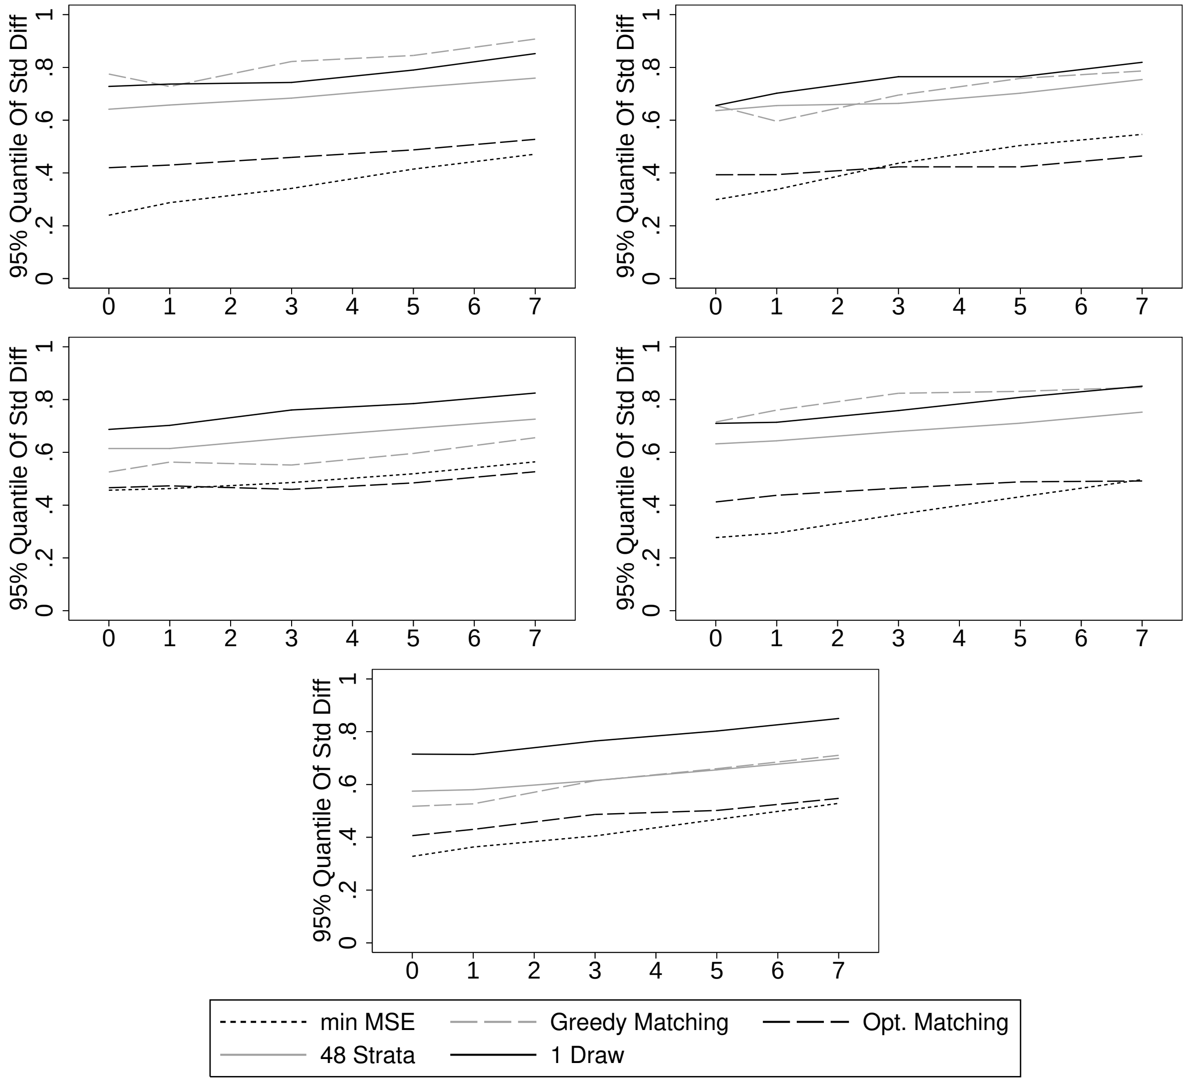

Note: Attrition. Evolution of the largest 95% quantile of differences in group means between treatment groups among the group of variables to consider for treatment assignment. X-axis: Number of units that are randomly removed from the study. Distributions are based on 10,000 treatment assignments. Differences in group means are expressed in standard deviations. A lower (maximal) 95% quantile of differences in group means indicates a better balance with respect to equality of group means.

FIGURE C.6: Simulated Attrition: Evolution of the Largest 95% Quantile of the Differences in Group Means Among The Group of Baseline Variables, N=30

Bibliography

Abadie, A., and G. W. Imbens. 2006. "Large sample properties of matching estimators for average treatment effects." *Econometrica* 74 (1): 235–267.

Abdellaoui, M. 2000. "Parameter-free elicitation of utility and probability weighting functions." *Management Science* 46 (11): 1497–1512.

Abdellaoui, M., C. Barrios, and P. P. Wakker. 2007. "Reconciling introspective utility with revealed preference: Experimental arguments based on prospect theory." *Journal of Econometrics* 138 (1): 356–378.

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

Aguila, E., O. Attanasio, and X. Quintanilla. 2010. "Cobertura del sistema privado de capitalización en Chile, Colombia y México." Working Paper WR-642. RAND Corporation.

Aiyagari, S. R. 1994. "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics* 109 (3): 659–684.

Aldrin, M. 2006. "Improved predictions penalizing both slope and curvature in additive models." *Computational Statistics & Data Analysis* 50 (2): 267–284.

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

Andrabi, T., J. Das, and A. I. Khwaja. 2015. "Delivering education: A pragmatic framework for improving education in low-income countries." Chap. 6 in *Handbook of International Development and Education,* edited by P. Dixon, S. Humble, and C. Counihan, 75:85–130. Cheltenham, UK: Edward Elgar Publishing.

Arlot, S., and A. Celisse. 2010. "A survey of cross-validation procedures for model selection." *Statistics Surveys* 4:40–79.

Banerjee, A. V., and E. Duflo. 2010. "Giving credit where it is due." *Journal of Economic Perspectives* 24 (3): 61–80.

Belisle, C. J. P. 1992. "Convergence theorems for a class of simulated annealing algorithms on $\mathbb{R}^d$." *Journal of Applied Probability* 29 (4): 885–895.

Bertsimas, D., M. Johnson, and N. Kallus. 2015. "The power of optimization over randomization in designing experiments involving small samples." *Operations Research* 63 (4): 868–876.

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.

———. 2002. "Loss aversion and scale compatibility in two-attribute trade-offs." *Journal of Mathematical Psychology* 46 (3): 315–337.

Bollaerts, K., P. H. C. Eilers, and I. Mechelen. 2006. "Simple and multiple P-splines regression with shape constraints." *British Journal of Mathematical and Statistical Psychology* 59 (2): 451–469.

Booij, A. S., and G. van de Kuilen. 2009. "A parameter-free analysis of the utility of money for the general population under prospect theory." *Journal of Economic Psychology* 30 (4): 651–666.

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

Bruhn, M., and D. McKenzie. 2009. "In pursuit of balance: Randomization in practice in development field experiments." *American Economic Journal: Applied Economics* 1 (4): 200–232.

Camerer, C. F., and T.-H. Ho. 1994. "Violations of the betweenness axiom and nonlinearity in probability." *Journal of Risk and Uncertainty* 8 (2): 167–196.

Camerer, C. F., A. Dreber, E. Forsell, T.-H. Ho, J. Huber, M. Johannesson, M. Kirchler, J. Almenberg, A. Altmejd, T. Chan, et al. 2016. "Evaluating replicability of laboratory experiments in economics." *Science* 351 (6280): 1433–1436.

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

Cochran, W. G., and D. B. Rubin. 1973. "Controlling bias in observational studies: A review." *Sankhyā: The Indian Journal of Statistics, Series A* 35 (4): 417–446.

Cohen, J. 1992. "Statistical power analysis." *Current Directions in Psychological Science* 1 (3): 98–101.

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.

Cox, D. R. 1957. "The use of a concomitant variable in selecting an experimental design." *Biometrika* 44 (1/2): 150–158.

Crainich, D., and L. Eeckhoudt. 2008. "On the intensity of downside risk aversion." *Journal of Risk and Uncertainty* 36 (3): 267–276.

Crainich, D., L. Eeckhoudt, and A. Trannoy. 2013. "Even (mixed) risk lovers are prudent." *The American Economic Review* 103 (4): 1529–1535.

De Boor, C. 1987. *A practical guide to splines.* Applied Mathematical Sciences. Vol. 27. New York, NY: Springer.

De Mel, S., D. McKenzie, and C. Woodruff. 2008. "Returns to capital in microenterprises: Evidence from a field experiment." *The Quarterly Journal of Economics* 123 (4): 1329–1372.

Deaton, A. 2010. "Instruments, randomization, and learning about development." *Journal of Economic Literature* 48 (2): 424–455.

Deck, C., and H. Schlesinger. 2010. "Exploring higher order risk effects." *The Review of Economic Studies* 77 (4): 1403–1420.

Dercon, S. 1996. "Risk, crop choice, and savings: Evidence from Tanzania." *Economic Development and Cultural Change* 44 (3): 485–513.

———. 2010. "Risk, poverty, and human development: What do we know, what do we need to know?" Chap. 1 in *Risk, Shocks, and Human Development: On the Brink,* edited by R. Fuentes-Nieva and P. A. Seck, 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.

Dirección de Gestión de Conocimiento, Cámara de Comercio de Bogotá. 2014. *Tablero de indicadores de Bogotá y Cundinamarca 2014 No. 2.* http://hdl.handle.net/11520/8554. Accessed: 2017-07-06.

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

Ebert, S., and D. Wiesen. 2011. "Testing for prudence and skewness seeking." *Management Science* 57 (7): 1334–1349.

Eckel, C. C., and P. J. Grossman. 2002. "Sex differences and statistical stereotyping in attitudes toward financial risk." *Evolution and Human Behavior* 23 (4): 281–295.

Eeckhoudt, L., and H. Schlesinger. 2006. "Putting risk in its proper place." *American Economic Review* 96 (1): 280–289.

Eilers, P. H. C., and J. J. Goeman. 2004. "Enhancing scatterplots with smoothed densities." *Bioinformatics* 20 (5): 623–628.

Eilers, P. H. C., and B. D. Marx. 2010. "Splines, knots, and penalties." *Wiley Interdisciplinary Reviews: Computational Statistics* 2 (6): 637–653.

Eilers, P. H., and B. D. Marx. 1996. "Flexible smoothing with B-splines and penalties." *Statistical Science* 11 (2): 89–102.

Eilers, P. H., B. D. Marx, and M. Durbán. 2015. "Twenty years of P-splines." *SORT-Statistics and Operations Research Transactions* 39 (2): 149–186.

Etchart-Vincent, N. 2004. "Is probability weighting sensitive to the magnitude of consequences? An experimental investigation on losses." *Journal of Risk and Uncertainty* 28 (3): 217–235.

Fagereng, A., L. Guiso, and L. Pistaferri. 2017. "Firm-related risk and precautionary saving response." *American Economic Review: Papers & Proceedings* 107 (5): 393–397.

———. Forthcoming 2017. "Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk." *Review of Economics Studies.*

Falk, A., and J. J. Heckman. 2009. "Lab experiments are a major source of knowledge in the social sciences." *Science* 326 (5952): 535–538.

Fennema, H., and M. Van Assen. 1998. "Measuring the utility of losses by means of the tradeoff method." *Journal of Risk and Uncertainty* 17 (3): 277–296.

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.

Fishburn, P. C., and G. A. Kochenberger. 1979. "Two-piece von Neumann-Morgenstern utility functions." *Decision Sciences* 10 (4): 503–518.

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.

———. 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.

Fisher, R. A. 1935. *The design of experiments.* Edinburgh, Scotland: Oliver & Boyd.

Friedman, M. 1957. "The permanent income hypothesis." Chap. 3 in *A theory of the consumption function,* 20–37. Princeton, NJ: Princeton University Press.

Fuchs-Schündeln, N., and M. Schündeln. 2005. "Precautionary savings and self-selection: Evidence from the German reunification experiment." *The Quarterly Journal of Economics* 120 (3): 1085–1120.

Gilboa, I. 1989. "Expectation and variation in multi-period decisions." *Econometrica: Journal of the Econometric Society* 57 (5): 1153–1169.

Gonzalez, R., and G. Wu. 1999. "On the shape of the probability weighting function." *Cognitive Psychology* 38 (1): 129–166.

Greevy, R., B. Lu, J. H. Silber, and P. Rosenbaum. 2004. "Optimal multivariate matching before randomization." *Biostatistics* 5 (2): 263–275.

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

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

Hansen, B. B., and J. Bowers. 2008. "Covariate balance in simple, stratified and clustered comparative studies." *Statistical Science* 23 (2): 219–236.

Harrison, G. W., M. I. Lau, and M. B. Williams. 2002. "Estimating individual discount rates in Denmark: A field experiment." *The American Economic Review* 92 (5): 1606–1617.

Harrison, G. W., and J. A. List. 2004. "Field experiments." *Journal of Economic Literature* 42 (4): 1009–1055.

Heckman, J. J., and J. A. Smith. 1995. "Assessing the case for social experiments." *The Journal of Economic Perspectives* 9 (2): 85–110.

Hofmann, W., M. Friese, and A. Roefs. 2009. "Three ways to resist temptation: The independent contributions of executive attention, inhibitory control, and affect regulation to the impulse control of eating behavior." *Journal of Experimental Social Psychology* 45 (2): 431–435.

Holt, C. A., and S. K. Laury. 2002. "Risk aversion and incentive effects." *American Economic Review* 92 (5): 1644–1655.

———. 2005. "Risk aversion and incentive effects: New data without order effects." *American Economic Review* 95 (3): 902–904.

Ibanez, M., and S. Schneider. 2017. "Unemployment, precautionary savings, and loss aversion: An empirical test." Mimeo.

Imai, K., G. King, and C. Nall. 2009. "The essential role of pair matching in cluster-randomized experiments, with application to the Mexican universal health insurance evaluation." *Statistical Science* 24 (1): 29–53.

Imbens, G. W. 2004. "Nonparametric estimation of average treatment effects under exogeneity: A review." *The Review of Economics and Statistics* 86 (1): 4–29.

———. 2011. "Experimental design for unit and cluster randomid trials." Technical Report. Harvard University.

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

Kallus, N. Forthcoming (2017). "Optimal a priori balance in the design of controlled experiments." *Journal of the Royal Statistical Society: Series B (Statistical Methodology).*

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.

Kasy, M. 2016. "Why experimenters might not always want to randomize, and what they could do instead." *Political Analysis* 24 (3): 324–338.

Keenan, D. C., and A. Snow. 2002. "Greater downside risk aversion." *Journal of Risk and Uncertainty* 24 (3): 267–277.

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

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. 1983. "Optimization by simulated annealing." *Science* 220 (4598): 671–680.

Klar, N., and A. Donner. 1997. "The merits of matching in community intervention trials: A cautionary tale." *Statistics in Medicine* 16 (15): 1753–1764.

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

Kohavi, R. 1995. "A study of cross-validation and bootstrap for accuracy estimation and model selection." In *Proceedings of the 14th international joint conference on Artificial intelligence,* 2:1137–1143.

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

Lee, J.-J., and Y. Sawada. 2010. "Precautionary saving under liquidity constraints: Evidence from rural Pakistan." *Journal of Development Economics* 91 (1): 77–86.

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

Levitt, S. D., and J. A. List. 2009. "Field experiments in economics: The past, the present, and the future." *European Economic Review* 53 (1): 1–18.

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

Loewenstein, G. 1987. "Anticipation and the valuation of delayed consumption." *The Economic Journal* 97 (387): 666–684.

Loewenstein, G. F., and D. Prelec. 1993. "Preferences for sequences of outcomes." *Psychological review* 100 (1): 91.

Lu, B., R. Greevy, X. Xu, and C. Beck. 2011. "Optimal nonbipartite matching and its statistical applications." *The American Statistician* 65 (1): 21–30.

Maier, J., and M. Rüger. 2011. "Experimental Evidence on Higher Order Risk Preferences with Real Monetary Losses." Munich Discussion Paper 2010-33. Volkswirtschaftliche Fakultät, Ludwig-Maximilians-Universität München.

Marx, B. D., and P. H. C. Eilers. 2002. "Multivariate calibration stability: A comparison of methods." *Journal of Chemometrics* 16 (3): 129–140.

McPeak, J., S. Chantarat, and A. Mude. 2010. "Explaining index-based livestock insurance to pastoralists." *Agricultural Finance Review* 70 (3): 333–352.

Modica, S., and M. Scarsini. 2017. "A note on comparative downside risk aversion." *Journal of Economic Theory* 122 (2): 267–271.

Moffitt, T. E., L. Arseneault, D. Belsky, N. Dickson, R. J. Hancox, H. Harrington, R. Houts, et al. 2011. "A gradient of childhood self-control predicts health, wealth, and public safety." *Proceedings of the National Academy of Sciences* 108 (7): 2693–2698.

Morgan, K. L., and D. B. Rubin. 2012. "Rerandomization to improve covariate balance in experiments." *The Annals of Statistics* 40 (2): 1263–1282.

———. 2015. "Rerandomization to balance tiers of covariates." *Journal of the American Statistical Association* 110 (512): 1412–1421.

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.

Pennings, J. M. E., and A. Smidts. 2003. "The shape of utility functions and organizational behavior." *Management Science* 49 (9): 1251–1263.

Pratt, J. W. 1964. "Risk aversion in the small and in the large." *Econometrica* 32 (1/2): 122–136.

Reinsch, C. H. 1967. "Smoothing by spline functions." *Numerische Mathematik* 10 (3): 177–183.

Rosenbaum, P. R. 1989. "Optimal matching for observational studies." *Journal of the American Statistical Association* 84 (408): 1024–1032.

Rubin, D. B. 1973. "Matching to remove bias in observational studies." *Biometrics* 29 (1): 159–183.

———. 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology* 66 (5): 688–701.

———. 1977. "Assignment to treatment group on the basis of a covariate." *Journal of Educational Statistics* 2 (1): 1–26.

Runge, C. 1901. "Über empirische Funktionen und die Interpolation zwischen äquidistanten Ordinaten." *Zeitschrift für Mathematik und Physik* 46:224–243.

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.

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

Schunk, D., and C. Betsch. 2006. "Explaining heterogeneity in utility functions by individual differences in decision modes." *Journal of Economic Psychology* 27 (3): 386–401.

Secretaría de Planeación, Alcaldía Mayor de Bogotá. 2014. *Índices de Ciudad.* http://www.sdp.gov.co/PortalSDP/ciudadania/Publicaciones_SDP. Accessed: 2017-07-06.

Shao, J. 1993. "Linear model selection by cross-validation." *Journal of the American Statistical Association* 88 (422): 486–494.

Shea, J. 1995. "Myopia, liquidity constraints, and aggregate consumption: A simple test." *Journal of Money, Credit and Banking* 27 (3): 798–805.

Tarazona-Gomez, M. 2004. "Are individuals prudent? An experimental approach using lottery choices." Technical Report. EHESS Toulouse.

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.

Tversky, A., and D. Kahneman. 1992. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.

Al-Ubaydli, O., and J. A. List. 2015. "On the generalizability of experimental results in economics." Chap. 20 in *Handbook of Experimental Economic Methodology*, edited by G. Frechette and A. Schotter, 420–462. Oxford, UK: Oxford University Press.

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.

Viceisza, A. C. G. 2016. "Creating a lab in the field: Economics experiments for policymaking: Creating a lab in the field." *Journal of Economic Surveys* 30 (5): 835–854.

Wakai, K. 2008. "A model of utility smoothing." *Econometrica* 76 (1): 137–153.

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

Wakker, P., and D. Deneffe. 1996. "Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown." *Management Science* 42 (8): 1131–1150.

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

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

World Bank. 2013. *World development report 2014: Risk and opportunity—managing risk for development.* Washington, DC: World Bank.

———. 2014. *Global financial development report 2014: Financial inclusion.* Vol. 2. Washington, DC: World Bank.