

Preference learning for policy analysis

Submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Engineering and Public Policy

Cristóbal De La Maza

B.S. Industrial Engineering, Pontifical Catholic University of Chile

M.S., Engineering and Public Policy, Carnegie Mellon University

M.S., Machine Learning, Carnegie Mellon University

Carnegie Mellon University

Pittsburgh, PA

December, 2018

© *Cristóbal De La Maza, 2018*
All rights reserved

Acknowledgments

I acknowledge funding from Center for Climate and Energy Decision Making, Grant number SES-1463492, through a cooperative agreement between National Science Foundation and Carnegie Mellon University; National Science Foundation, Decision Risk and Management Science, Award number 1530479 to Cleotilde Gonzalez, from CONICYT/PFCHA/DOCTORADO BECAS CHILE/2014/72150553 and CMU EPP Department. The survey in Chapter 4, was carried out with the aid of a grant from the Latin American and Caribbean Environmental Economist Program (LACEEP). For their comments and suggestions, I am in debt with Emmanouil Konstantinidis, Kelly Klima, Jihoon Min, Matt Babcock, Hau-yu Wong, Fei Lu, Baruch Fischhoff, John Sterman, Granger Morgan, Cosma Shalizi, Luis Rizzi and Luis Cifuentes. An special thanks to Julian Arellana, Claudia Aravena, Orsolya Kovacs and Brian Sergi for their help with the survey instruments. To my dissertation committee Alex Davis (chair), Inês Azevedo (co-chair), Cleotilde Gonzalez and Aarti Singh, thanks for your guidance on this journey. Alex, thank you so much for being a mentor and a friend, you are sincerely a great example. Coty, thank you so much for being a guide both in research and in life. I am also deeply grateful with my children Diego, Maria Gracia and Gabriel; my parents Rodrigo and Eugenia; my godparents Isabel and Enrique; my parents in law Miguel, Ivon and Sara; my grandparents Mario, Marisa, Lorenzo and Rose; and with my friends Jorge Gomez, Sebastián Elgueta, Lynn Kaack, Rahul Ladhania, Daniela Mahan, Rolando Bettancourt, Nicole Olivieri, Abelino Jimenez, Pia Coppa, Christian Garcia, Soledad Diaz, Jose Fernando Prada, Claudia Mahave and Ron Samuels. An special mention to my wife Militza. Militza, this dissertation also belongs to you, thanks for all your love, patience and compassion. I love you with all my soul. May I be happy. May my family be happy. May all beings be happy.

Abstract

In policy analysis, individual preferences are used to measure welfare across the population. Nonetheless, individuals are heterogeneous in both what they want or their preference content, and whether they know what they want or their preference structure. Prior work typically restricts preference heterogeneity analysis to differences in preference content. This dissertation explores the intersection of public policy analysis with preference heterogeneity along these two dimensions. We present a general framework for analyzing and discovering preference content and structure from choice data. Our framework extends welfare measurement to fully account for preference heterogeneity and can help to better understand the welfare impacts of new policies for sub-populations. As heterogeneity in preference can be related to judgment structure, we first study how heterogeneity in preference content is related to heuristic judgment. We establish the relationship between judgment and choice for cumulative flood risks. Second, we propose a model that can directly determine differences in both preference content and structure for individual decision makers empirically using graph matching methods. Finally, we measure heterogeneity in preference content across sub-populations. We develop and test a method to uncover relevant sub-populations in a choice model automatically using machine learning tools. We illustrate the approach discovering relevant socioeconomic covariates in a recent and real decision facing the Chilean government about the environmental impacts of electricity generation. Our framework can help to design policy interventions tailored for the heterogeneous preferences of the public.

To Militza,
Diego, Maria Gracia and Gabriel

Contents

1	Introduction	1
2	Understanding cumulative risk perception from judgments and choices	5
2.1	Introduction	6
2.2	Study design	10
2.2.1	Judgment task	12
2.2.2	Choice task	13
2.3	Results	16
2.3.1	What heuristics do people use for cumulative risks judgments?	16
2.3.2	Can heuristics affect flooding protection choices?	20
2.4	Discussion	25
2.5	Conclusions	29
3	Learning preference structure from choices with graph matching	30
3.1	Introduction	31
3.2	Learning preference structure	33
3.2.1	Preference representation as graphs	33
3.2.2	Graph matching and dissimilarity estimation	36
3.2.3	Lower dimensional dissimilarity embedding	38
3.2.4	Clustering allocation and merging	39
3.2.5	Within-cluster modeling and prediction	40

3.3	Results	41
3.3.1	Overview	41
3.3.2	Transitivity in risky choices	43
3.3.3	Anomalies in risky choices	45
3.3.4	CO_2 mitigation choices	47
3.3.5	Classification of decision rules	49
3.4	Discussion	50
3.5	Conclusions	53
4	Welfare analysis using the sparse multinomial logit model	54
4.1	Introduction	55
4.2	Model description and parameter estimation	57
4.3	Simulation	60
4.4	Empirical analysis	64
4.4.1	Survey design	66
4.4.2	Modeling results	70
4.4.3	Willingness to pay	74
4.5	Discussion	77
4.6	Conclusions	80
5	Summary and conclusions	81
6	References	87
7	Appendix: Cumulative risk perception from judgments and choices	105
8	Appendix: Learning preference structure with graph matching	108
9	Appendix: Welfare analysis using the sparse multinomial logit model	112

List of Tables

2.1	Choice stimuli	14
2.2	Multinomial (MNL) and latent class (LC) logit models results	23
4.1	Sparse multinomial logit simulation results	63
4.2	Experimental design	68
4.3	Modeling results	72
4.4	Willingness to pay for hybrid adaptive sparse multinomial logit model	75
8.1	Problems used for the choice experiments	109
8.2	Problem ID in adjacency matrix.	110
8.3	Linear utility models transitivity task	110
8.4	Willingness to pay per cluster CO_2 task	111
8.5	Latent decision rule model anomalies task	111
8.6	Proportion anomalies per cluster	111
9.1	Summary statistics	114
9.2	Factor analysis qualitative scales	115
9.3	Environmental impacts per technology	116
9.4	Private costs per technology	116

List of Figures

1.1	Preference modeling scheme	3
2.1	Experimental design schematic	12
2.2	Proportion using different heuristics	17
2.3	Probability judgment densities by heuristic use	19
2.4	Proportion that choose the no insurance coverage option	21
2.5	Probability of choosing option with no insurance coverage	24
3.1	Method summary	33
3.2	Tournament graphs for four alternatives and adjacency matrices	35
3.3	Autoencoder summary	39
3.4	Choice sets examples and embeddings in two dimensions	42
3.5	Clustering results transitivity task	44
3.6	Clustering results anomalies task	46
3.7	Clustering results CO_2 task	48
3.8	Hierarchical clustering expected adjacency matrices	50
4.1	Regularization path and bias	62
4.2	Support recovery	62
4.3	Sample average willingness to pay values across simulations	65
4.4	Choice set example (original in Spanish)	69
4.5	Estimated coefficients ASML model	74

4.6	Average externalities per technology	77
7.1	Expected loss for each option and risk premium	106
7.2	Indifference curves for individuals and judgments	107
9.1	Correlation matrix survey	113

1

Introduction

Policy analysis is often concerned with the determinants of preference and choice. In policy analysis, individual preferences are used to measure welfare across the population. Thus, individual preferences are first modeled and later aggregated to evaluate different policy schemes and forecast their consequences for society. Nonetheless, individuals are heterogeneous in both what they want or their preference content, and whether they know what they want or their preference structure. Prior work typically restricts preference heterogeneity analysis to differences in preference content. This dissertation explores the intersection of public policy analysis with preference heterogeneity along these two dimensions. We present a general framework for analyzing and discovering preference content and structure from choice data. Our framework extends welfare measurement to fully account for preference heterogeneity and can help to better understand the welfare impacts of new policies for sub-populations. Our framework can help to design policy interventions tailored for the heterogeneous preferences of the public.

For over a century, the dominant paradigm for preference analysis has been a set of

axioms that are necessary and sufficient for behavior to be consistent with the maximization of a well-behaved utility function, an idea dating back to the nineteenth century theorist Jeremy Bentham [21]. This paradigm requires decision-makers to be able to consistently rank any set of alternatives that they come across, restricting preference structure to a total order [177]. Following this paradigm, discrete choice modeling has been the workhorse for preference learning, requiring a complete, monotonic and transitive preference map [166]. In sum, discrete choice models assume individuals map attributes that describe each alternative to a real-valued scale, then choose the alternative with the highest utility, up to noise. Even more, it is required that this preferences map is continuous. Continuity in preferences refers to the condition that individuals take into account all available information when making a choice in a compensatory process [166].

Humans are largely heterogeneous, with preferences that vary over time [237]. Decision-makers are also susceptible to subtle but inconsequential changes in how the alternatives are described or made available such as framing effects, context effects, or reference dependence [248, 122, 246, 23]. Therefore, other cyclic preference structures besides a total order have been observed both in lab and in the wild. Hereafter, we provide methods on how to address these other structures in policy analysis. Further, psychological research finds that the cognitive burden of selecting the best alternative is often impossible to overcome, and to cope with the complexities of choice, individuals use simple heuristics to make a decision [89]. For example, continuity often does not hold, with decision-makers using choice rules that are non-compensatory [178, 238], ignoring some information [178, 109].

From a practical standpoint, in a policy analysis, the modeler has to design the experimental paradigm, test if the axioms of the theory are fulfilled, choose the parameters to include in the model, the proper functional form and select the appropriate error structure. This task is often daunting and will force the use of oversimplified methods and models that can lead to biased parameter estimates, resulting in harmful distortions when used for decision-analysis

and policy evaluation. Figure 1.1 organizes the challenges faced by the preference modeler. First, decision-makers might not be impervious to changes in the elicitation procedure. For example, participants can fail to understand the information provided, changing their answer depending on how information is presented. The modeler can address this issue by controlling for those factors in the experimental design and later in the model specification. Next, choice data needs to comply with several requirements established by the choice theory to allow the use of modern discrete choice methods. When those requirements are not fulfilled, results can be seriously misleading. A preference modeler must be able to detect those inconsistencies and propose alternatives to address them. Further, the modeler needs to select the proper choice model specification ex-ante. Finally, statistical model selection is performed and out-of-sample predictive accuracy of the model is tested.




Figure 1.1: Preference modeling scheme

In this dissertation, we provide new methods for the first three topics exposed in Figure 1.1, with three case studies. As heterogeneity in preference can be related to judgment structure, we first studied how heterogeneity in preference content is related to heuristic judgment. In the second chapter, we study failures of procedural invariance associated with different

perceptions of cumulative risks depending on risk information, and its influence in insurance decisions. Catastrophic events, such as floods, earthquakes, hurricanes, and tsunamis are rare, yet the cumulative risk of an event occurring at least once over an extended time period can be quite substantial. We study how failures in cumulative risk judgments can bias individual choices and how to correct these potential failures, improving risk communications. We used flood risks as a most relevant application.

In the third chapter, we propose a model that can directly determine differences in both preference content and structure for individual decision makers empirically using graph matching methods. Our method can discover non-traditional (and traditional) preference structures empirically using the directed graph representation for a sequence of choices, constrained optimization graph matching methods, and kernel-based clustering algorithms. We are able to identify individuals with similar preferences in both content and structure from pairwise comparison data. Our method helps analysts discover differences in choice patterns without any prior assumptions. We apply the approach to identify intransitivity or discontinuity in preferences in new implementations of classic experiments as well as newer stated preference studies.

Finally, we measured heterogeneity in preference content across sub-populations. We developed and tested a method to uncover relevant parameters in a choice model automatically using machine learning tools, the sparse multinomial logit model. The approach is flexible enough to model a wide variety of phenomena while retaining a parametric form that is readily interpretable, making it useful for aiding individual and societal decisions. We illustrate the approach discovering relevant socioeconomic covariates with multinomial logit models in a recent and real decision facing the Chilean government about the environmental impacts of electricity generation. We believe enhancing models of preference can benefit policy-makers by helping them understand the welfare impacts of new policies, and help design policy interventions customized for the demands of the population.

*"A wealth of information creates
a poverty of attention"*

Herbert Simon, 1971

2

Understanding cumulative risk perception from judgments and choices

Catastrophic events, such as floods, earthquakes, hurricanes, and tsunamis are rare, yet the cumulative risk of each event occurring at least once over an extended time period can be substantial. In this work we assess the perception of cumulative flood risks, how those perceptions affect the choice of insurance, and whether perceptions and choices are influenced by cumulative risk information. We find that participants' cumulative risk judgments are well represented by a bimodal distribution, with a group that severely underestimates the risk and a group that moderately overestimates it. Individuals who underestimate cumulative risks make more risk-seeking choices compared to those that overestimate cumulative risks. Providing explicit cumulative risk information for relevant time periods, as opposed to annual probabilities, is an inexpensive and effective way to improve both the perception of cumulative risk and the choices people make to protect against that risk.

2.1 Introduction

Probabilistic information about a hazard, such as the harmful side effect of a medication [221], the likelihood of a flood [59], or the chance of a stroke [83], is typically provided in the form of annual percentage rates or base rates (e.g. in chances per year) [254]. This form of risk presentation focuses attention on the probability of suffering an event in each period, obscuring the fact that the probability of at least one event occurring over many time periods becomes large over repeated exposure [216]. In general, people have a poor understanding of how risks accumulate over time, with many perceiving no accumulation at all [70, 243]. A possible reason for this misperception is that people use simple heuristics to cope with the complex computations required to accurately assess cumulative risks [244]. Although under some circumstances these heuristics can lead to accurate risk assessments [89], they can also lead to an underestimation or overestimation of the risk, impacting individual choices [216].

There are two main streams of work on cumulative risks perception, using either judgments or choices. Many previous studies have examined judgments of cumulative risk [59, 216, 134, 163, 162, 215, 199, 147, 225, 226, 127, 223]. In most of these studies, participants are directly asked to compute cumulative risks over a specified number of events (time periods) when given information about the probability of each event occurring. For example, Doyle provided participants with annual probabilities of a flood (0.5%, 1%, 2%, 3%, 5%, or 8%) and asked them to estimate the probability of being hit by a flood at least once over a certain time period (either 1, 5, 10, 25, or 50 years) [59]. The total cumulative probability of suffering an event X at least once during T years is given by:

$$P(X \geq 1) = P\left(\bigcup_{t=1}^T X_t = 1\right) = \sum_{t=1}^T p \cdot (1-p)^{t-1} = 1 - (1-p)^T \quad (2.1)$$

where p is the annual probability of an event and t is an index for each time period. People do not generally compute those probabilities when asked about cumulative risks. Instead,

Juslin *et al.* proposed that people use one of two heuristics to assess cumulative risks [127]. The first is an additive heuristic where decision-makers sum up the probabilities over each time period t ($\sum_{t=1}^T p_t \cdot t$) [127], which is the same as a multiplicative rule $p \times T$ when p_t is constant over time [59, 199]. A linear model of this judgment heuristic would have an interaction of the base rate p and time period T , with slope one and intercept at zero [156]. Someone using an additive heuristic would give a cumulative risk assessment of 30% when faced with a 1% annual flooding risk for 30 years. The second is a mean heuristic, where decision-makers give cumulative probabilities that correspond to the average of all periods ($1/T \cdot \sum_{t=1}^T p_t$) [127], which is the same as a constant probability if the annual probability p_t is constant over time [59, 199]. Someone who uses a mean heuristic would give a cumulative risk assessment over 30 years of 1% when faced with a 1% annual flooding risk. For flood risks, Doyle finds that the dominant heuristic used by people is the additive heuristic, which can be fairly close to the actual cumulative risk computed using small probabilities or for few time periods [59].

Other studies have focused on choices, rather than judgments, to understand the relationship between heuristic use and cumulative risk perceptions [214, 15, 6, 46, 45, 47, 48, 229, 34]. Choice studies provide participants with a set of alternatives for them to choose among, and identify bias in cumulative risk judgments based on violations of expected utility theory [256]. Most choice studies have identified an underestimation of cumulative risk [15, 6, 46]. For example, Bar-Hillel asked subjects to choose between two lotteries, a simple lottery depending on a single draw from a random variable (e.g., roll of a die), and a compound lottery depending on the outcomes of multiple random variables in a sequence that represent cumulative risk (e.g., the roll of one die, and if successful, the roll of another) [15]. Although the probabilities and outcomes were formally equivalent in the simple and compound lotteries, participants generally preferred the simple lottery, suggesting an underestimation of the probability of winning, and violation of the axiom of expected utility theory that lotteries

over the same outcomes with formally equivalent probability distributions are equivalent (the reduction of compound lotteries axiom) [256, 15].

In this paper, we aim to understand how the heuristics people use when assessing cumulative risks affect long-term decisions. The relationship between perception of risk and choice should be strong according to normative theories like expected utility theory, where risks form the basis of probability assessments which are then multiplied by the utility of outcomes to form a measure of the desirability of an alternative [256]. However, prior evidence suggests that “risk perceptions have significant, albeit small, associations with both intentions and behavior” [202], where the use of choice rules other than expected utility would diminish the impact of risk judgments on choice. In this work, we use previously described judgment strategies as a system to label and classify judgments into groups and measure their effect on choices. We use the classification system from Juslin *et al.*, expecting that decision-makers use one of two heuristics to assess cumulative risk [127]: 1) a mean heuristic (i.e. the probability is constant over time periods) resulting in underestimation of risk, and 2) an additive heuristic (i.e. adding up the event probabilities over the time periods) resulting in overestimation of risk.

We use flooding risk as a familiar and concrete example. Flooding corresponds to one of the most devastating hazards in the U.S. [219], yet demand for flood insurance is woefully inadequate despite subsidized premiums [171]. Although the magnitude of its effect is yet to be fully understood, the tendency to disregard low probability risks (a possible failure to account for cumulative risks), could in part explain the “underinsurance problem” [42]. Hence, we aim to identify one possible mechanism of low demand for flood insurance. Further, we aim to identify how cumulative flood risks might be better communicated to those who may be exposed to the hazard. To do so, we test the effect of providing cumulative risk information directly.

We extend prior work by proposing two novel hypotheses. The first is that presenting

cumulative risk information will help those who use the mean heuristic recognize their faulty risk assessment. By providing decision-makers with the correct cumulative probability of an event for a 1% annual base rate, we expect those who use the mean heuristic to realize that their cumulative risk assessments can't possibly be correct. Yet, recognizing the time-dependence of cumulative risk does not mean decision-makers figure out what the dependence should be, only that there should be *some* dependence [69]. One likely outcome is that they switch from the use of a mean heuristic to the use of an additive heuristic. This switch as a result of information should generally improve cumulative risk assessments, but can lead to large inaccuracies for long time periods, where the additive heuristic fails (and when used blindly, can yield probabilities greater than one).

To test this hypothesis we randomize participants to one of two conditions, either an example of how annual risks accumulate over time, called “information” hereafter (“Please note that a 1% chance of flooding each year is the same as a 26% chance within 30 years”), or no cumulative risk information. This manipulation contrasts with probabilistic information usually provided by governmental agencies to the public in the form of the chance of flooding occurring per year [121], with only a few efforts offering cumulative risk information [121, 65]. Our manipulation, although heavy-handed, provides one additional point on the curve relating base rates and the time period of exposure (1% and 30-years) to cumulative risk [69]. The manipulation may help respondents think of that curve, providing insight into how risks accumulate over time. One alternative explanation is that instead of helping decision-makers deduce the correct relationship between base rates, time period, and cumulative risk, the single piece of information may lead them to abstract an alternative rule [164], such as the additive heuristic, which would lead to an overestimation of cumulative risk [127]. A second alternative is that the manipulation may provide a new anchor that respondents adjust depending on the base rate and time period, without invoking the additive heuristic or learning the true relationship [231, 164]. The anchoring explanation would lead to increasing inaccuracy of risk

assessments as they deviate from the 1% and 30 year time period, although those deviations would not be systematic overestimates, as in the case of the additive heuristic.

Our second hypothesis is that errors in cumulative risk assessments are reflected in choices to insure against flooding. We propose that choices over risky prospects first involve an assessment of the probability of each outcome, followed by integration of those probabilities with other information to yield a choice. This two-stage process can be understood as the editing phase and evaluation phase described in Prospect Theory [129], where a decision maker first edits risk information with a judgment function $\pi(p, T)$ that depends on base rate p and time period T [240], then later integrates that judgment with outcomes to yield a choice [240]. This leads to the prediction that participants who use a mean heuristic, who underestimate the risk, will make more risk-seeking choices, while those that use an additive heuristic, who overestimate the risk, will make more risk-averse choices.

To test these hypotheses we asked participants to make cumulative flood risk assessments directly, along with choices about insurance policies, imagining that they own a house in a flood-prone area. In our judgment task, participants are asked questions such as “imagine that the risk of a flood is $p\%$ each year. Please estimate the risk for a period of T years”. In the choice task, risk perception is inferred through a series of binary choices between a risky and a safe option, representing no insurance coverage and full insurance coverage respectively. We then compare the heuristics implied by their judgments and choices, as well as modeled choices as a function of the judgment heuristics they used. In the next section we provide more details regarding the design of the study, and the judgment and choice tasks.

2.2 Study design

Our survey consists of three sections: a judgment task, a choice task, and a general section with follow-up questions and socioeconomic information. A full version of the questionnaire and the survey data is available online (osf.io/uqp25). A randomly assigned group is provided

the following sentence (“information condition”) during both their judgment and choice tasks: “Please note that a 1% chance of flooding each year is the same as a 26% chance within 30 years”. All participants were asked to complete both judgment and choice tasks, with the order randomized between subjects, resulting in a 2×2 between-subjects design, with order of the tasks as the first factor, and cumulative risk information (either provided or not) as second factor. In Figure 2.1 we provide a schematic of our experimental design.

We recruited participants using Amazon Mechanical Turk (MTurk), a platform for recruiting that is comparable to laboratory experiments [52]. Inclusion criteria were age of at least 18 years, IP address in the U.S., and completion of more than 1,000 hits with an approval rate of 95% or higher. We provided a payment of \$0.5 per participant and a \$1 bonus payment if the participant answered three of four attention checks correctly. These questions were two choice sets with a dominated alternative, and two memory questions about the hypothetical scenario. Approximately 90% of the sample answered three questions correctly and were paid the bonus. Our main analysis includes all participants, and is robust to exclusion of those who failed the attention checks.

We administered the survey between December 18th and December 20th 2015, recruiting 997 participants. Participants were on average 37 years old (standard deviation of 11) and 49% were female. Participants had on average 15 years of education (standard deviation of 4) and 34% reported that they finished high school with grade A in mathematics. Although 21% stated they had experienced a flood, only a few had been physically injured in a flood (1%) or experienced financial loss (9%). Further, 7% of the sample had experience with online flood risk calculators and 6% had flood insurance.




Figure 2.1: Experimental design schematic. Participants were randomly assigned to the information or no-information condition. Then they were also randomly assigned to receive the judgment task first or the choice task first. We initially hypothesized that, if the choice task was focused only on a 30-year period, a judgment task for both a 20 and a 30-year period could provide a sense of inconsistency across the experiment. Hence, we randomly assigned participants in the no-information condition that received the judgment task first to make risk assessments about a 20 and 30-year period, or to make assessments only for a 30-year period. Judgments and choices were not significantly different in these two groups, so we grouped the responses.

2.2.1 Judgment task

In the judgment task we ask participants to assess the cumulative probability of a flood over periods of 20 or 30 years when the base occurrence rates are 1%, 2% or 4%. We selected risk levels based on the definition of high risk area in the U.S. [65]. A high-risk area is exposed to a chance of at least 1% of a catastrophic flood in any given year. The 1% benchmark is commonly used in brochures to communicate flood risks [121, 65]. Base rates of catastrophic flood events higher than 4% are uncommon and may have made the task implausible, so we

did not use numbers greater than 4% [71]. Specifically, participants were asked the following (below we show example for base rate of 1%):

“Imagine that you have just moved to a city that may experience floods. You plan to buy a house for \$300,000 on a 20 or a 30-year mortgage. If a flood hits your property, the house will be entirely destroyed. Suppose there is a 1% chance each year that your house will be destroyed by a flood. If you buy a house with a 20-year mortgage (30-year mortgage), what is the percent chance that a flood will destroy your house during that period?”

2.2.2 Choice task

For the choice task, participants were asked to make choices about different insurance options given varying cumulative flooding risks. Participants made 10 choices between two options, A and B, with one representing full insurance coverage and the other no insurance coverage, both for a fixed decision time period of 30 years. Two choices were dominated (1 and 10 in Table 2.1), where one of the alternatives was strictly better on all dimensions. For example, in choice 1, we offered full insurance coverage at zero cost in Option B, which should always be preferred over no insurance coverage in Option A. We randomized the order of options across participants. Participants were asked to make choices given the following hypothetical scenario:

“Imagine that you have just moved to a city that may experience floods. You plan to buy a house for \$300,000 on a 30-year mortgage. If a flood hits your property, the house will be entirely destroyed. To protect your house from the flood risk during your mortgage period, your bank offers you a 30-year insurance policy bundled into your mortgage payments. For an additional yearly fee, the insurance policy will pay you for some of the house’s value if it is destroyed by a flood.”

The options described the probability of a flood each year (with annual base rate probabilities of 1%, 2% or 4%), the monetary loss when a flood does (\$300,000) or does not occur (\$0), and the insurance premium that the participant needs to pay each year throughout the duration of a 30-year contract. Insurance premiums used in the study are designed to provide more than actuarially fair premiums, corresponding to a contract where total premiums are equal to expected losses. If the annual flood risk was 1%, the annual premium was either \$1,000 or \$1,500. If the annual flood risk was 2%, the premium was either \$1,500, \$3,000 or \$5,000. Finally, if the annual risk was 4% the premium was \$2,000, \$5,000 or \$8,000. Table 2.1 presents stimuli for the choice task. The first column presents choice set ID and flood annual risk in parenthesis, columns two to four present cumulative risks under different heuristics, columns five and six present insurance yearly premiums, and columns seven to ten present both expected loss and expected cost in each option.

Table 2.1: Choice stimuli. Each option is composed of the monetary loss at stake, the chance of a flood each year, and the insurance premium for a 30-year contract. Option A (no coverage) implies the loss of the \$300,000 house in case of a flood and a zero insurance premium. Option B (full coverage) implies no loss in case of a flood at the stated insurance premium. Additive heuristic judgments are truncated at 100% when necessary. Expected loss is computed with a 0% discount rate.

	Cumulative risk (% 30-year period)			Insurance premium (\$ per year)		A (no coverage) (\$ per 30-years)		B (full coverage) (\$ per 30-years)	
	Accurate	Mean	Additive	A	B	E(Loss)	E(Cost)	E(Loss)	E(Cost)
1 (1%)	26%	1%	30%	0	0	78K	0	0	0
2 (1%)	26%	1%	30%	0	1,000	78K	0	0	30K
3 (1%)	26%	1%	30%	0	1,500	78K	0	0	45K
4 (2%)	45%	2%	60%	0	1,500	135K	0	0	45K
5 (2%)	45%	2%	60%	0	3,000	135K	0	0	90K
6 (2%)	45%	2%	60%	0	5,000	135K	0	0	150K
7 (4%)	71%	4%	100%	0	2,000	213K	0	0	60K
8 (4%)	71%	4%	100%	0	5,000	213K	0	0	150K
9 (4%)	71%	4%	100%	0	8,000	213K	0	0	240K
10 (4%)	71%	4%	100%	0	2,000	213K	0	213K	60K

Each insurance contract corresponds to a screening device designed to separate individuals according to their risk perception [222, 191, 220], where premiums are selected to attract individuals either exposed to a high or a low risk [191]. In our case, risk exposure is the same across individuals, but cumulative risk judgments heuristics can potentially lead to different risk perceptions [220]. For example, a decision maker that faces a 2% chance of suffering a catastrophic flood (or equivalently a 45% chance during a 30-year mortgage) should prefer full insurance coverage at a premium of \$1,500 per year (an expected cost of \$45,000 in a 30-year period) over an expected loss of \$135,000 if no insurance coverage is purchased. The same applies if risk is perceived as higher. In contrast, if a decision maker follows a mean heuristic, expected loss is estimated only at \$6,000, and no insurance coverage should be the choice.

According to expected utility theory, when offered a fair insurance premium, a risk-averse individual with accurate judgment should prefer a full insurance coverage contract [252]. Moreover, risk-averse individuals should be willing to pay an additional amount or risk premium on top of the actuarially fair premium (e.g., the risk premium is \$74,000 for the example, if we assume a square-root constant relative risk aversion (CRRA) utility function). The same applies to individuals that overestimate the risk by using an additive heuristic. In contrast, if individuals severely underestimate cumulative risk by using a mean heuristic, they would consider themselves worse off buying insurance, instead preferring no insurance coverage. Thus, depending on the judgment heuristic participants use, individuals might choose to seek more or less coverage than expected utility maximization would prescribe. In Figures 7.1 and 7.2 in the appendix section we graphically illustrate the screening device and the effect of risk aversion.

Risk aversion is present when decision makers frame the purchase for insurance in terms of gains with respect to the worst-case scenario (a catastrophic flood) [258]. However, the purchase of insurance could be also be framed as a loss with respect to best-case scenario (no

flood) [258]. When choice is framed as a monetary loss, Prospect Theory predicts subjects are risk-seeking [129], which in our case would lead to preferences for no insurance over full coverage. This presents a challenge to the inference that only individuals using a mean heuristic (i.e. those that underestimate the risk) would prefer the no insurance coverage option.

Prospect Theory also predicts overweighting of low probabilities and underweighting of high probabilities in choice [129, 241]. Accordingly, a decision-maker following an additive heuristic, tending to overestimate cumulative risks, when combined with the Prospect Theory probability weighting function, should have that overestimation dampened in choice, and vice versa for those who use a mean heuristic. As a result, Prospect Theory preferences will tend to mitigate the effect of risk assessment heuristics. Finally, our analyses assume the insurance policy would be paid during the entire contract period. Many insurance contracts forfeit future payments after an event. If participants assumed insurance would not be paid for the entire 30 years, the expected value difference between no insurance coverage and full insurance coverage would be harder to distinguish.

2.3 Results

2.3.1 What heuristics do people use for cumulative risks judgments?

Based on their judgments, we classified participants into three groups. The first two were defined according to their similarity with an additive heuristic or a mean heuristic. A third group was identified as being accurate (i.e. close to the real cumulative risk value). The three groups were identified as follows: suppose π_i is the subjective cumulative probability [240] over time period T provided by participant i for an event with annual probability p . Then, if $\pi_i - (1 - (1 - p)^T) \geq 0$ and $|\pi_i - p \times T| < |\pi_i - (1 - (1 - p)^T)|$, the participant's cumulative

risk assessment π_i is closer to an additive heuristic ($p \times T$) than to accurate judgment, and we categorize the judgment as following an additive heuristic. If instead $|\pi_i - (1 - (1 - p)^T)| < 0$ and $|\pi_i - p| < |\pi_i - (1 - (1 - p)^T)|$, judgments are categorized as following a mean heuristic. Otherwise, judgments are categorized as accurate. In Figure 2.2, we show the proportion of participants in each group, for judgments with 1%, 2% and 4% base rates. Non-monotonic judgments (i.e. chance of a flood greater for a 20-year period than for a 30-year period) were excluded from the analysis (67 participants) [199].




Figure 2.2: Proportion using different heuristics and 90% confidence intervals. 20 subjects provided decimal responses, which we transformed to percentages. Me: Mean, Acc: Accurate, Add: additive.

We used a series of chi-squared tests of independence $\chi^2(k)$ with Bonferroni correction and $k = 1$ degrees of freedom to evaluate differences in pairs of proportions between conditions [260]. Contrary to previous findings [59], we observed the mean heuristic as the most prevalent strategy in the no-information condition for all annual base rates (all $\chi^2(1) \geq 140$, $p < 0.01$). Indeed, for all cases where no cumulative risk information is provided, approximately 70% of our participants used a mean heuristic. Moreover, the proportion using a mean heuristic was not statistically different for all annual base rates (all $\chi^2(1) \leq 2.87$, $p \geq 1$). When additional information was provided to participants, the proportion using a mean heuristic was significantly reduced (all $\chi^2(1) \geq 183$, $p < 0.01$). Surprisingly, even when the correct information was provided, the proportion of participants underestimating the risk was significantly greater than zero ($\chi^2(1) = 48$, $p < 0.01$). The use of the base rate as an anchor to provide risk judgments might be harder to overcome than initially expected.




Figure 2.3: Probability judgment densities by heuristic use. Kernel density estimates are presented separately for 1%, 2% and 4% base rates, and for a 30-year time period.

When the cumulative risk information was available, the majority of participants were accurate for the 1% and 2% base rates (74% and 68% respectively), but this was not true for the 4% base rate (36%), where the additive heuristic fails to give accurate risk assessments. Accuracy was significantly reduced as the base rates increased (all $\chi^2(1) \geq 3.4$, $p < 0.03$). Figure 2.3 shows kernel density plots of risk assessments for 1%, 2% and 4% base rates for a 30-year period. The left panel in Figure 2.3 corresponds to the no-information condition and the right panel corresponds to the information condition. For the no information condition, we found that judgments were well represented by a bimodal distribution, with a group that severely underestimates the risk and a group that moderately overestimates risk. Judgment distributions varied significantly between the no-information and information condition, where the additional information improved accuracy. We reject the null hypothesis of equal judgment distributions using a Kolmogorov-Smirnov test (all $D_{614,383} > 0.5$, $p < 0.01$ with Bonferroni correction), where u and v in test statistic $D_{u,v}$ are sample sizes for each empirical distribution. For the accurate group, participants likely use other judgment heuristics. For example, those in the information condition classified as with accurate judgment, seemed to be anchoring on the 26% value, then adjusting based on the base rate information. The mode for this group closely resembles a multiplicative factor of the anchor (26%) and the ratio between the base rates (e.g. $\frac{2\%}{1\%} \times 26\% = 2 \times 26\% = 52\%$ and $\frac{4\%}{1\%} \times 26\% = 4 \times 26\% = 104\%$).

2.3.2 Can heuristics affect flooding protection choices?

In Figure 2.4 we compare the proportion of participants that chose no insurance coverage according to their judgment heuristic. For the choice task, we classified participants into one of the three groups (mean heuristic, additive heuristic, accurate) by majority rule. Because there were three base rates, this meant that participants were assigned to the group that best fit two of the three judgments. As expected, as insurance gets more expensive, more participants

prefer no insurance coverage. Participants that follow a mean heuristic were more likely to prefer no insurance coverage. This conclusion holds across all insurance premium levels and base rates, with the exception of the stimuli of a \$1000 insurance premium combined with a 1% base rate (all $\chi^2(1) > 2.3, p < 0.05$). Individuals that underestimated cumulative risk in their judgments (i.e. those that use the mean heuristic) tended to be more risk-seeking than other individuals (i.e. they were more frequently not willing to pay for insurance coverage). The proportion of risky choices (i.e. no insurance coverage) for participants that followed an additive heuristic was not significantly different from the choices of those with accurate judgments (all $\chi^2(1) \leq 2.9, p \geq 0.09$).




Figure 2.4: Proportion that choose the no insurance coverage option with 90% Wald confidence interval. At the top of the plot, we show the insurance yearly fee. Me: Mean, Acc: Accurate, Add: additive.

In order to better understand our participants' choices, we used multinomial (MNL) and latent class (LC) logit models [235, 98]. In all models, as we assumed utility is a linear function, the probability of choosing no insurance coverage depends on the subjective expected value in the no insurance and insurance alternatives $P_A = \frac{e^{SEV_A}}{e^{SEV_A} + e^{SEV_B}}$ [8]. Weighted additive approximations for the subjective expected utility of each option are: $SEV_A \approx (\alpha + \beta Mean) \cdot p + \sum_{l=1}^L \gamma_l S_l + \sum_{k=1}^K \delta_k D_k + A$ and $SEV_B \approx \rho \cdot y$, where p is the probability of a flood and y is the insurance premium each year. Here S_l are L socioeconomic variables, D_k are K dummy variables controlling for experimental factors such as order of the task (*Order*) or treatment condition (*Condition*), and the sub-index indicates the alternative (A or B). Notice that in alternative A, monetary loss in case of a flood is invariant (\$300K) and hence omitted from the equation. Likewise, in alternative B full insurance coverage limits the loss to zero and hence expected utility is unaffected by the risk. *Mean* is a dummy variable accounting for the use of a mean heuristic as an interaction term with base rate p , and the constant term A measures an unexplained tendency to prefer no insurance coverage. A , γ_l , δ_k , α , β and ρ are coefficients to be estimated. A latent class logit models allows inferring preference heterogeneity, and hence different coefficients, in different sample segments assuming a finite mixture of q classes. An LC model where $q = 1$ is equivalent to a multinomial logit model. Table 2.2 provides the results of those regressions.

Model MNL 1 in Table 2.2 shows participants are sensitive to increments in risk (base rate p), where a negative coefficients shows that increasing base rates will reduce the attractiveness of no insurance coverage. If we disregard treatment effects and the constant term, MNL 1 model predicts that using a mean heuristic will make preference for no insurance coverage less affected by base rates with $SEV_A \propto -14.9 \times p$, compared to an additive heuristic where $SEV_A \propto -33.8 \times p$. Figure 2.5 presents logit probabilities. In the first panel in Figure 2.5, we observe that individuals using a mean heuristic (in a red solid line) are more likely to prefer the risky option than individuals using an additive heuristic (in a black dashed line).

Table 2.2: Multinomial (MNL) and latent class (LC) logit models for probability of no coverage.

	<i>Dependent variable:</i>			
	P(no coverage)			
	(MNL 1)	(MNL 2)	(LC:Class 1)	(LC:Class 2)
Base rate (p)	-33.8*** (3.1)	-34.1*** (3.1)	-52.8*** (6.8)	-47.1*** (6.0)
Insurance premium (y)	-3.4×10^{-4} *** (1.5×10^{-5})	-3.4×10^{-4} *** (1.5×10^{-5})	-5.8×10^{-4} *** (3.2×10^{-5})	-7.2×10^{-4} *** (5.6×10^{-5})
Constant (A)	-1.0*** (0.1)	-1.1*** (0.2)	-2.8*** (0.2)	0.2 (0.2)
Base rate (p):Mean	18.9*** (2.0)	19.3*** (2.0)	8.3** (3.9)	18.7*** (5.8)
Condition (No information)	-0.4*** (0.1)	-0.4*** (0.1)	-0.7*** (0.2)	-0.5*** (0.2)
Order (Judgment first)	0.2** (0.1)	0.2** (0.1)	-0.1 (0.2)	0.1 (0.2)
Condition:Order	4.4×10^{-3} (0.1)	0.03 (0.1)	0.9*** (0.3)	0.3 (0.3)
Age (years)		-0.001 (0.002)		
Gender (female)		0.02 (0.1)		
Education (years)		0.002 (0.01)		
Math grade in high school (A)		0.2*** (0.1)		
Experienced a flood		0.02 (0.1)		
Injured in a flood		-0.6* (0.3)		
Experienced financial loss in a flood		-0.03 (0.1)		
Flood insurance		-0.3** (0.1)		
Constant (Class 2)				-0.6*** (0.03)
Observations	7,658	7,650		7,658
Log Likelihood	-4,608	-4,592		-3,507
BIC	9,218	9,188		7,018

*p<0.1; **p<0.05; ***p<0.01

Further, a higher insurance premium y , will make SEV_B lower, where a negative coefficient indicates that full insurance coverage becomes less attractive as the premium increases [232].

Also in the first panel in Figure 2.5, increasing insurance premiums from \$2,000, to \$5,000

and \$8,000 per year systematically increases the chance of choosing no insurance coverage. Further, model MNL 2 in Table 2.2 shows participants that had been injured in a flood or that currently hold flood insurance, with a negative coefficient, had a lower probability of choosing no insurance coverage. Strikingly, participants with higher numeracy, with a positive coefficient, had a higher probability of preferring no insurance coverage. A logistic regression model showed that among several covariates, only numeracy skill was significantly related with a lower probability of risk underestimation in the judgment task ($p < 0.1$).




Figure 2.5: Probability of choosing option with no insurance coverage with 90% parametric bootstrapped confidence interval with 10,000 draws from MNL model and a LC model.

Finally, our latent class model in Figure 2.5 uncovers two classes that may be interpreted as a framing process, where Class 1 individuals (35%) are risk-averse with a negative constant term A favoring preference for insurance, and Class 2 individuals (65%) are risk-seeking with

a positive constant term A (although not significant) favoring preference for no insurance coverage. Class 1 individuals are unaffected by judgment heuristics as evidenced in Figure 2.5 by red and black dashed lines being very close to each other in the middle panel.

2.4 Discussion

In this study we examined the heuristics people use when assessing cumulative flooding risks, the effect of those assessments on their choices, and a simple method of improving those judgments and choices. Given prior work showing that respondents use either a mean heuristic that leads to underestimation of cumulative risk or an additive heuristic that overestimates cumulative risk, we hypothesized that respondents that use the mean heuristic will make risk-seeking insurance choices, while those that use the additive heuristic will make risk-averse insurance choices. Furthermore, we hypothesized that providing a simple piece of information with the correct cumulative risk for one scenario (1% and 30-years) would lead to less reliance on the mean heuristic overall, and as a result, more accurate cumulative risk perception, and more risk averse choices.

To test our hypotheses, MTurk respondents completed an online survey experiment with both judgment and choice tasks. In line with prior studies, we found that cumulative risk judgments could be represented by a bimodal distribution, with a group that severely underestimates the risk and a group that moderately overestimates the risk. Contrary to previous findings, where the majority of individuals used an additive heuristic and overestimated the risk [59], we found that the majority of individuals used a mean heuristic to assess the cumulative probability of a catastrophic flood, leading to systematic underestimation of the risk [127]. For comparison, Doyle presented flood base rate risks as both the percentage chance per year and the frequency of occurrence (number of houses affected per 1,000 units), finding overestimation [59]. It is matter of future investigation to discover if the dominant

judgment heuristic depends on the metric used to present risks (such as relative frequencies versus probabilities).

Individuals that used a mean heuristic were also less inclined to pay for insurance, inappropriately taking their chances with the cumulative flood risks that they underestimated. The simple provision of cumulative risk information (in addition to annual probabilities) substantially reduced that tendency. If applied more broadly, such a simple approach could significantly improve insurance decisions among homeowners not legally bound to pay for insurance [118]. Similar results might hold for other protection decisions that could be paid out on a long term loan, such as increasing house elevation, investing in wet or dry flood proofing, and building flood barriers. However, consistent with the finding that individuals generally fail to infer the normative principles of judgment and decision-making from examples [69], the benefit of providing the correct cumulative risk for a 1% annual probability was seriously diminished when decision-makers faced a 4% annual probability, suggesting that cumulative risks would have to be provided for all levels of risk, not just for a single annual probability.

As previously proposed [59, 199, 225, 226, 127], results from the judgment task suggest that participants are likely using an anchoring and adjustment heuristic [244]. We propose that participants first intuitively consider the base rate p as a plausible answer. If bias is not detected [128], minor adjustments occur through an addition process $p + \alpha$, leading to responses that are close to the base rate. On the other hand, if information with the correct cumulative risk probability is provided, bias is detected and participants search for a new heuristic [128]. A variant of an additive rule, $p \times T$ corresponds to the most likely available rule [223, 227, 242]. Interestingly, several participants using an additive rule truncated their answer to 100% when asked for cumulative risk judgments with 4% base rate and 30-year period [59, 199]. It seems that when this heuristic leads to unreasonable values, participants improve responses by an anchoring and subtraction process ($p \times T - \beta$), rather

than recognizing their flawed reasoning and adjusting the heuristic itself [199]. Alternatively, participants may have used the additional information (26% for a 1% annual base rate) as a new anchor, then when faced with a new base rate, tried to adjust that anchor in a seemingly reasonable way, by multiplying by the new base rate ($\frac{p}{1\%} \times 26\% + \gamma$, with $\gamma < 0$ when the heuristic apparently yields unreasonable results). An open question is how many points on a cumulative risk curve a decision-maker requires before that curve can be reliably learned [164].

Participants showed significant heterogeneity in insurance decisions even after accounting for the judgment heuristic they used. One explanation for this result is that participants may have framed the decision differently, perceiving the insurance option to be either a gain or loss relative to an unobserved individual-specific reference point [129]. If the reference point corresponds to the best-case scenario (no flood), paying for insurance is seen as a loss, leading to risk-seeking behavior [129, 258, 241], and preference for no insurance coverage over full insurance coverage. If the reference point corresponds to the worst-case scenario (a catastrophic flood), paying for insurance is seen as a gain [129, 241]. Although predicting framing processes at an individual level is difficult [67], our latent class model uncovers two hidden classes that may be interpreted as a framing process, where Class 1 individuals are risk-averse and Class 2 individuals are risk-seeking. Interestingly, Class 1 individuals are unaffected by judgment heuristics.

With the exception of numeracy, we did not find a significant effect of different covariates on cumulative risk estimation. However, the work presented here is not a comprehensive analysis of all the factors that may influence cumulative risk assessments. For example, Keller *et al.* showed that affect is related to flood risk perception by presenting photographs depicting houses during a flood, observing an increase in perceived risk [134]. Likewise, Knäuper *et al.* tested whether people underestimate the cumulative risk of HIV infection of potential attractive partners, showing an underestimation of risk associate with a motivational bias [137].

Other studies have also shown a direct relation between emotions and flood protective behavior [230, 265, 205]. These results are in line with the use of an “affect heuristic” (i.e. judgments are based on emotions) [217]. “Anticipatory emotions” such as “fear, worry and anxiety” [151], or “anticipated emotions” such as “regret, guilt or shame” [151] under a catastrophic scenario might also explain bias in cumulative risk judgments [147, 202, 151]. Strikingly, participants with higher numeracy had a higher probability of preferring no insurance coverage. Although high numeracy individuals are more likely to select the appropriate mathematical rule, they may extract more “affective meaning from numbers”, clouding the decision process [179].

Experience plays an important role in risk perceptions and decisions [263]. This is also true for flood hazards [38, 132]. Previous research has shown that past experience is a crucial factor influencing risk perceptions [206, 192, 120, 139, 133, 134]. This phenomenon is related with the availability heuristics [242], where perceived risk is intensified when past events can be easily remembered. Past experiences also influence mitigation behavior [101, 204, 138], where a history of flood-related damage can heavily impact flood insurance demand [17, 266, 28, 146]. However, previous studies have not addressed how experience affects cumulative risk perceptions. Our results show that experience with a flood is a significant covariate in the choice task, but not in the judgment task. Nonetheless, only participants that had been injured in a flood were more likely to prefer full insurance coverage. It may be that the severity of the consequences of past flood events is what truly molds preferences for preventive measures [257, 228, 101, 205]. A possible explanation is that an event such as a personal injury may be easier to retrieve from memory, heightening risk perception [134, 131]. Further, theories of experiential based choice found that the frequency and recency of an event are the most relevant determinants of both risk judgments and choice [95, 115]. This is true for flood events [16, 39] and could potentially impact cumulative risk perceptions as well. Furthermore, our results may be domain specific [88]. It has been shown that psychometric risk-domain characteristics such as control or familiarity can affect risk appraisals [72, 262].

Future work should test if our conclusions are driven mainly due to the flood risk context or whether the errors apply more generally.

2.5 Conclusions

Respondents in our study demonstrated a poor understanding of how risks accumulate over time, with many perceiving no accumulation at all. The typical approach of providing information about the annual risk of an adverse event, such as a catastrophic flood, is unlikely to help, as decision-makers are unable to reliably transform that information into cumulative risks, resulting in judgments that underestimate the actual risk, and choices that fail to protect against that risk. Explicit cumulative risk information is an inexpensive and effective way to improve both the perception of cumulative risk (measured using judgment), and the choices people make to protect against that risk. When provided for the appropriate time period of actual decisions, such a strategy has the potential to improve both the public's perception of the cumulative risks of natural hazards, and their choices in the face of uncertain outcomes.

*"It is by logic that we prove, but
by intuition that we discover"*

Henri Poincaré, 1908

3

Learning preference structure from choices with graph matching

Approaches that elicit preferences from choices people make assume decision-makers know what they want. That is true if decision-makers can consistently order available alternatives, yielding transitive preferences, and are not susceptible to subtle changes in how alternatives are described. We leverage recent advances in graph matching and non-linear embeddings, to cluster decision-makers based on what they want or preference content, and whether they know what they want or preferences structure. Across three pairwise comparison experiments, including classic studies of risky choice and a two-attribute study about electricity generation portfolios, we are able to characterize heterogeneity of both the content and structure of preferences. Decision-makers frequently choose in a way consistent with utility maximization, yet some decision-makers make choices consistent with heuristic rules, while others appear to be uncertain about their preferences. As a generalization of traditional preference analysis, the approach can be used to make recommendations for those with consistent preferences, uncover complex choice rules, and suggest paths toward clarification for those who are uncertain.

3.1 Introduction

The relationship between preference and choice is one of the most important topics in Psychology [209, 239, 68, 40, 189], Economics [21, 177, 195, 3, 129, 250, 245, 167], and the Decision Sciences [211, 241, 247, 66]. The recent emergence of massive amounts of data on individual choices, from product purchases in online marketplaces, to voting in local or national elections, has led to the development of sophisticated statistical models that aim to determine the basic attributes that people use to make their choices. For example, many recommender systems use the similarity of individuals' past choices to make suggestions about products [91], the value of the statistical life is modeled using econometric estimates of the compensation individuals require in exchange for doing a job that has a higher risk of death [253, 5, 103], and votes are tallied under the assumption that each vote contains a well-defined expression of the voter's preference [187].

In each case, individual choice behavior reveals preferences that are consistent with utility maximization only if decision-makers can order the available alternatives [256, 11], and are not susceptible to subtle but inconsequential changes in how the alternatives are described or presented (framing effects, context effects, reference dependence) [246, 23]. If these conditions hold, it is possible to define a rank ordering of the alternatives according to the decision-maker's preferences, and there exists an ordinal utility function corresponding to that ranking.

Researchers in the decision sciences have found that, in many circumstances, preferences are not always well-behaved [246, 23]. One reason for these deviations is that the burden of selecting the best alternative among a large set, considering the potential costs and benefits of each alternative, is too difficult [66], forcing individuals to use short-cuts or heuristics to make their choices [210, 89, 178]. For example, one psychologically plausible way to deal with complex choices is to use only the most important attribute unless the alternatives are

psychologically indistinguishable on that attribute. Tversky’s lexicographic semiorder is such a process [239], and can lead to intransitive behavior.

In simple decisions between two alternatives, each with few attributes, cognitive overload is less likely to occur. However, decision-makers may still behave in a manner that is inconsistent with utility maximization if they are unsure about what they want. For example, a prospective homeowner may begin searching based on square footage, but, after touring a few homes, decide that the number of full bathrooms is the more important attribute. Such changes in decision rules, whether systematic or random, will lead to inconsistent choices and an inability to construct a proper ranking over alternatives. If that inconsistency arises from random fluctuations in preference, there is a substantial literature around stochastic transitive preferences that can be used to model choice data [160, 56]. If inconsistency arises only in the short-run, then giving decision-makers more time or more opportunities to choose will lead to stable preferences [40]. Yet not all choice inconsistencies can be characterized as random deviations from well-ordered preferences or failures to reach long-run stability [246, 23].

To separate decision-makers that know what they want from decision-makers that are uncertain or use choice heuristics, we develop a statistical approach that uses pairwise comparisons to cluster decision-makers based on the content and structure of their preferences. To do this we leverage recent advances in graph matching and non-linear embeddings. We first represent individual choices as preference graphs, then compute the distance between preference graphs for a sample of decision-makers, embed these distances into a lower dimensional space, and finally clusters decision-makers based on these embeddings. In what follows we describe the method and we apply the approach to uncover clusters of decision-makers with different preference content and structure across three experimental tasks.

3.2 Learning preference structure

In the following section we describe our approach. Figure 3.1 summarizes the four steps of our method for a simulated sample of 100 decision-makers. As aforementioned, first we construct preference graphs for each decision maker, then compute dissimilarities in both content and structure between all pairs of decision-makers. Next, we estimate a lower dimensional embedding for each dissimilarity matrix. Finally, we find clustering allocations and propose a decision rule for each cluster.




Figure 3.1: Method summary. The schema summarizes the four steps of our method for a simulated sample of 100 decision-makers. First, we represent choices as preference graphs. Next, we compute dissimilarities on both content and structure. Further, we estimate a lower dimensional embedding for each dissimilarity matrices. Finally, we find clustering allocations.

3.2.1 Preference representation as graphs

For step 1 in Figure 3.1, our basic unit of analysis is an individual decision-maker’s *preference graph* $G = (V, E)$, which consists of a set of vertices V and edges E where vertices represent

alternatives and edges represent binary *preference relations* between alternatives. For all pairs of alternatives a and b , one and only one of the following three preference relations holds [30]: i) if $a \succ b$, the decision maker strictly prefers a over b , then there is a directed edge $a \rightarrow b$ and not $b \rightarrow a$ in the graph (*strict preference* or aPb). ii) If $a \sim b$, the decision maker is indifferent between a and b , then a and b are connected by an undirected edge $a - b$ (*indifference* or aIb). iii) If a is incomparable with b , then no edge between a and b exists (*incomparability* or aJb). Although the preference graph representation is quite general, we focus on *tournaments* [174], where all alternatives are compared and all preference relations are strict. The four types of tournament structures possible for four alternatives are shown in Figures 3.2a, 3.2b, 3.2c, and 3.2d [55]. Preference graphs can also be represented in terms of their *adjacency matrices* A , where each cell A_{ij} in the matrix is a 1 if alternative i is strictly preferred to j , and 0 otherwise. Adjacency matrices are shown in Figures 3.2e, 3.2f, 3.2g, and 3.2h, with reflexive preferences (along the main diagonal) omitted.

Preference graphs can be oriented such that alternatives with a higher *score* are placed closer to the top, where the score for an alternative is the number of times it is preferred to each other alternative [174]. Arrows are omitted when going from top to bottom in the graph if transitivity holds, and curved upward arrows show transitivity violations. With four alternatives, the maximum score is 3 (an alternative that is preferred to all others), and the minimum is zero (an alternative preferred to no others). A score vector of $s = [3, 2, 1, 0]$ is a complete ranking of the alternatives, or a *chain*, shown in Figure 3.2i, and is consistent with classical utility maximization [249, 2]. In contrast, the lexicographic semiorder is a preference graph that can contain cycles [239], such as those shown in Figures 3.2j, 3.2k, and 3.2l, where the exact structure of the cycle is determined by the alternatives and their attributes.

The value of the preference graph approach is apparent when considering the decision analysis that an individual with each preference structure must undertake. Given a choice between any subset of four alternatives, a decision-maker with a chain provides a ranking




Figure 3.2: Tournament graphs for four alternatives. First row, preference graphs. Second row, adjacency matrices. Third row, unlabeled tournament structures over four alternatives can be defined by their score vectors: chain ($s = [3, 2, 1, 0]$) (i), cycle at bottom ($s = [3, 1, 1, 1]$) (j), cycle at top ($s = [2, 2, 2, 0]$) (k), long cycle ($s = [2, 2, 1, 1]$) (l). Adjacency matrices show ones in black and zeros in grey.

consistent with the global ranking over four alternatives. A decision-maker with a cycle at the top can consistently rank only the worst alternative, and likewise, the decision-maker with a cycle at the bottom can consistently rank only the best alternative. A decision-maker with the long cycle has a consistent ranking over any subset of alternatives, but no global ranking.

3.2.2 Graph matching and dissimilarity estimation

To complete step 2 in Figure 3.1, our primary analytical tool is a method of calculating the *distance* between preference graphs for both content and structure.

The content of preferences

A common distance metric between two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, is the minimum number of edges that need to be rearranged to make them equal, known as the *Hamming distance* $d_H(G_1, G_2) = \|\text{vec}(G_1) - \text{vec}(G_2)\|_1$ [102]. Decision-makers that have a small Hamming distance between their preference graphs tend to choose similar alternatives, or have similar *preference content*. For a sample of n individuals, the dissimilarity between all pairs of decision-makers can be represented in an $n \times n$ dissimilarity matrix D where D_{ij} contains the Hamming distance between the preference graphs of decision-maker i and decision-maker j . We use standard graph similarity tools to identify clusters of graphs with similar content, which in the case of ordinal multidimensional scaling, is equivalent to Coombs' multidimensional unfolding [49].

The structure of preferences

Preference structure cannot be obtained from these Hamming distance calculations. A chain graph $A \rightarrow B \rightarrow C$ has the same structure as the chain graph $C \rightarrow B \rightarrow A$ but their

Hamming distance is equal to the number of distinct pairs (3). To capture the notion of structure, we use a measure of the *structural distance* between two preference graphs, which will be zero if and only if two preference graphs are isomorphic [1], meaning there is a bijection $f : V_1 \rightarrow V_2$ such that the edges of all pairs of vertices $u, v \in V_1$ in G_1 have the same edges for $f(u), f(v) \in V_2$ in G_2 (and vice versa). The *automorphism group* $Aut(G)$ of a graph G contains all the graphs that are isomorphic to it [14], making it possible to test whether two graphs are isomorphic by determining whether their automorphism groups intersect. This is a well studied problem in computer science, called the *graph isomorphism problem* [14]. The minimum Hamming distance between two graphs across their automorphism groups then gives their *structural distance* d_S [41]: $d_S(G_1, G_2) = \min(d_H(Aut(G_1), Aut(G_2)))$. If two graphs are similar (but not isomorphic), their structural distance should be small.

With a few alternatives the structural distance between graphs can be quickly calculated using exhaustive search. However, the problem is NP-hard [1], requiring approximation techniques for large graphs with more than 8 alternatives. We recast the structural distance calculation as an *inexact graph matching problem* [150], where the objective is to find the permutation matrix P^* over the set of permutations that makes two adjacency matrices A_1 and A_2 as similar as possible. The objective function is [1, 150, 255]:

$$P^* = \underset{P \in \mathcal{P}}{\operatorname{argmin}} f(P) = \operatorname{dis}_{A_1 \rightarrow A_2}(P) = \|A_1 - P^T A_2 P\| \quad (3.1)$$

where A_1, A_2 are the adjacency matrices for the preference graphs of two decision-makers, and P is in the set of permutation matrices \mathcal{P} . If the squared Frobenius (L_2) norm is used, the problem is known as quadratic assignment (QAP) with non-deterministic polynomial time complexity [140]. Because the solution set \mathcal{P} is not convex, a common approach is to replace \mathcal{P} by its convex hull \mathcal{D} , the set of doubly stochastic matrices (all entries greater than equal to zero and each row and column sums to 1). This relaxation leads to a quadratic

program (QCV), solvable in polynomial time [149, 1]. Because this relaxation can lead to inaccurate results [1], we instead use Vogelstein’s approach (rGM) [255] that replaces the objective function $f(P)$ by the identity $-tr(A_1PA_2^TP^T)$, leading to a non-convex problem, where $\nabla^2f(P) = B \otimes A_1 + A_2^T \otimes A_1^T$ is not positive definite [255]. Vogelstein *et al.* proposed to solve this problem sequentially with Frank-Wolfe algorithm [78, 255]. We initialized the optimization with the QCV solution [157].

3.2.3 Lower dimensional dissimilarity embedding

For n decision-makers, the $n \times n$ matrix D_H of pairwise Hamming distances contains information about the content of decision-maker preferences, while the matrix D_S of pairwise structural distances carries information about their structure. Our approach aims to classify decision-makers into groups with similar preference content and structure simultaneously, so in step 3 in Figure 3.1, we first embed D_H and D_S into lower dimensional spaces with dimension $n \times d_1$ and $n \times d_2$, respectively, then concatenate the embeddings into an $n \times d$ matrix D that carries information about both the content and structure of preference, where $d = d_1 + d_2$. To construct the embeddings, we convert dissimilarities in D_H and D_S to values between zero and one using a radial basis kernel, with σ_H and σ_S fixed at the median of the respective dissimilarities [136, 130]. Next we train an autoencoder to embed each $n \times n$ kernel dissimilarity matrix into an $n \times d_1$ and $n \times d_2$ space [96], seeking to minimize the reconstruction error:

$$\min_{W,b,c} L(x) = - \sum_j x_j \log(\hat{x}_j) + (1 - x_j) \log(1 - \hat{x}_j)$$

The autoencoder encodes the input space x into a lower dimensional space $h(x)$ at its output layer, then reconstructs (decodes) the original input space as $\hat{x}(h)$ [96]. We used the non-linear sigmoid activation function for both the encoder $h(x) = Wx + b$ and the decoder

$\hat{x} = Wh + c$, where W is a matrix with weights and b and c are bias vectors. We used a 0.1 learning rate and 1,000 epochs. In Figure 3.3, we show the autoencoder network. We selected the number of dimensions in the embedding d using the elbow method, and we pretrained the autoencoder with a Restricted Boltzmann Machine [119, 259].




Figure 3.3: Autoencoder summary. Each original dissimilarity matrix is encoded into a lower dimensional space minimizing reconstruction error [259].

3.2.4 Clustering allocation and merging

In step 4 in Figure 3.1, we use clustering techniques on the $n \times d$ dissimilarity embedding matrix, with the main assumption being that decision-makers with small distances between each other indicate a common pattern of preference in a population of decision-makers, partially masked by noise. We use the k-medians algorithm to determine cluster allocation [212], solving the following optimization problem:

$$\min_{\mu, C} J(\gamma, \mu) = \sum_i^n \sum_j^k \gamma_{ij} \|x_i - \mu_j\|_1$$

We initialized the algorithm with centroids from a prior hierarchical k-means solution [106, 155, 9]. Here γ is a binary allocation matrix, k is the apriori defined number of clusters, C is the cluster allocation, and μ the vector with medians for each group. We used the gap-statistic to determine the number of clusters k [234]. If necessary, clusters are merged to provide a more general solution.

3.2.5 Within-cluster modeling and prediction

To understand the choice rules decision-makers use within each cluster, we take a simple modeling approach, allowing us to compare within-cluster behavior to prior work. We use the multinomial logit (MNL) model to approximate decision rules within each cluster. The MNL model assumes the probability that an individual in cluster q chooses alternative $i \in J$ is $P_{iq} = \frac{e^{V_{iq}}}{\sum_{j \in J} e^{V_{jq}}}$ where $V_{iq} = \sum_l \beta_{lq} \times x_l$ is a (usually linear) utility function [166], with x a vector of attributes and l an index for the elements of x . We tested both multi-attribute (compensatory) and single attribute (non-compensatory) utility functions, where other attributes are disregarded. We predict out-of-sample choices using a mixture of the within-cluster multinomial logit models, where the choice of alternative $i \in J$ has a probability $P_i = \sum_q \pi_q P_{iq}$, with π_q as the probability that an individual belongs to cluster q . In a purely predictive approach, where no information about a decision-maker m 's choices are available, predictions about the new decision-maker's behavior are simply the weighted average behavior of individuals within each cluster in the training sample, where cluster weights π_q are the in-sample proportion of individuals in each cluster. If T choices for the new decision-maker m are available, then we can place more weight on the clusters that are most consistent with the decision-maker's behavior using Bayes' Rule:

$$\pi_{q|T} = P(m \in q|T) = \frac{\prod_{t \in T} P_{tq} \times \pi_q}{\sum_q \prod_{t \in T} P_{tq} \times \pi_q}$$

Predictions about a new decision-maker's choices are also a weighted average, but where the weights are posterior probabilities $P_i|T = \sum_q \pi_{q|T} P_{iq}$ given the decision-maker's T choices.

3.3 Results

3.3.1 Overview

We collected choice data from Amazon Mechanical Turk (MTurk) workers for three stated preference tasks: 1) choices between two risky options based on a classic study by Tversky [239] (*transitivity task*), 2) choices between risky prospects similar to those used in a recent choice prediction competition [60] (*anomalies task*), and 3) choices between electricity generation options for one’s state that trade-off CO_2 emissions and electricity bill impacts [197] (*CO_2 task*). For each task, we recruited 200 MTurk participants, using inclusion criteria of: age of at least 18 years, IP address in the U.S. and completion of more than 100 hits with an approval rate of 95% or higher. A full version of the questionnaires and survey data is available online (osf.io/pf7jn). We provided a payment of \$1 per participant and a \$0.5 bonus if the participant answered an attention check correctly. The attention question was a choice set with a deterministically dominated alternative. Figure 3.4 shows an example of each choice tasks in the first column. For both the transitivity and the anomalies task probabilities are presented as pie charts. In the CO_2 task attributes are also presented graphically. The approach uncovered six, two, and seven clusters in the three tasks, respectively. The second column in Figure 3.4, presents the visualization of the dissimilarity embeddings and clustering using the t-Distributed Stochastic Neighbor algorithm (t-SNE) to project the embeddings to two dimensions [158]. Clusters are indicated using different colors and shapes, with Voronoi polygons used to show cluster separation.





Figure 3.4: Choice set examples in each task. Transitivity task: Choice set example alternative b vs. c. CO₂ task: Choice set example a vs. b. Anomalies task: Choice set example a vs. k. Visualization of embeddings in two dimensions with t-Distributed Stochastic Neighbor Embedding (t-SNE) [158]. Clusters are described with different numbers, colors and shapes. Voronoi polygons are plotted to show cluster separation.


3.3.2 Transitivity in risky choices

In the *transitivity task*, participants chose between the pairs of gambles included in Table 8.1(a) in the appendix section, from Tversky’s classic paper on intransitive preferences [239], along with five additional gambles. As shown in Figure 3.4(a), probabilities were presented as pie charts without numeric information. Participants were presented all pairs of alternatives (45 pairs in total), with three repetitions for each pair (in a randomized order), yielding a total of 135 choices per participant. Almost all participants (95%) were paid the bonus for passing the attention check. Our approach yielded six clusters: four with chain structures, one with a small cycle, and one with multiple cycles.


In Table 8.3 in the appendix section, we test three decision rules to explain decision-maker choice behavior in each cluster: 1) maximize expected value $V_{i1} = \beta_1 \times EV_i$, 2) maximize probability of winning $V_{i2} = \beta_2 \times P(\text{winning})_i$, 3) and maximize payoffs $V_{i3} = \beta_3 \times \text{Payoff}_i$. Figure 3.5a shows the expected adjacency matrices for the preference graphs in each cluster, where alternatives are arranged so a lower triangular adjacency matrix indicates choices based strictly on probabilities, and an upper triangular adjacency matrix indicates choices based strictly on payoffs. Significant heterogeneity can be seen in both the content of preferences (with most choosing based only on payoffs), and the structure (with clusters 1-4 showing clear chain structures, and clusters 5-6 with one or more intransitive cycles). Figure 3.5b shows the predicted probabilities from the logistic regressions that fit the data the best in each cluster, finding that for clusters 1-4, a decision rule based on a single attribute (either probabilities or payoffs) fit the data better than an expected value rule. Decision-makers in clusters 1, 2 and 3 preferred the alternative with a higher probability in 87%, 96% and 100% of choices (respectively). Decision-makers in cluster 4 almost always chose the alternative with the higher payoff (93% of the time). Although decision-makers in cluster 5 and 6 showed a cyclic structure, the proportion of choices favoring the option with the higher probability of



a) Expected adjacency matrix per cluster



b) Logit probabilities



c) Model accuracy

Figure 3.5: Clustering results transitivity task. First row, weighted expected adjacency matrix in each cluster for the transitivity task. We used a color scale to easy ease interpretation with adjacency matrices colored from one in darker tones and zeros in lighter tones. We also present moon graphs to explicitly differentiate preference structure. The proportion of the sample in each cluster is presented last. Second row on left, logit probabilities $P(A)$ of choosing the alternative with a higher probability of winning (A) per cluster. Second row on right, model accuracy on 1,000 bootstrapped samples as more choices are observed from participants. Observed choice sets are order according with their mutual observation with respect to a vector with the cluster assignments.

winning is significantly different from 50%, suggesting their choices were not entirely random. Figure 3.5c shows that using a mixture of multinomial logit models based on our clustering approach performs as well as a pooled multinomial logit model fit on all the data, when no choices for a decision-maker are observed. However, prediction accuracy rapidly improves for our mixture approach when just a few choices are observed, because those choices sort individuals into clusters with common preference content and structure. In sum, the majority of decision-makers in the sample used a single-attribute choice rule (clusters 1-4), simplifying the task, and leading to transitive preferences within-cluster. Decision-makers whose choices could not be easily explained by a single attribute were also more likely to have intransitive preferences (clusters 5 and 6).

3.3.3 Anomalies in risky choices

Over the past 50 years, one of the most important findings from the decision sciences is that decision-makers exhibit systematic deviations from behavior predicted by expected utility theory [32, 26, 24, 40, 189, 60]. In the *Anomalies task*, we selected five anomalies that formed the foundation for Prospect Theory [129, 241, 60]: 1) the certainty effect, 2) the reflection effect, 3) overweighting of rare events, 4) loss aversion and 5) risk aversion. In Table 8.1(b) in the appendix section, we present the alternatives in the experiments and the expected preference relation for all anomalies. We used 11 gambles resulting in 55 pairs with three replications per pair with their order randomized. The attention check was a choice between a lottery with a 50% chance of winning \$1,000 and 50% chance of winning \$500, against a lottery that offered \$450 for sure. Only 55% of the sample passed the attention check, suggesting they either did not understand the task, were not paying attention, or have more severe violations of expected utility theory than previously considered (dominance).

As shown in Figure 3.6, two groups emerged from our structure learning approach, both

with a chain structure. In Table 8.5 in the appendix section, we modeled decision-maker behavior within each cluster with four decision rules [60]: 1) maximize expected value and minimize variance [256, 144]; 2) maximize the probability of a better outcome [251]; 3) maximize a weighted additive function of outcomes [178]; and 4) maximize the probability of winning the high outcome [60]. Both clusters chose largely based on the probability of the better outcome [251]. Although participants in both clusters deviate from expected utility maximization, both clusters are transitive in expectation. Given that choice patterns in the two clusters are very similar, we do not pursue further analysis. At a first glance, it seems the anomalies emerged from the same process.





Figure 3.6: Clustering results anomalies task. In the first row, weighted expected adjacency matrix in each cluster for the anomalies task. We used a color scale to easy ease interpretation with adjacency matrices colored from one in darker tones and zeros in lighter tones. We also present moon graphs to explicitly differentiate preference structure.

3.3.4 CO_2 mitigation choices

Policy-focused researchers have used multi-attribute discrete choice models to estimate policy-relevant quantities, such as the market share of existing and new products [114, 97], substitution patterns [112], implicit discount rates [172], willingness-to-pay [165, 108], and consumer’s surplus [218, 264]. In the CO_2 task, we collected data based on a recent paper by Sergi *et al.* [197], who estimated willingness to pay for CO_2 emission reductions. In our extension of their experiment, participants are asked to make trade-offs between higher (or lower) impacts of electricity generation on climate change and a higher (or lower) electricity bill. As shown in Table 8.1(c) in the appendix section, participants were presented all pairs from 8 alternatives (28 pairs) with no repetitions. In this task, 97% of the 200 participants passed the attention check. Here our approach yielded seven clusters: five with chain structures, and two with multiple cycles.

To model behavior in each cluster, in Table 8.4 in the appendix section, we used a weighted additive linear utility model with no intercept over both attributes (bill and CO_2) as $V_i = \beta_{Bill} \cdot Bill_i + \beta_{CO_2} \cdot CO_{2i}$. Figure 3.7a shows the expected adjacency matrices per cluster. Alternatives are arranged so a lower triangular adjacency matrix indicates choices based strictly on electricity bill savings and an upper triangular adjacency matrix indicates choices based strictly on CO_2 . As we observe in Figure 3.7a, decision-makers tended to focus either on CO_2 , or the electricity bill. Decision-makers in cluster 1 chose strictly based on a lower electricity bill. Decision-makers in cluster 4 chose only based on lowering CO_2 emissions. Decision-makers in clusters 2, 3, 5 and 6 were willing to trade-off a higher bill for reductions in CO_2 emissions. Almost 30% of the sample has intransitive cycles in their preference structure in expectation, indicating some level of incoherence. Decision-makers in cluster 7 showed multiple cycles and are clearly uncertain about what they want. In Figure 3.7b we present coefficient values for the weighted additive linear utility model in



a) Expected adjacency matrix per cluster (Cl: Cluster)




Figure 3.7: Clustering results CO_2 task. First row, weighted expected adjacency matrix in each cluster for the CO_2 task. We used a color scale to ease interpretation with adjacency matrices colored from one in darker tones and zeros in lighter tones. We also present moon graphs to explicitly differentiate preference structure. The proportion of the sample in each cluster is presented last. Second row on left, coefficients for both attributes assuming a weighted additive linear utility model with no intercepts ($V_j = \beta_{bill} \cdot Bill - \beta_{CO_2} \cdot CO_2$). Given their similarities we merged clusters 2 and 3; and clusters 5 and 6. Second row on right, model accuracy on 1,000 bootstrapped samples as more choices are observed from participants. Choice sets are order according with their mutual observation with respect to a vector with cluster assignments.

each cluster for both bill and CO_2 . Clusters 1 and 7 are insensitive to changes in CO_2 with coefficients close to zero. Finally, Figure 3.5c shows again that a mixture of multinomial logit models performs better than a pooled multinomial logit model, when a few choices are used to assign cluster membership. In short, more than half the participants do not have well-behaved preferences and are either using simplifying choice heuristics based on a single attribute (clusters 1 and 4) or have intransitive preferences (clusters 3 and 7).

3.3.5 Classification of decision rules

To better understand the relationship between the clusters for each task, we use a hierarchical clustering approach [135]. As shown in Figure 3.8a for the transitivity task, the hierarchical clustering sorts decision-makers according to the primary attribute they used to make their decisions, with clusters 1, 2, 3, 5 and 6 deciding based on probabilities and cluster 4 deciding based on payoffs. Next, decision-makers varied on the degree to which they could discriminate between the probabilities, which were shown only in graphical form [25], where those in clusters 1, 2, and 3 had high discrimination, and those in 5 and 6 had low discrimination. Those with low discrimination also tended to have intransitive cycles in their preferences. For the CO_2 task, as shown in Figure 3.8b, individuals either were "Greens" (clusters 4, 5 and 6), focusing on CO_2 , or "Bills" (1, 2, 3 and 7), focusing on saving money. At the next level in the hierarchy, decision-makers in clusters 2, 3, 4, and 6 tended to use a compensatory decision rule (giving weight to both attributes), versus those in clusters 1, 4, and 7 who used a non-compensatory decision rule (using only either CO_2 or bill to make their choice). Finally, some individuals were uncertain about their preferences, as in clusters 3 and 7, who had intransitive cycles.




Figure 3.8: Hierarchical clustering on the expected adjacency matrices for each cluster in the transitivity and CO_2 tasks.

3.4 Discussion

In this paper we use a graph representation of preference to uncover heterogeneity in the content and structure of preferences across three samples of decision-makers. The approach first represents the choices of individual decision-makers as graphs, then computes the Hamming and structural distance between graphs for all pairs of decision-makers, embeds those distance matrices into a smaller dimensional spaces, clusters decision-makers based on their proximity in those spaces, and finally uncovers the underlying decision rule within each cluster. We explore the approach with two new empirical implementations of classic experiments in decisions under risk and a new policy-relevant stated preference task. In each experiment we exploit regularities in choice patterns to identify individuals using similar choice rules. In a classic experimental design by Tversky we find that the vast majority of the sample uses a single attribute (up to noise) to choose, undermining the plausibility of other more complex rules like expected value calculations. In a set of preference anomaly choice tasks, we find two groups that although deviate from expected utility maximization,

are transitive in expectation. Lastly, in a policy-relevant choice task that asked respondents to choose between savings on their electricity bill and CO_2 emissions, we find non-compensatory behavior in about 40% of the sample, with almost all of them (34%) unwilling to pay some cost to avoid climate change.

In the transitivity task, a large proportion of decision makers showed choice patterns consistent with a single attribute decision rule [32, 87, 213, 74]. Decisions based on simple rules save time and effort required in the task, making them an attractive approach [198, 178, 86]. In our experiment, 65% of the sample chose only based on the probability of winning a gamble, and 9% chose only based on the gamble's payoff [27, 145, 27, 25]. The data are consistent with a lexicographic order, where decision-makers use only one attribute unless there are exact ties [74], rather than a lexicographic semiorder which allows for inexact ties within a just-noticeable-difference [239]. None of the decision-makers made choices consistent with subjective expected value, which would be some function of the product of the probability and payoff attributes [256, 75]. Decision-makers in most clusters had transitive preferences [181], likely because a simplifying single-attribute lexicographic decision rule makes consistency (and transitivity) a foregone conclusion.

For the anomalies task, although as shown in Table 8.6 in the appendix section, decision-makers in our sample were loss and risk averse. We were not able to replicate overweighting of rare events, the certainty effect, or the reflection effect [129]. Even more troubling was the inability of almost half of our respondents to choose a dominated alternative, suggesting either severely irrational preferences or an inability to understand the task. Strikingly, even though decision-makers tended to be loss averse, their preferences were transitive in aggregate for both clusters, highlighting that although expected utility theory may not be descriptively valid, it may still be possible to construct an ordinal representation of decision-maker preferences.

Lastly, in the CO_2 task the majority of the sample (60%) used a compensatory decision rule and were willing to make trade-offs between economic costs and climate protection, but

a full 40% did not. A naive approach would fit a single multinomial logit model (MNL) with two attributes on the full sample, leading to distorted policy analysis [169]. For example, willingness to pay corresponds to the *marginal rate of substitution* (MRS) between an attribute k and the cost of each alternative $MRS_{kc} = \frac{\partial u_i}{\partial x_k} / \frac{\partial u_i}{\partial c_i}$. If a model was fitted assuming the utility of each alternative is linear in its attributes, the function would be $V_i = -8.3 \cdot \text{Bill} - 4.5 \cdot \text{CO}_2$ on the whole sample, giving a willingness to pay (WTP) of $WTP = 30 \cdot -4.5 / -8.3 = 16$ % increment in the monthly electricity bill for a 30% percent reduction in CO_2 emissions. That is, the population is willing to pay to avoid CO_2 emissions. This calculation assumes homogeneity in both the content and structure of preference [166]. A very different picture emerges from our preference clusters, where many are unwilling to make the trade-off implied by the marginal rate of substitution, or do not even have coherent preferences that could be characterized by a utility function. Analysis of willingness to pay in aggregate would imply trade-offs that much of the population is unwilling to make.

Our approach is able to separate decision-makers based on whether their choice patterns are consistent with a specific theory of decision-making, such as utility theory or a lexicographic order. It is also able to aid policy analysis, allowing subgroups with heterogeneous preference content and structure to express the trade-offs that they are willing (or not willing) to make. The approach can also improve predictive accuracy. In Figures 3.5 and 3.7, using 1,000 bootstrapped samples from the original observations for both the transitivity task and the CO_2 task, a mixture of multinomial logit models for each cluster with linear utility functions, with individual mixing probabilities conditional on the observed choices, yields a higher accuracy than a pooled multinomial logit model. Accuracy increases as more information is available to estimate cluster membership. Predictive accuracy increased around 10% when only a few choices were available to estimate mixing probabilities.

Finally, we highlight some limitations of our method. Clustering always has some arbitrariness. For example, the number of dimensions to embed the dissimilarity matrices

in a lower dimensional space was defined using the elbow method, a useful heuristic [96]. Allowing for the number of dimensions to be determined automatically by the data in the optimization process would be an important improvement [50]. The data requirements also present an important challenge. The number of pairwise comparisons required to complete a tournament grows quadratically with the number of alternatives, increasing the risk of decision-maker fatigue.

3.5 Conclusions

We present a general framework for analyzing and discovering preference content and structure from choices. The approach can suggest new theories to decision researchers, or confirm old ones, and lend strength to welfare analysis, or undermine it. Policy decisions ought to be based on trade-offs decision-makers would make between different private or public goods [12]. A major challenge faced by policy analysts is to identify decision makers that are not willing to make such trade-offs. Respondents in our studies showed heterogeneous patterns of choice, with a large proportion not willing to compromise. Our approach can identify those groups and uncover heterogeneity in preference structure without requiring any prior knowledge of those structures. Practitioners will be able to use this approach to classify decision-makers according to their preference content and structure. This can inform decision-makers themselves through decision analysis, as well as help policy-makers better understand the welfare impacts of new policies, and design policy interventions that meet the demands of the public.

"The purpose of computing
is insight, not numbers"

Richard Hamming, 1962

4

Welfare analysis using the sparse multinomial logit model

Variable selection approaches in econometric models are typically ad-hoc, do not select the correct variables asymptotically, or yield biased parameter estimates. In this work we address the variable selection problem for discrete choice models using a modification of the multinomial logit model. We show how to debias the modified multinomial logit model with an adaptive Lasso penalty, allowing asymptotically consistent variable selection and unbiased parameter estimates that are useful for welfare analysis. We demonstrate the approach with a real decision facing the Chilean government about what types of energy to produce, and where production should take place to minimize environmental externalities.

4.1 Introduction

Discrete choice models are widely used in economics and psychology to understand individual preference [235]. The multinomial logit model is a canonical example, in part because of its simplicity and ability to capture preference heterogeneity among measured subgroups of decision-makers. Variation across these subgroups on key metrics, such as willingness-to-pay, is of economic interest for both revealed preference studies and stated preference surveys [168]. Applications of discrete choice models to understand welfare impacts on subgroups include the valuation of insurance contracts [29, 203], the environmental and social quality of products [188, 85, 54, 194, 172], quality of life improvements from health care [33, 44], and the value of safety [104, 184, 185]. Subgroup analyses are also important for understanding public policy impacts, such as mortality risks reductions [113, 141], morbidity effects [148, 126, 125], urban noise [84], and environmental [197, 105] and landscape impacts [186].

Although researchers sometimes have specific hypotheses about subgroup behavior, there are often many more subgroups of interest than can be reasonably included in one statistical model. Traditionally, choice models have included individual socioeconomic characteristics such as age, income, gender, or marital status as interactions with observed attributes in the model specification [167]. Other individual traits such as perceptions, beliefs, motivations, and prior experience can also shape individual preferences [167]. For example, Madanat *et al.* [159] included attitudes and perceptions, together with other socioeconomic variables, in a discrete choice model of intentions to avoid traffic congestion, using the scores for different psychometric scales as interactions with attributes of alternatives.

In this work we use a data driven approach for subgroup analysis by modifying the multinomial logit model, called the *sparse multinomial logit*, allowing the most relevant effects to emerge directly from the statistical model. Such an approach can improve both the predictive power of the multinomial logit, as well as its interpretability, helping researchers

focus only on the subgroups that are best for forecasting and have the strongest effects. The approach can provide a check on theory-driven model specifications, help avoid researcher degrees of freedom issues that arise from repeated model specification [208], and prevent suboptimal model selection issues from insufficient specification search or use of simplifying heuristics to choose subgroups.

The desirable characteristics of a statistical inference approach are traditionally lack of bias, asymptotic consistency, and efficiency. Those characteristics typically say nothing about which variables should enter the model, only that if the right ones are included, these desirable properties hold. Approaches that conduct automatic covariate selection, on the other hand, should have the *oracle property*, where the model asymptotically selects only the variables that have truly non-zero coefficients (those parameters that are in the *support* of the true model) [61]. Common approaches to covariate selection, such as forward or backward variable selection, do not have this oracle property [233], nor does penalizing the model negative log-likelihood by the absolute magnitude of the regression coefficients, as in the Lasso (although such an approach does zero out some coefficients in a process known as *denoising* [107, 81, 233, 201]). A suitable modification of the latter penalization approach, called the *adaptive Lasso*, does have the oracle property [268].

While using the adaptive Lasso penalty combined with maximum likelihood estimation has the oracle property, the price of that property is paid in bias, where regression coefficients are too close to zero [80, 82]. Adding bias can actually reduce forecasting error [82], but in economic and policy analysis it is usually the coefficients themselves, not forecasts, that are of interest. For example, willingness-to-pay is a common metric used in economic and policy analysis, and in a linear utility model, can be expressed as the ratio of an estimated coefficient relative to the price coefficient [235]. Bias in the estimated coefficients would then lead to biased policy analysis. Although penalized models have gained interest in economics [19, 20, 63], and more recently in non-parametric discrete choice models [90, 236, 64], sparsity-induced

bias has limited their usefulness for welfare analysis, such as the valuation of public goods.

In this work we describe an approach that uses state-of-the-art methods to debias the coefficients that result from applying the adaptive Lasso penalty to the multinomial logit model, enabling them to be used for economic and policy analysis. We use the valuation of environmental public goods as the test case in a recent, and real, decision facing the Chilean government. The next section of the paper describes the model. Then we detail convex optimization techniques that allow model estimation and bias correction. We then use simulations and a survey to illustrate the value of the sparse multinomial logit model.

4.2 Model description and parameter estimation

Our objective is to develop an approach that automatically selects subgroups in a population that are heterogeneous in their preferences, but without introducing bias into estimated coefficients. We begin with the multinomial logit model, then describe the adaptive Lasso penalty, which has the oracle property in its automatic selection of those sub-populations. In the multinomial logit model, the utility of alternative i for individual n , $U_{ni} = V_{ni} + \epsilon_{ni}$, is defined as the sum of a deterministic part V_{ni} and a stochastic part ϵ_{ni} with a Type-I extreme value distribution [166]. The deterministic part of the utility $V_{ni} = V_{ni}(\mathbf{x}_{ni}, \mathbf{s}_{ni})$ depends on the attributes of each alternative \mathbf{x}_{ni} and interaction effects between those attributes and characteristics of a decision maker \mathbf{z}_n (i.e. $\mathbf{s}_{ni} = \mathbf{x}_{ni} \times \mathbf{z}_n$). The multinomial logit model assumes that the probability that an individual n chooses alternative i depends linearly on the deterministic utility assigned to each alternative $V_{ni} = \omega_i + \sum_j^p \alpha_j x_{nij} + \sum_j^p \sum_r^k \beta_{jr} x_{nij} z_{nr}$, where $\theta = \{\omega, \alpha, \beta\}$ is a vector of coefficients, ω_i is an alternative specific constant, p is the number of attributes in each alternative and k is the number of subgroup covariates (e.g., socioeconomic characteristics and other decision-maker variables).

A typical model will have main effects for each attribute and two-way interactions between

attributes and socioeconomic variables, meaning the number of coefficients $|\theta|$ that need to be estimated is $|\theta| = p + k \times p$. For example, for a choice model where each alternative has $p = 5$ attributes, with generic coefficients for main effects and no alternative specific constants, adding interaction effects for $k = 15$ covariates will involve estimating $|\theta| = 5 + 5 \times 15 = 80$ coefficients. If we increase the number of covariates to $k = 25$, the number of coefficients will increase to $|\theta| = 5 + 5 \times 25 = 130$. Including so many coefficients can harm the model's inferential ability, forecasting skill, and interpretability. To address this, consider how the model is estimated. For a multinomial logit model, the conditional probability P_{ni} that individual n chooses alternative i from a set of alternatives J is derived using the sigmoid (or softmax) function [166]. The model is estimated by minimizing a loss function $f(\theta)$, usually the negative log-likelihood function $\min_{\theta \in \Theta} -l(\theta)$ over parameter space $\Theta = R^{|\theta|}$, where $l(\theta)$ is:

$$l(\theta) = \sum_{n=1}^N \sum_{i=1}^J I_{ni} \log(P_{ni}) = \sum_{n=1}^N \sum_{i=1}^J I_{ni} \log\left(\frac{e^{V_{ni}}}{\sum_{j=1}^J e^{V_{nj}}}\right) \quad (4.1)$$

and I_{ni} is an indicator variable for whether alternative i was chosen by decision-maker n . If we introduce a penalty $Q_\lambda(\theta)$ into the loss function $f(\theta)$ we can induce *sparse* models, that set some coefficients to zero [233]. The loss function $f(\theta)$ for a sparse multinomial logit model will correspond to a Lasso penalized log-likelihood function as in [233]:

$$\min_{\theta \in \Theta} f(\theta) = \min_{\theta \in \Theta} \left[-l(\theta) + Q_\lambda(\theta) \right] = \min_{\theta \in \Theta} \left[-l(\theta) + \lambda \sum_{i=j}^p \gamma_j |\theta_j| \right] \quad (4.2)$$

where $\lambda \geq 0$ is a tuning parameter that determines how much the loss function is penalized for having a larger magnitude parameter vector, and is selected by cross validation [261]. The traditional Lasso weights each parameter equally, with $\gamma_j = 1$ for $j > 0$, except the intercept (with $\gamma_0 = 0$) [80]. The adaptive Lasso penalty scales λ by $\gamma_j \geq 0$, giving the oracle property [268], meaning that the estimator is consistent in variable selection for variables in the true model support S (i.e., $\lim_{n \rightarrow \infty} P(S_n = S) = 1$, where S_n is the support of a model with n

observations and S is the true model support), namely those elements in θ that should not be mapped to zero [61, 268]. The weighting $\gamma_j = (1/|\theta^{MLE}|)^\delta$ has this property, where θ^{MLE} is the maximum likelihood estimate of the θ vector, and δ is an additional tuning parameter that can also be selected by (two-dimensional) cross-validation [268]. Intuitively, the weighting scheme penalizes coefficients more if their MLE is closer to zero [268]. The adaptive Lasso is only asymptotically unbiased [268], but a straightforward debiasing technique uses linear interpolation between the unpenalized simulated maximum likelihood estimator (MLE) and the penalized solution [170]. A similar approach uses the adaptive Lasso for variable selection and the MLE for estimation on the selected variables, an approach known as hybrid estimation [123].

To estimate the model, we notice $-l(\theta)$ is convex, but $Q_\lambda(\theta)$ is not continuously differentiable [61], ruling out second order methods like BFGS that use the Hessian of the objective function [176, 92, 76, 37, 200]. For a Lasso regularized multinomial logit model we can decompose $f(\theta) = g(\theta) + h(\theta)$, where $g(\theta) = -l(\theta)$ and $h(\theta) = \lambda \sum_{i=1}^p \gamma_j |\theta_j|$. That is, $f(\theta)$ can be decomposed into the sum of two parts, a convex and differentiable function $g(\theta)$ and a convex non-differentiable function $h(\theta)$. The problem can then be efficiently solved using proximal gradient descent [31], where we use a second order approximation of g without modifying h , and define a proximal mapping $\text{prox}_t(x) = \arg \min_z \frac{1}{2t} \|x - z\|_2^2 + h(z)$ [93]. It follows that the proximal gradient descent update will be $\theta_{(l+1)} = \text{prox}_t(\theta_{(l)} - t_l \nabla g(\theta_{(l)}))$. The gradient of $g(\theta)$ can be obtain analytically deriving the log-likelihood function. The update for the j -th coordinate for the proximal operator will be:

$$S_{\lambda\gamma_j}(\theta_j, t) = \begin{cases} \theta_j - \lambda\gamma_j t \text{ sign}(\theta_j) & |\theta_j| > \lambda\gamma_j t \\ 0 & |\theta_j| \leq \lambda\gamma_j t \end{cases} \quad (4.3)$$

Details of this derivation can be seen in [82]. The proximal operator can be written more succinctly as the soft-thresholding operator with its j -th component as $S_{\lambda\gamma_j}(\theta_j, \lambda\gamma_j t) = \text{sign}(\theta_j) \cdot \max\{|\theta_j| - \lambda\gamma_j t, 0\}$ [143]. The proximal update will be $\theta_{j(t+1)} = S_{\lambda\gamma_j}(\theta_{j(t)} - t_l \nabla g(\theta_{j(t)}, t_l)$. This algorithm is known as iterative soft-thresholding, or ISTA, with an $O(1/\epsilon)$ convergence rate [18]. The step size t_l can be selected using backtracking line search to satisfy a sufficient decrease condition (the Armijo condition) [176]. With minor modifications, a convergence rate of $O(1/\sqrt{\epsilon})$ can be achieved by extending the Nesterov acceleration method [175] to composite functions, known as the FISTA method [18].

4.3 Simulation

As a proof of concept we apply the sparse multinomial logit model to simulated data from a conjoint analysis experiment. Conjoint analysis is a stated preference method [153] where respondents are asked to choose between hypothetical alternatives that differ in the level of one or more attributes, including price, allowing researchers to infer the willingness-to-pay, or the marginal rate of substitution between each attribute and price for each attribute $MRS_k = \theta_k / \theta_{price}$ [58, 105].

We simulated an experiment where a sample of 200 pseudo-individuals faced 10 choice sets with three alternatives, drawn from an optimal orthogonal design with five attributes and four levels, for a total of 2,000 choices [4]. We also simulated $k = 25$ uncorrelated covariates that varied only across decision-makers, drawn from a standard multivariate normal distribution. We assumed utility functions had coefficients for each attribute and two-way interactions between decision-maker characteristics and each attribute, giving a total of $q = p + p \times k = 5 + 125 = 130$ parameters. We used a vector of coefficients for each generic attribute $\beta = \{4, 1, -1, 2, -2\}$, where the first coefficient $\beta_1 = 4$ corresponds to the cost attribute. We also generated a vector of 125 coefficients for each covariate α , drawing

randomly from a uniform distribution between -2 to 2 with an 80% chance of being equal to zero.

In Figure 4.1, the maximum likelihood estimator fails to set most of the coefficients to zero. Next, we estimate coefficients by varying λ , known as the regularization path [80]. Increasing the penalty term λ shrinks the coefficients towards zero as expected, also increasing bias $\|\hat{\theta} - \theta\|_1$. We compute the Bayesian Information Criterion (BIC) for each value of λ in the sequence, assuming degrees of freedom can be approximated by the number of non-zero coefficients $|S(\hat{\theta})|$ [269]. As an alternative (but slower) approach, we select the penalty via cross validation, using the value for λ that gives a classification error one standard deviation above the minimum ($\lambda \approx 45$) [79]. As shown in Figure 4.1 on right, BIC is also minimized for values of λ in the range of those found with cross-validation [180].

In Figure 4.2, we represent the support recovered by each of the models tested using a grey-scaled heat-map. A higher absolute value is represented by darker colors and ivory represents coefficients equal to zero. We used estimates from the standard multinomial logit model via maximum likelihood (MNL), the sparse multinomial logit model with a lasso penalty (SMNL), and the adaptive version of the sparse multinomial logit (ASMNL). The maximum likelihood estimator is clearly not sparse, with all coefficients different than zero. The support recovered by the Lasso penalty is close to the original support, with a penalty $\lambda \approx 45$ selected with the one standard deviation rule. Nonetheless, the sparse multinomial logit models fails to zero-out variables that are not in the original support in 12% of the cases. Models that hold the oracle property overcome this issue, where the adaptive sparse multinomial logit models recovers the true support with 100% accuracy with $\delta = 1$ [268].

We extended the simulation experiment by varying the number of covariates k and the correlation ρ between the covariates. In Table 4.1, we present results from several alternative models in terms of the number of non-zero coefficients $|S(\hat{\theta})|$ and bias $\|\hat{\theta} - \theta\|_1$. In Table 4.1 we first present our simulation results for an oracle estimator, namely the maximum




Figure 4.1: Regularization path and bias. 2,000 choices of simulated data from a multinomial logit model with five alternatives and 25 covariates. Increasing the penalty term λ shrinks the coefficients towards zero. Increasing the penalty term λ allows for sparse models at the cost of a higher bias.




Figure 4.2: Support recovery for $\lambda \approx 45$. Original support, maximum likelihood estimator (MNL), sparse multinomial logit (SMNL), adaptive sparse multinomial logit (ASMNL). Low values were jittered to ease comprehension. Darker colors indicate larger coefficients in absolute value.

likelihood estimator assuming that the true support $S(\theta)$ is known apriori (*Oracle*). Second, we present results for the standard multinomial logit model (MNL). Third, we present results

for our sparse multinomial logit model (SMNL) and its adaptive version (ASMNL). Finally, we present results from an hybrid approach (HASMNL) where the adaptive model is used to discover the right support $S(\hat{\theta})$ and next a standard multinomial logit model is estimated with all features with non-zero coefficients in $S(\hat{\theta})$. We simulated 100 samples with the same specifications described above and $\lambda = 45$. Increasing correlation between covariates deteriorates the model variable selection performance to some extent. In the scenario with $k = 25$, if we increase correlation between covariates to $\rho = 0.5$ it becomes harder to discover the original support. The error rate increases from %1 to %6 of the cases, where the model assumed variables were zero when they were not (NZ) or it assumed variable were different than zero when they should be zero-out (Z).

Table 4.1: Sparse multinomial logit simulation results. 100 simulated samples with different number of covariates k and increasing correlation ρ . We fixed tuning parameter $\lambda = 45$ and $\delta = 1$. MNL: multinomial logit, SMNL: sparse multinomial logit, ASMNL: adaptive sparse multinomial logit and HASMNL: hybrid adaptive sparse multinomial logit. NZ : variables different than zero are zero-out, (Z): zero variables are not zero-out.

	$\rho = 0$				$\rho = 0.5$			
	$ S(\hat{\theta}) $	$\ \hat{\theta} - \theta\ _1$	Z	NZ	$ S(\hat{\theta}) $	$\ \hat{\theta} - \theta\ _1$	Z	NZ
$k = 15$								
Oracle	23 (0.0)	2.2 (1.2)	0.0	0.0	23 (0.0)	2.4 (1.0)	0.0	0.0
MNL	80 (0.0)	9.3 (2.7)	57	0.0	80 (0.0)	11.2 (3.2)	57	0.0
SMNL	27.4 (2.2)	21.2 (0.4)	4.4	0.0	27.8 (2.1)	21.9 (0.5)	5.2	0.4
ASMNL	23 (0.1)	18.9 (0.6)	0.0	0.0	22.5 (0.7)	19.7 (0.7)	0.0	0.5
HASMNL	23 (0.1)	2.2 (1.2)	0.0	0.0	22.5 (0.7)	2.9 (1.3)	0.0	0.5
$k = 25$								
Oracle	34 (0.0)	4.5 (2.4)	0.0	0.0	34 (0.0)	4.4 (1.9)	0.0	0.0
MNL	130 (0.0)	32.3 (9.6)	96	0.0	130 (0.0)	31.3 (9.9)	96	0.0
SMNL	40.7 (2.6)	34.1 (0.4)	6.8	0.2	37.9 (2.6)	34.6 (0.4)	5.7	1.8
ASMNL	33.7 (0.6)	29.8 (0.9)	0.0	0.4	31.8 (1.4)	30.4 (1.0)	0.0	2.2
HASMNL	33.6 (0.6)	4.3 (2.2)	0.0	0.4	31.8 (1.4)	6.0 (2.5)	0.0	2.2

To show how estimation bias in penalized models transfers to willingness to pay (WTP),

we focused our analysis on the 25 uncorrelated covariates. Assuming the first attribute corresponds to the cost vector, we computed the ratio of coefficients with respect to the four non-cost attributes. All socioeconomic covariates were mean-centered, so are neglected in our computations. Figure 4.3 presents the distribution for willingness to pay across simulations for attributes 2 and 4. The red line shows the true WTP value. WTP for attributes 1 and 3 show the same pattern, but with the opposite sign. As it is unbiased, the maximum likelihood estimator recovers the true willingness to pay, while both the sparse and the adaptive sparse multinomial logit models introduce bias into the WTP (15% lower on average for the Lasso model and 10% lower on average for the adaptive Lasso model). The hybrid approach (HASMNL), on the other hand, debiases coefficient estimates, providing a solution with both sparsity and accurate WTP estimates.

4.4 Empirical analysis

Finally, we use data from a conjoint analysis survey designed to estimate willingness to pay (WTP) to avoid environmental impacts associated with different electricity generation technologies. There have been several efforts to estimate willingness to pay values for externalities of electricity generation. Previous experiments have valued wind power generation externalities [7, 10], and hydroelectric generation externalities [94, 224]. However, electricity supply is composed of several sources with different external costs. The simultaneous valuation of several environmental goods is a complex empirical matter, as electricity generation can account for numerous environmental externalities and any survey effort can deal at most with only a small number of those impacts; otherwise the cognitive burden imposed on respondents would be prohibitive. In this case, the conjoint analysis technique provides a simple way of controlling simultaneously for the presence of several externalities of electricity generation [154]. Only a handful of studies have investigated simultaneously willingness to pay values

for external costs of different electricity sources. Landscape impacts, air pollution, wildlife impacts and employment effects are assessed in [22] and greenhouse gas emissions, security of energy supply and employment effects in [152]. Thus, the present work contributes to the extension of the application of choice experiments to this subject.

We implemented the conjoint analysis valuation survey in Santiago, the capital of Chile in January 2013, where samples were distributed among Santiago’s boroughs to mimic the population income distribution. A full version of the questionnaire and the survey data is available online (osf.io/uqtjb). At the time of the survey, a broad public debate was taking place on how Chile should double electricity generation in the next 20-years to meet steadily increasing demand. Chile has tremendous potential to develop water resources, and




Figure 4.3: Sample average willingness to pay values across simulations for attributes 2 and 4. 100 simulations with 25 uncorrelated covariates. True WTP values in red lines. WTP for attributes 1 and 3 shows the same pattern, but with the opposite sign. We added a scalar 0.1 for attribute 4 to avoid overlapping between the figures.

a large hydro-power dam project was planned for the Aysen Region, a world heritage area for biodiversity in the southern part of the country. Large coal-fired power plants in the Atacama Region in northern Chile could also provide a good option for electricity generation at a reasonable private cost. Nevertheless, a share of non-traditional renewable sources, such as wind or solar, could replace both thermoelectric and hydroelectric facilities if required. Although Chilean authorities were considering a mix of these alternatives for future power generation, none of them was universally accepted by the public. Large hydro-power projects in the south of the country were subject to strong criticism from several environmental NGOs and the civil society at large; coal-fired power plants near demand sites were opposed by nearby residents due to health risks from pollution. Even nontraditional sources, such as wind, met with resistance from local communities due their landscape impacts. Prior to this study, political decisions focused only on private costs, without a thorough consideration of externalities. Willingness to pay values to avoid the impacts of electricity generation would help public decision-making in this hotly contested topic.

4.4.1 Survey design

In our survey, participants were asked to choose between alternative hypothetical public electricity generation programs that differ in their level of four environmental impacts in specific areas of the country. The hypothetical scenario was presented as a strategic program seeking to understand household's preferences regarding the future Chilean electricity mix and its associated environmental local externalities [110]. Three alternatives were presented to household heads: the electricity generation development plan at the moment of the survey (the status quo), and two alternative plans, A and B. Each alternative presented the environmental impacts of different scenarios for the electricity mix. In each scenario, individuals chose their preferred option. The first option represented traditional electricity sources considered to

be constructed in the business-as-usual scenario (BAU), with its associated environmental impacts at zero additional cost. Alternatives A and B represented different development plans with higher costs at a lower environmental impact. Respondents traded-off the level of environmental impact with a higher or lower electricity cost, allowing us to infer willingness to pay for each attribute [105].

The four environmental impacts considered in the scenarios were: destruction of native forest (ha), respiratory and heart diseases due to air pollution (emergency room visits per year), use of space (ha), and location either in pristine or non-pristine areas. Respondents traded-off the level of environmental impact with a higher or lower electricity cost, allowing inference about the willingness to pay for each attribute [105]. To select the alternatives presented in the task, an efficient statistical design was chosen using the software NGENE [190]. This design comprised twenty-four choice sets in two blocks of twelve choice sets per individual. Figure 4.2 summarizes attributes, metrics, levels and magnitudes selected for the survey. In Figure 4.4 we present an example of a choice set used to introduce the choice task for each participant.

Our sample included 486 heads of household in Santiago, Chile. Interviewers were specially trained for the occasion, and detailed socioeconomic and perception data were gathered during the interviews. We selected the following variables as potential covariates for our model: dummy variables that indicate if participants had directly seen any of the alternatives for electricity generation, such as dams, power plants or other non-traditional sources, if participants had visited any potentially affected area, or participants had family in those areas. Additionally, we controlled for participants gender, age, ethnic origin, membership in an environmental NGO, monthly electricity bill, income level, and family composition. Four qualitative scales adapted from [35, 36] were included as candidate variables: a social values scale with four factors (individual responsibility, responsibility of others, pro-social and altruistic), a trust in government scale with two factors (integrity and competence), an

Table 4.2: Experimental design. Levels for some of the attributes are also described graphically illustrating with reference the area of Santiago, Chile. For example, 6,000 hectares are equivalent to an area covered by five of Santiago’s boroughs (Providencia, Ñuñoa, Santiago, Recoleta and Estación Central). The average monthly consumption of a household in Santiago in 2014 was approximately 200 KWh. At a cost per KWh of *CLP* 125, the average monthly electricity bill was *CLP* 25,000 or *USD* \sim 40 per month. The highest cost scenario of *CLP* 5,000 or *USD* 8, would imply a 20% increment in the monthly electricity bill. Each emergency room visit is introduced as a lottery with a 25% chance of being hospitalized and a 2.5% chance of dying.

Attribute	Metric	BAU	(1)	(2)	(3)	(4)	(5)	(6)
Native Forest	Hectares	6,000	4,000	2,000	0			
								
Morbidity	Events per year	150	100	50	0			
								
Use of Space	Hectares	20,000	15,000	10,000	5,000			
								
Location		Pristine	Non-pristine					
Cost	MCLP month	0	0.5	1	2	3	4	5

environmental beliefs scale with three factors (ecologist, economist, and trust in technology), and environmental behavior scale with 3 factors (efficient, green, and activist). Scores were constructed for each scale using factor analysis from a set of statements with polychoric correlations [182]. Scale construction details are presented in Table 9.2 in the appendix section. We also include a variable about participant belief on the percentage of the population willing to pay to protect the environment.

The sample consisted of 56% female interviewees who were on average 48 years old. One per cent of the sample declared they belonged to an environmental NGO, and four percent declared an ethnic origin. Household average income was approximately *USD*16,000 per year, a little below the national purchasing power parity adjusted GDP per capita for 2014 of



Figure 4.4: Choice set example (original in Spanish)

USD23,000 (data.worldbank.org). 27% of the sample had seen a dam, 13% a power plant and 18% a wind farms. To avoid multicollinearity issues we finally selected the most meaningful variables with a correlation $\rho = 0.4$ or lower in absolute value. A statistical summary and a correlation matrix for the selected covariates can be found in Table 9.1 and Figure 9.1 in the appendix section, respectively. To control for hypothetical bias [173], we included a “cheap talk” budget reminder adapting those in [53] and [161]. Further, participants were randomly assigned to two conditions where either they were required or not to commit via a signed oath to a truthful answer [124]. The oath was presented before the choice task and stated: “I solemnly swear that during the following questions I will always answer truthfully, responsibly, according to my personal opinions and considering my family budget”. More details of the survey design and data collection can be found in De la Maza *et al.* [57].

4.4.2 Modeling results

We use the hybrid adaptive sparse multinomial logit model for estimation. We compare our results with a standard multinomial logit model and a mixed-logit model [235]. Multinomial and mixed logit models were estimated using mlogit R-package [51]. In the formulation, *Forest* stands for native forest destroyed, *Morbidity* represents the number of emergency room visits for respiratory or cardiovascular diseases, *Use of space* stands for land use and *Location* stands for a dummy variable indicating pristine (1) or non-pristine (0) areas. Henceforth, we defined the utility function for alternative i as:

$$\begin{aligned}
 V_i = & \alpha_i + \sum_j^p \alpha_j x_{ij} + \sum_j^p \sum_l^k \beta_{jl} x_{ij} z_{nl} \\
 & + \beta_{81} \times \text{Seen hydro} \times \text{Forest} \\
 & + \beta_{82} \times \text{Seen hydro} \times \text{Forest} \times \text{Location} \\
 & + \beta_{83} \times \text{Seen power} \times \text{Morbidity} \\
 & + \beta_{84} \times \text{Seen renewable} \times \text{Use of space} \\
 & + \beta_{85} \times \text{Public support} \times \text{Cost}
 \end{aligned} \tag{4.4}$$

Where:

x_i is {Forest, Forest:Location, Morbidity, Use of space, Cost}

z_n is {Signed oath, Gender, Age, Ethnicity, NGO, Electricity Bill, Income, Family, Children, Past visits, Responsibility, Pro-social, Altruistic, Green, Trust Technology, Trust Integrity}

The alternative specific constant for alternative B was set to zero to allow model identification [235]. If we account for main effects for all attributes the number of total features

q will be equal to 92: i) 5 coefficients for main effects including cost, ii) 2 coefficients for alternative specific constants excluding alternative B and iii) 85 coefficients for the selected interactions between attributes and covariates ($q = 2 + p + p \times k = 2 + 5 + 5 \times 17 = 92$). We did not penalize main effects or alternative specific constants. Modeling results are presented in Table 4.3. Our sparse model recovers 16 relevant two-way interactions between main effects and covariates with the penalty that minimizes BIC equal to $\lambda = 12$. Taste heterogeneity is well modeled using the approach, as the BIC from the sparse model is better than the mixed logit model, which treats subgroup variation as unobserved heterogeneity (random effects) [235]. Most of the reported coefficients are statistically significant ($p \leq 0.05$).

As shown in Figure 4.5, all main effects are negative. This indicates that increments in native forest destruction, use of space or morbidity effects will result in a lower utility level for the decision maker and hence a detriment to the attractiveness of an alternative that offers those increments. Additionally, as the alternative specific constant for choosing alternative A is positive, there is positive tendency to choose the alternative located at the middle beyond any other explanation offered by relevant covariates [186]. This tendency could reflect a heuristic method to reduce the cognitive burden posed by the task [186]. No relevant covariates were associated with impacts in native forest in non-pristine areas and use of space. Participants who had seen a power plant assigned an extra premium to avoid health impacts, possibly caused by an increased sense of risk [10]. Participants with a higher electricity bill assigned an extra premium to avoid health impacts, recognizing the relation between consumption and production. Moreover, participants that had visited the potentially affected areas were willing to pay more to avoid health risks, accounting for a potential use or option value [13]. Interestingly, older participants extracted a higher utility level from increments in health impacts and hence from more thermoelectric generation, suggesting generational differences on the definition of progress [35].

Regarding native forest destruction on pristine areas, membership in an environmental

Table 4.3: Modeling results

	<i>Dependent variable:</i>		
		choice	
	(MNL)	(MXL)	(ASMNL)
Forest	-0.04*** (0.01)	-0.05*** (0.01)	-0.04*** (0.01)
Morbidity	-0.35*** (0.06)	-0.48*** (0.08)	-0.28*** (0.07)
Use of space	-0.01** (0.003)	-0.01** (0.004)	-0.01** (0.004)
Cost	-0.25*** (0.01)	-0.54*** (0.05)	-0.23*** (0.02)
SQ	0.10 (0.14)	-0.15 (0.18)	0.05 (0.15)
A	0.13*** (0.04)	0.18*** (0.05)	0.14*** (0.04)
Forest:Location	-0.07*** (0.02)	-0.10*** (0.02)	-0.04* (0.02)
sd(Forest)		0.03 (0.07)	
sd(Morbidity)		-0.02 (0.30)	
sd(Use of space)		-0.03 (0.03)	
sd(Cost)		0.62*** (0.07)	
sd(Forest:Location)		-0.02 (0.06)	
Morbidity:Seen power			-0.16** (0.07)
Morbidity:Age			0.12*** (0.03)
Morbidity:Electricity Bill			-0.08*** (0.03)
Morbidity:Past visits			-0.30*** (0.07)
Cost:Signed oath			-0.08*** (0.02)
Cost:Public support			0.06*** (0.01)
Cost:Responsibility			0.05*** (0.01)
Cost:Pro-social			0.03*** (0.01)
Forest:Location:Ethnicity			0.13*** (0.03)
Forest:Location:NGO			-0.17*** (0.06)
Forest:location:Electricity Bill			-0.03*** (0.01)
Forest:Location:Responsibility			-0.02*** (0.01)
Forest:Location:Trust Technology			0.05*** (0.01)
Forest:Location:Trust Integrity			0.04*** (0.01)
Forest:Location:Family			-0.08*** (0.02)
Forest:Location:Past visits			-0.13*** (0.02)
Observations	4,395	4,395	4,395
Log Likelihood	-4,599	-4,575	-4,348
BIC	9,258	9,251	8,922

*p<0.1; **p<0.05; ***p<0.01

NGO played a relevant role, showing a strong rejection from members for interventions on the Patagonia regions (Aysen). Participants with a higher electricity bill and an elevated sense of individual responsibility for the impacts, also assigned an extra premium to avoid the destruction of native forest in pristine areas. Participants that had visited or had family in the potentially affected areas were willing to pay more for their protection, again suggesting a potential use or option value [13]. Not surprisingly, participants with a favorable trust in government's integrity to act in favor of the community or the belief that technology can provide effective environmental protection, would accept a higher intervention in pristine areas than the average population [36]. Finally, participants from indigenous ethnic groups would extract a positive utility level from an intervention of the native forest in the Patagonia, a land that was not occupied by their ancestors, perhaps foreseeing a withdraw of industry pressure on their own native lands.

Regarding costs, head of households that forecast a large public support for an environmental fee, that held strong pro-social values and with strong feelings of individual responsibility, assigned less importance to an increment in the electricity bill than the average population in exchange for environmental protection. Further, in accord with results found in [124], participants who committed themselves through a signed oath to a truthful answer assigned more importance to the cost attribute than on average. As claimed by [124], this so-called "truth telling commitment device" could induce respondents to bind themselves to their budget constraints. That is to say, a higher importance of costs might bear relation to a higher commitment to the survey instrument. The "oath effect", namely the commitment made when abiding to the written oath [124], has not yet being thoroughly explored in this kind of experiments and it presents an interesting alternative for reducing hypothetical bias.




Figure 4.5: Estimated coefficients ASML model

4.4.3 Willingness to pay

Finally we estimate willingness to pay (WTP) values. WTP represents the marginal rate of substitution between an specific attribute k and the cost attribute as $MRS_k = \theta_k / \theta_{cost}$ [58]. In Table 4.4 we present in the first column marginal WTP values in monthly figures, followed by sample mean values for each relevant covariate, where continuous covariates were mean-centered. In the third column, we present average WTP values in yearly figures. Finally, social willingness to pay for each impact is computed expanding average willingness to pay values to all households in the country (approximately 6 million) [10]. To account for uncertainty, HASML estimates were assumed to distribute asymptotically normal [268]. Confidence intervals were constructed with 10,000 simulations of parametric bootstrap drawing the required coefficients from a multivariate normal distribution with parameters those from

the model [207]. We extended this analysis to present the average external costs of electricity generation by technology. To do so, we account for the impacts for each generation technology as the average of the impact reported by companies in their environmental impact assessment. Externalities in terms of average costs (USD/MWh) are estimated as follows:

$$WTP [USD/MWh] = \frac{WTP [USD-metric/year] \cdot Impact [metric/MW]}{8760 \cdot CapacityFactor [\%]} \quad (4.5)$$

Table 4.4: Willingness to pay for hybrid adaptive sparse multinomial logit model

Impact	Mg.WTP (USD/unit-month)	Population mean	Avg. WTP (USD/unit-year)	Social WTP (MUSD/unit-year)
<i>Use of space (ha)</i>	5.2×10^{-5}		6.2×10^{-4}	3.7 [1.8, 5.7]
<i>Forest (ha)</i>	2.6×10^{-4}		3.2×10^{-3}	18.9 [13.1, 24.9]
<i>Forest:Location (ha)</i>	2.4×10^{-4}		2.9×10^{-3}	20.9 [9.3, 32.8]
-:Ethnicity	-5.9×10^{-4}	4.3%	-3.0×10^{-4}	
-:NGO	1.3×10^{-3}	1.2%	1.9×10^{-4}	
-:Electricity Bill	4.0×10^{-4}	0	0	
-:Responsibility	3.6×10^{-4}	0	0	
-:Trust Technology	-8.7×10^{-5}	0	0	
-:Trust Integrity	-2.1×10^{-5}	0	0	
-:Family	7.4×10^{-4}	12.4%	1.1×10^{-3}	
-:Past visits	1.1×10^{-3}	20%	2.6×10^{-3}	
<i>Morbidity (event-year)</i>	1.7×10^{-2}		2.0×10^{-1}	1,380 [950, 1,800]
-:Seen power	2.7×10^{-2}	17%	5.5×10^{-2}	
-:Age	9.6×10^{-3}	0	0	
-:Electricity Bill	2.2×10^{-2}	0	0	
-:Past visits	3.5×10^{-2}	20%	8.5×10^{-2}	

Cost coefficient for participants who committed themselves through a signed oath to a truthful answer.

In Figure 4.6, we present environmental impacts per technology. We complemented the

analysis including the social cost of carbon as a triangular distribution with mode the current carbon tax in Chile of USD 5 and minimum and maximum values of USD 0 and USD 20 per Ton of CO_2 [183]. We also present in Table 9.3 and Table 9.4 in the appendix section, the environmental impacts and private costs per technology, collected from stakeholders in a participatory process [193]. Figure 4.6 shows that there is a large difference between social costs and private costs for each technology. Coal-fired power plants account for the highest external costs with roughly half of the total impacts associated with climate change. For other thermoelectric technologies such as biomass, externalities are mainly associated with local health impacts. Regarding hydro-power, externalities associated with dam projects are one order of magnitude higher than those related to the “run-of-the river” technology. Indeed, the main difference corresponds to native forest destruction impacts. Location can vary this conclusion considerably as destruction of native forest in pristine areas more than doubled base impacts. As an important caveat, health impacts were estimated for low population areas ($\sim 10,000$ inhabitants). If projects are constructed in sites that had been commonly used for this purpose in the past, externalities could amount from ten to a hundred times the ones presented here. We can observe in Figure 4.6 on right that as the exposed population increased, the external impacts of coal also increased. If the population affected is approximately ten times larger ($\sim 100,000$ inhabitants) a solar project would hold a lower social cost. These caveat extends to other external costs, as the average impact in native forest could easily be exceeded three times depending on river gorge features. In the case of non-traditional renewable sources, although use of space represents the main impact, total external costs are valued well below the ones associated with traditional sources, a fact which confirms the advantage of promoting this type of technologies.




Figure 4.6: Average externalities per technology. HDP: Health damaging pollutants. We complemented the analysis including the social cost of carbon as a triangular distribution with mode the current carbon tax in Chile of USD 5 and minimum and maximum values of USD 0 and USD 20 per Ton of CO₂ [183]. Coal-fired power plants account for the highest external costs, followed by the externalities associated with dam projects. On left we observe that as the exposed population increased, the external impacts of coal also increased. Based on values reported in Table 9.3 and Table 9.4 in the appendix section.

4.5 Discussion

We developed and tested the sparse multinomial logit model that is able to automatically identify relevant covariates in a choice process with an adaptive lasso penalty. As the model can capture preference heterogeneity for different socioeconomic characteristics of decision makers in a parsimonious fashion, within the estimation process, it provides the means to recognize the stylized elements of subgroup behavior in a choice task. The approach is flexible enough to model a wide variety of phenomena while retaining a parametric form that is readily interpretable, making it useful for aiding individual and societal decisions. We show how to estimate our model with state-of-the-art convex optimization methods. Our model can systematically recover the true support at the cost of inducing bias in coefficients. We also show a simple hybrid method on how to debiased our sparse model results to

make them useful for welfare analysis, using an oracle estimator for variable selection and later the unbiased maximum likelihood estimator on the selected variables. Other popular regularization alternatives that give unbiased estimates and also hold oracle properties are non-convex. Some examples are the smoothly clipped absolute deviation penalty (SCAD) [61], the minimax concave penalty [267], or Bridge penalty [77]. The use of non-convex penalties can be understood as different types of adaptive weighting schemes for the lasso penalty [62]. Future work should evaluate the benefits of using non-convex penalties.

We used data from a conjoint analysis survey designed to estimate willingness to pay to avoid environmental impacts of electricity generation to test our model. The effect on willingness to pay of experience with the impacts, ties with the affected sites, age, electricity bill, ethnicity, membership in an environmental NGO, pro-social behavior, trust in government, trust in technology, the sense of individual responsibility and perception of public support for environmental protection was uncovered by the model. Regarding willingness to pay estimation, according to our calculations, use of space is valued on average in USD 4,000 [2,000, 6,000] per ha-year. Average social willingness to pay to avoid the destruction of native forest was estimated at USD 19,000 [13,000, 25,000] per ha-year, which roughly two-folds if native forest is in a pristine area increasing in USD USD 21,000 [9,000, 33,000] per ha-year. Our results are in line with [94] results, that valued use of space in approximately in USD 6,000 per ha-year and destruction of native forest in USD 21,000 per ha-year, accounting for the time value of money. Health impacts were valued in MMUSD 1,4 [0.9, 1.8] per event-year. Following [186], we introduced health impacts as a chain of events in which each subsequent outcome has a higher level of illness intensity at a decreasing probability of occurrence, ranging from an emergency room visits to an hospitalization, with death as the worst possible outcome. Each emergency room visit is introduced as a lottery with a 25% chance of being hospitalized and a 2.5% chance of dying. As our health endpoint is more comprehensive than other studies, comparison becomes harder. Nonetheless, focusing only in mortality

impacts can provide a counterfactual. In [183], the value of statistical life (VSL) for Chile was assumed to be in a range from MMUSD 0.2 to MMUSD 2. If we use the higher end in that range and a 2.5% chance of death, a lower bound for the value of our endpoint will correspond to MMUSD 0.05, 30 times lower than our estimation. As VSL values in Chile are commonly estimated in the context of motor vehicle accidents, this large differences could be explained by differences in risk perception, as health risks associated with air pollution from electricity generation might be perceived as uncontrollable and involuntary [73].

In order to develop an electricity mix that considers economic, environmental and social aspects, authorities must thoroughly take into account external costs. Our estimated values could be used to assess costs and benefits of alternative electricity grid scenarios as compared with the status quo. Results indicate that external costs associated with thermoelectric power plants are the highest among the evaluated technologies, followed by hydro-power and other non-traditional renewable sources. Externalities related with health outcomes represent the greater impact, followed by greenhouse emissions, native forest destruction and use of space. Externalities can vary depending on particular technological specifications designed to reduce environmental impacts or regarding location. These factors must be considered in defining the electricity mix that maximizes social welfare. For example, health effects for thermoelectric projects vary depending on the exposed population. If projects were developed in areas closer to urban developments, externalities could be two orders of magnitude higher those estimated. In addition, it should be noted that if new projects are developed in pristine areas, it would imply an additional social welfare loss. Although these outcomes allow drawing some conclusions, supplementary inquiries must be undertaken in order to determine willingness to pay values for environmental impacts of electricity generation with more precision. In this study, respondents do not represent those directly affected. On the contrary, estimations correspond mainly to altruistic or option values of the main urban consumers of electricity in the country, as one third of Chile's population corresponds to Santiago residents. Hence,

further work is required to extend these figures at a national level. Nevertheless, results represent a good first approximation and could be applied to estimate social benefits for alternative electricity mix scenarios vis-a-vis Business-as-usual, simulating dispatch merit based on social costs. Henceforth, our results can provide valuable insights to determine the mix that maximizes social welfare.

4.6 Conclusions

Discrete choice models are used in important domains, ranging from consumer product recommendations, to valuing of environmental impacts, weighing the course of action in individual and social decisions. In a typical modeling exercise, the modeler needs to select the parameters to include in the model ex-ante. The model specification task can be challenging when many socioeconomic variables can be relevant and could lead to the use of simple heuristics to choose relevant subgroups in the sample. Oversimplified models can then lead to biased parameter estimates, resulting in harmful distortions when used for policy evaluation. Our method can aid the preference modeler to identify relevant covariates for the choice model automatically, reducing reliance on the modeler's judgment. Practitioners of value elicitation methods will be able to use our approach to better understand taste heterogeneity across a sample of decision-makers. The method can also inform policy-makers, to help them better understand the welfare impacts of new policies for subgroups of the target population, and to design policy interventions tailored for the heterogeneous preferences of different communities.

5

Summary and conclusions

In this dissertation, we present a general framework to model preference heterogeneity in individual choices along two dimensions, preference content and structure. With our approach, we extend welfare analysis to recognize for these differences in the preferences of the public. We propose three new methods to improve preference learning for policy design. To evaluate the methods, we present three case studies where we combine experimental design methods from the decision sciences with machine learning algorithms. We first provide a new method to improve choice experiments, avoiding judgment bias when risks accumulate on time. Thus, in the first case study we assessed the perception of cumulative flood risks, how those perceptions

affect insurance decisions, and whether those risk perceptions can be influenced providing simple cumulative risk information. We found that participants' cumulative risk judgments were well represented by a bimodal distribution, with a group that severely underestimates the risk and a group that moderately overestimates it. We also found that individuals who underestimate cumulative risks made more risk-seeking choices. Our results show that a common approach of providing information about the annual risk of an adverse event (rather than cumulative risk) could expose public to a level of harm that they would not be willing to accept when fully informed. Instead, materials aimed at helping decision makers improve their choices should include cumulative risk information directly. Results from the judgment task suggest that participants are likely using an anchoring and adjustment heuristic where they first intuitively consider the base rate risk as a plausible answer. If the judgment error in this response is detected, other heuristics are searched for an applied. An additive rule such as multiplying the base rate risk by the number of periods is likely to be selected [227]. As we provide more information the cumulative risk functional form can be discovered by participants. Nevertheless, additional information can also be used as a new anchor, to construct cumulative risk judgments, without yet recognizing the true relation between time of exposure and risk. Further research should show how many new points on a cumulative risk curve are required to faithfully learn that curve [164]. Other factors influencing both cumulative risk judgments and choices, such as emotions [217], numeracy skills [179] and experience [95, 115], should be also better understood in the future. Furthermore, empirical tests in other domain can uncover the generality of our findings and encourages a fruitful line of future research.

Next, we designed a method to determine preference structure from choice data even when those data are inconsistent with the axioms required by the choice theory the model relies on. We leveraged recent advances in graph matching and non-linear embeddings, combined with pairwise comparison choice data, to cluster decision-makers based on what

they want (the content of their preferences) and whether they know what they want (the structure of their preferences). Across three experiments, including classic studies of risky choice and a two attribute study about state-level electricity generation portfolios, we found significant heterogeneity in both the content and structure of decision-maker preferences. Decision-makers most frequently choose in a way consistent with utility maximization, yet some decision-makers make choices consistent with heuristic rules, while others appear to be uncertain about their preferences. As a generalization of traditional preference analysis, the approach can be used to make recommendations for people who know what they want, uncover complex choice rules, and suggest paths toward clarification for those who are uncertain. Although several individuals appear to be consistent with utility theory, we discovered a large proportion of individuals with choice patterns consistent with a lexicographic order behavior, choosing based on a single attribute. Although some individuals have cyclic preferences, most groups were transitive in expectation. Nonetheless, some groups showed evidence of intransitivity patterns both at individual and group level. Welfare calculations on decision-makers who have intransitive preferences or transitive preferences that result from simplifying heuristics can be misleading. The most harmful condition occurs when individuals use simple heuristics that mimic utility maximization, resulting in overestimation of the willingness to compromise in a population.

Our approach involves applying different procedures sequentially. We first represent individual choices as graphs, then compute dissimilarities between graphs for all pairs of decision-makers, embed those dissimilarities into a lower dimensional space, and later clusters decision-makers based on their proximity in those embeddings. Henceforth, improvements to our method can come from automatic selection of hyper-parameters in each stage such as the number of clusters and the number of dimensions in the embedding process. Casting hyper-parameter selection within the objective function used in each optimization procedure within the algorithm could reduce reliance on expert judgment. Further, methods that can

discover a hierarchy structure between the uncovered clusters can be useful to summarize even further the structure of preferences. Sequential application of k-medians and hierarchical clustering algorithms showed promising results, but future work should automate this process. The experimental design also provides a challenge. A new experimental design method should identify the minimal number of pairwise comparisons required to recover preference structure. Empirical applications should be also extended to include more attributes and to other experimental domains. A future challenge involves discovering the content and structure of preferences with a unique statistical model. Methods that can perform both clustering and classification simultaneously provide an attractive starting point.

Finally, we designed a method to aid the preference modeler to identify relevant covariates for the choice model, improving the model specification task. In the third case, we used tools from machine learning to discover relevant covariates in a discrete choice model. We developed and tested a method for high-dimensional preference modeling, where irrelevant coefficients are set to zero with an adaptive lasso penalty. Our model can systematically recover the true support at the cost of inducing bias in coefficients. In economic and policy analysis, bias in coefficients is a serious concern, as can lead to incorrect conclusions and policy distortions. We show how to debias our sparse model results with state-of-the-art convex optimization methods. Next we tested the model on a real data set. We used data from a conjoint analysis survey designed to estimate willingness to pay to avoid environmental impacts associated with different electricity generation technologies. The effect on willingness to pay of trust, values, beliefs and socioeconomic characteristics was uncovered by the model. Taste heterogeneity was well captured by the model, as our it improves upon the mixed logit model, that mimics these differences as random effects.

Extensions of our sparse model can be used to automatically learn the nested structure for the alternatives in a choice task with a nested logit model [169]. A nested logit model can be approximated with a set of random alternative-specific constants [116]. The nesting

structure will then be reflected in the variance-covariance matrix of those constants or the error structure. A random effects approximation for a nested logit model requires knowing the error structure ahead of time. Nevertheless, if utility functions can be written in terms of random variables whose parameters determine the model structure, setting some of those parameters to zero using sparsity constraints has the effect of selecting the appropriate error structure.

Further, the model can be used to uncover different attribute non-attendance strategies used by groups of decision makers with a sparse latent class model. A common violation of preference continuity corresponds to ignore a subset of alternative attributes during the choice, namely attribute non-attendance [43, 196, 117, 111, 142]. In a latent class model [99, 100], individuals are assigned to different classes, based on observable features. The probability that an individual chooses an alternative is then modeled as a finite mixture of classes [99]. Each class can also represent an specific attribute non-attendance strategy, where some attributes are ignored [117, 111, 142]. Enforcing sparsity constraints to the coefficients in each class can then uncover attribute non-attendance patterns.

The three methods offered here can affect policy research in the decision-sciences, with applications to health and environmental decision-making, as well as fundamental studies of human cognition. Policy decisions ought to be based on individuals preferences and trade-offs decision makers would make between different private or public goods [12]. A major challenge faced by policy analysts is to design models that can capture individual preferences, even when descriptions can be misleading; that can uncover heterogeneity in preference structure to classify decision makers and extract insights from choice data even when data is inconsistent with prior models; and that can better understand taste heterogeneity across a sample of decision-makers, even when relevant parameters are unclear. Practitioners will be able to use these methods to control judgment bias during the experimental task in the model specification, classify decision-makers according to their preference structure, answering

first if they know what they want and henceforth what they want or the content of those preferences, and to build more parsimonious models that can identify the stylized elements in a choice task. Our methods can be used in important domains, ranging from consumer product recommender systems used by technology companies to sell products, to revealed preference studies used by researchers to infer quantities like the value of a statistical life or the value of other environmental externalities, to the design of voting schemes that determine the trajectory of nations, weighing the course of action in individual and social decisions. Modelers of preferences will be able to use these methods to inform decision-makers themselves through decision analysis or to help policy-makers to better understand the welfare impacts of new policies. Further, our methods can help to design policy interventions customized for the demands that each subgroup of the population mandates, and encourage better public decisions. We trust our contribution can promote a more sustainable common future.

6

References

- [1] Yonathan Aflalo, Alexander Bronstein, and Ron Kimmel. “On convex relaxation of graph isomorphism”. In: *Proceedings of the National Academy of Sciences* 112.10 (2015), pp. 2942–2947.
- [2] Sidney N Afriat. “Efficiency estimates of production functions”. In: *International Economic Review* 13 (1972), pp. 568–598.
- [3] Sidney N Afriat. “The construction of utility functions from expenditure data”. In: *International Economic Review* 8 (1967), pp. 67–77.
- [4] Hideo Aizaki and Kazushi Nishimura. “Design and analysis of choice experiments using R: a brief introduction”. In: *Agricultural Information Research* 17.2 (2008), pp. 86–94.
- [5] Anna Alberini, Alistair Hunt, and Anil Markandya. “Willingness to pay to reduce mortality risks: evidence from a three-country contingent valuation study”. In: *Environmental and Resource Economics* 33.2 (2006), pp. 251–264.
- [6] Francesco Alberoni. “Contribution to the study of subjective probability. I”. In: *The Journal of general psychology* 66.2 (1962), pp. 241–264. ISSN: 0022-1309.
- [7] Begoña Álvarez-Farizo and Nick Hanley. “Using conjoint analysis to quantify public preferences over the environmental impacts of wind farms. An example from Spain”. In: *Energy policy* 30.2 (2002), pp. 107–116.
- [8] Steffen Andersen et al. “Estimating subjective probabilities”. In: *Journal of Risk and Uncertainty* 48.3 (2014), pp. 207–229. ISSN: 0895-5646.
- [9] Kohei Arai and Ali Ridho Barakbah. “Hierarchical K-means: an algorithm for centroids initialization for K-means”. In: *Reports of the Faculty of Science and Engineering* 36.1 (2007), pp. 25–31.

- [10] Claudia Aravena, W George Hutchinson, and Alberto Longo. “Environmental pricing of externalities from different sources of electricity generation in Chile”. In: *Energy economics* 34.4 (2012), pp. 1214–1225.
- [11] Kenneth J Arrow. *Social choice and individual values*. John Wiley Sons, Inc: New York, 1951.
- [12] Kenneth J Arrow et al. “Is there a role for benefit-cost analysis in environmental, health, and safety regulation?” In: *Science* 272.5259 (1996), pp. 221–222.
- [13] Diego Azqueta Oyarzun. *Valoración económica de la calidad ambiental*. 333.7 333.339. McGraw-Hill, 1994.
- [14] László Babai and Eugene M Luks. “Canonical labeling of graphs”. In: *Proceedings of the fifteenth annual ACM symposium on Theory of computing*. ACM. 1983, pp. 171–183.
- [15] Maya Bar-Hillel. “On the subjective probability of compound events”. In: *Organizational behavior and human performance* 9.3 (1973), pp. 396–406. ISSN: 0030-5073.
- [16] Julie Barnett and Glynis M Breakwell. “Risk perception and experience: Hazard personality profiles and individual differences”. In: *Risk Analysis* 21.1 (2001), pp. 171–178.
- [17] Duane D Baumann and John H Sims. “Flood insurance: Some determinants of adoption”. In: *Economic Geography* 54.3 (1978), pp. 189–196.
- [18] Amir Beck and Marc Teboulle. “A fast iterative shrinkage-thresholding algorithm for linear inverse problems”. In: *SIAM journal on imaging sciences* 2.1 (2009), pp. 183–202.
- [19] Alexandre Belloni and Victor Chernozhukov. *High dimensional sparse econometric models: An introduction*. Springer, 2011.
- [20] Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. “Inference for high-dimensional sparse econometric models”. In: *arXiv preprint arXiv:1201.0220* (2011).
- [21] Jeremy Bentham. *An introduction to the principles of morals and legislation*. Clarendon Press, 1879.
- [22] Ariel Bergmann, Nick Hanley, and Robert Wright. “Valuing the attributes of renewable energy investments”. In: *Energy policy* 34.9 (2006), pp. 1004–1014.
- [23] Sudeep Bhatia. “Associations and the accumulation of preference.” In: *Psychological Review* 120.3 (2013), pp. 522–543.
- [24] Michael H Birnbaum. “New paradoxes of risky decision making.” In: *Psychological Review* 115.2 (2008), pp. 463–501.
- [25] Michael H Birnbaum. “Testing Mixture Models of Transitive Preference: Comment on Regenwetter, Dana, and Davis-stober (2011)”. In: *Psychological Review* 118.4 (2011), pp. 675–683.
- [26] Michael H Birnbaum and Alfredo Chavez. “Tests of theories of decision making: Violations of branch independence and distribution independence”. In: *Organizational Behavior and human decision Processes* 71.2 (1997), pp. 161–194.

- [27] Michael H Birnbaum and Roman J Gutierrez. “Testing for intransitivity of preferences predicted by a lexicographic semi-order”. In: *Organizational Behavior and Human Decision Processes* 104.1 (2007), pp. 96–112.
- [28] WJ Wouter Botzen and Jeroen CJM van den Bergh. “Risk attitudes to low-probability climate change risks: WTP for flood insurance”. In: *Journal of Economic Behavior & Organization* 82.1 (2012), pp. 151–166.
- [29] WJ Wouter Botzen and Jeroen CJM Van Den Bergh. “Monetary valuation of insurance against flood risk under climate change”. In: *International Economic Review* 53.3 (2012), pp. 1005–1026.
- [30] Denis Bouyssou and Philippe Vincke. “Binary relations and preference modeling”. In: *Decision-making Process: Concepts and Methods* (2010), pp. 49–84.
- [31] Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- [32] Eduard Brandstätter, Gerd Gigerenzer, and Ralph Hertwig. “The priority heuristic: making choices without trade-offs.” In: *Psychological Review* 113.2 (2006), pp. 409–432.
- [33] John FP Bridges et al. “Conjoint analysis applications in health—a checklist: a report of the ISPOR Good Research Practices for Conjoint Analysis Task Force”. In: *Value in health* 14.4 (2011), pp. 403–413.
- [34] Joel Brockner et al. “Regulatory focus and the probability estimates of conjunctive and disjunctive events”. In: *Organizational behavior and human decision processes* 87.1 (2002), pp. 5–24. ISSN: 0749-5978.
- [35] Nicolás C Bronfman et al. “Understanding attitudes and pro-environmental behaviors in a Chilean community”. In: *Sustainability* 7.10 (2015), pp. 14133–14152.
- [36] Nicolás C Bronfman et al. “Understanding social acceptance of electricity generation sources”. In: *Energy Policy* 46 (2012), pp. 246–252.
- [37] Charles G Broyden. “The convergence of single-rank quasi-Newton methods”. In: *Mathematics of Computation* 24.110 (1970), pp. 365–382.
- [38] Philip Bubeck, WJ Wouter Botzen, and Jeroen CJM Aerts. “A review of risk perceptions and other factors that influence flood mitigation behavior”. In: *Risk analysis* 32.9 (2012), pp. 1481–1495.
- [39] Donald H Burn. “Evaluation of regional flood frequency analysis with a region of influence approach”. In: *Water Resources Research* 26.10 (1990), pp. 2257–2265.
- [40] Jerome R Busemeyer and James T Townsend. “Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment.” In: *Psychological Review* 100.3 (1993), pp. 432–459.
- [41] Carter T Butts and Kathleen M Carley. “Some simple algorithms for structural comparison”. In: *Computational & Mathematical Organization Theory* 11.4 (2005), pp. 291–305.

- [42] Colin F Camerer and Howard Kunreuther. “Decision processes for low probability events: Policy implications”. In: *Journal of Policy Analysis and Management* 8.4 (1989), pp. 565–592. ISSN: 1520-6688.
- [43] Danny Campbell, David A Hensher, and Riccardo Scarpa. “Non-attendance to attributes in environmental choice analysis: a latent class specification”. In: *Journal of environmental planning and management* 54.8 (2011), pp. 1061–1076.
- [44] Michael D Clark et al. “Discrete choice experiments in health economics: a review of the literature”. In: *Pharmacoeconomics* 32.9 (2014), pp. 883–902.
- [45] John Cohen and EI Chesnick. “The doctrine of psychological chances”. In: *British Journal of Psychology* 61.3 (1970), pp. 323–334. ISSN: 2044-8295.
- [46] John Cohen, EI Chesnick, and D Haran. “A confirmation of the inertial-psi effect in sequential choice and decision”. In: *British Journal of Psychology* 63.1 (1972), pp. 41–46. ISSN: 2044-8295.
- [47] John Cohen, EI Chesnick, and D Haran. “Evaluation of compound probabilities in sequential choice”. In: *Nature* 232.5310 (1971), pp. 414–416. ISSN: 1476-4687.
- [48] John Cohen and CEM Hansel. “The nature of decisions in gambling: Equivalence of single and compound subjective probabilities”. In: *Acta Psychologica* 13 (1957), pp. 357–370. ISSN: 0001-6918.
- [49] Clyde H Coombs and Richard C Kao. “On a connection between factor analysis and multidimensional unfolding”. In: *Psychometrika* 25.3 (1960), pp. 219–231.
- [50] Marc-Alexandre Côté and Hugo Larochelle. “An infinite restricted Boltzmann machine”. In: *Neural computation* 28.7 (2016), pp. 1265–1288.
- [51] Yves Croissant et al. “Estimation of multinomial logit models in R: The mlogit Packages”. In: *R package version 0.2-2.* (2012).
- [52] Matthew JC Crump, John V McDonnell, and Todd M Gureckis. “Evaluating Amazon’s Mechanical Turk as a tool for experimental behavioral research”. In: *PloS one* 8.3 (2013), e57410. ISSN: 1932-6203.
- [53] Ronald G Cummings and Laura O Taylor. “Unbiased value estimates for environmental goods: a cheap talk design for the contingent valuation method”. In: *American Economic Review* 89.3 (1999), pp. 649–665.
- [54] Kim Darby et al. “Decomposing local: A conjoint analysis of locally produced foods”. In: *American Journal of Agricultural Economics* 90.2 (2008), pp. 476–486.
- [55] Robert L Davis. “Structures of dominance relations”. In: *Bulletin of Mathematical Biology* 16.2 (1954), pp. 131–140.
- [56] Clinton P Davis-Stober. “Analysis of multinomial models under inequality constraints: Applications to measurement theory”. In: *Journal of Mathematical Psychology* 53.1 (2009), pp. 1–13.

- [57] Cristobal De La Maza. *Willingness to pay to avoid the environmental impacts of electricity generation*. Tech. rep. Working paper 58, LACEEP <http://www.laceep.org/>, 2013.
- [58] Juan de Dios Ortúzar. *Modelos econométricos de elección discreta*. Universidad Católica de Chile, 2000.
- [59] James K Doyle. “Judging Cumulative Risk”. In: *Journal of Applied Social Psychology* 27.6 (1997), pp. 500–524. ISSN: 1559-1816.
- [60] Ido Erev et al. “From Anomalies to Forecasts: Toward a Descriptive Model of Decisions Under Risk, Under Ambiguity, and From Experience.” In: *Psychological Review* 124.4 (2017), pp. 369–409.
- [61] Jianqing Fan and Runze Li. “Variable selection via nonconcave penalized likelihood and its oracle properties”. In: *Journal of the American statistical Association* 96.456 (2001), pp. 1348–1360.
- [62] Jianqing Fan and Jinchi Lv. “A selective overview of variable selection in high dimensional feature space”. In: *Statistica Sinica* 20.1 (2010), p. 101.
- [63] Jianqing Fan, Jinchi Lv, and Lei Qi. “Sparse high-dimensional models in economics”. In: *Annu. Rev. Econ.* 3.1 (2011), pp. 291–317.
- [64] Vivek F Farias, Srikanth Jagabathula, and Devavrat Shah. “Sparse choice models”. In: *Information Sciences and Systems (CISS), 2012 46th Annual Conference on*. IEEE, 2012, pp. 1–28.
- [65] FEMA. *Defining Flood Risks*. Web Page. Monday, 02-May-2016, 7:36 PM (EDT) 2016. URL: https://www.floodsmart.gov/floodsmart/pages/flooding_flood_risks/defining_flood_risks.jsp.
- [66] Baruch Fischhoff. “Cognitive processes in stated preference methods”. In: *Handbook of environmental economics* 2 (2005), pp. 937–968.
- [67] Baruch Fischhoff. “Predicting frames”. In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9.1 (1983), p. 103. ISSN: 1939-1285.
- [68] Baruch Fischhoff. “Value elicitation: Is there anything in there?” In: *American Psychologist* 46.8 (1991), pp. 835–847.
- [69] Baruch Fischhoff and Maya Bar-Hillel. “Focusing techniques: A shortcut to improving probability judgments?” In: *Organizational Behavior and Human Performance* 34.2 (1984), pp. 175–194. ISSN: 0030-5073.
- [70] Baruch Fischhoff, Ann Bostrom, and Marilyn Quadrel. “Risk perception and communication”. In: *Annual Review Public Health* 14 (1993), pp. 183–203.
- [71] Baruch Fischhoff and Alexander L Davis. “Communicating scientific uncertainty”. In: *Proceedings of the National Academy of Sciences* 111.Supplement 4 (2014), pp. 13664–13671. ISSN: 0027-8424.

- [72] Baruch Fischhoff et al. “How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits”. In: *Policy sciences* 9.2 (1978), pp. 127–152. ISSN: 0032-2687.
- [73] Baruch Fischhoff et al. “How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits”. In: *Policy sciences* 9.2 (1978), pp. 127–152.
- [74] Peter C Fishburn. “A study of lexicographic expected utility”. In: *Management Science* 17.11 (1971), pp. 672–678.
- [75] Peter C Fishburn. “Subjective expected utility: A review of normative theories”. In: *Theory and Decision* 13.2 (1981), pp. 139–199.
- [76] Roger Fletcher. “A new approach to variable metric algorithms”. In: *The computer journal* 13.3 (1970), pp. 317–322.
- [77] LLdiko E Frank and Jerome Friedman. “A statistical view of some chemometrics regression tools”. In: *Technometrics* 35.2 (1993), pp. 109–135.
- [78] Marguerite Frank and Philip Wolfe. “An algorithm for quadratic programming”. In: *Naval Research Logistics (NRL)* 3.1-2 (1956), pp. 95–110.
- [79] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. “glmnet: Lasso and elastic-net regularized generalized linear models”. In: *R package version 1.4* (2009).
- [80] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. “Regularization paths for generalized linear models via coordinate descent”. In: *Journal of statistical software* 33.1 (2010), p. 1.
- [81] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*. Vol. 1. Springer series in statistics Springer, Berlin, 2001.
- [82] Jerome Friedman et al. “Pathwise coordinate optimization”. In: *The Annals of Applied Statistics* 1.2 (2007), pp. 302–332.
- [83] Richard Fuller, Nigel Dudley, and Jon Blacktop. “Older people understanding of cumulative risks when provided with annual stroke risk information”. In: *Postgraduate medical journal* 80.949 (2004), pp. 677–678. ISSN: 1469-0756.
- [84] Patricia Galilea and Juan de Dios Ortuzar. “Valuing noise level reductions in a residential location context”. In: *Transportation Research Part D: Transport and Environment* 10.4 (2005), pp. 305–322.
- [85] Zhifeng Gao and Ted C Schroeder. “Effects of label information on consumer willingness-to-pay for food attributes”. In: *American Journal of Agricultural Economics* 91.3 (2009), pp. 795–809.
- [86] Gerd Gigerenzer. “Fast and frugal heuristics: The tools of bounded rationality”. In: *Blackwell handbook of judgment and decision making* (2004), pp. 62–88.

- [87] Gerd Gigerenzer and Daniel G. Goldstein. “Betting on one good reason: Take the best and its relatives”. In: *Simple Heuristics that Make Us Smart*. Ed. by Gerd Gigerenzer, P.M. Todd, and ABC Research Group. Oxford University Press, New York, 1999, pp. 75–95.
- [88] Gerd Gigerenzer, Ralph Hertwig, and Thorsten Pachur. *Heuristics: The foundations of adaptive behavior*. Oxford University Press, Inc., 2011. ISBN: 0199744289.
- [89] Gerd Gigerenzer and Peter M Todd. *Simple heuristics that make us smart*. Oxford University Press, USA, 1999. ISBN: 0195143817.
- [90] Maureen D Giron and Erniel B Barrios. “High Dimensional Nonparametric Discrete Choice Model”. In: *Proceedings of the 11th National Convention on Statistics (NCS), Mandaluyong City, Philippines* (2010).
- [91] Ken Goldberg et al. “Eigentaste: A constant time collaborative filtering algorithm”. In: *Information Retrieval 4.2* (2001), pp. 133–151.
- [92] Donald Goldfarb. “A family of variable-metric methods derived by variational means”. In: *Mathematics of computation* 24.109 (1970), pp. 23–26.
- [93] Tom Goldstein, Christoph Studer, and Richard Baraniuk. “A field guide to forward-backward splitting with a FASTA implementation”. In: *arXiv preprint arXiv:1411.3406* (2014).
- [94] Andres Gomez-Lobo, Javier Nunez, and Cristobal Ruiz-Tagle. *Choice complexity in a Stated Choice Experiment: valuing environmental resources in Chile*. Tech. rep. University of Chile, Department of Economics, 2003.
- [95] Cleotilde Gonzalez, Javier F Lerch, and Christian Lebiere. “Instance-based learning in dynamic decision making”. In: *Cognitive Science* 27.4 (2003), pp. 591–635. ISSN: 0364-0213.
- [96] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [97] William H Greene. *Econometric analysis*. Granite Hill Publishers, 2012.
- [98] William H Greene and David A Hensher. “A latent class model for discrete choice analysis: contrasts with mixed logit”. In: *Transportation Research Part B: Methodological* 37.8 (2003), pp. 681–698. ISSN: 0191-2615.
- [99] William H Greene and David A Hensher. “A latent class model for discrete choice analysis: contrasts with mixed logit”. In: *Transportation Research Part B: Methodological* 37.8 (2003), pp. 681–698.
- [100] William H Greene and David A Hensher. “Revealing additional dimensions of preference heterogeneity in a latent class mixed multinomial logit model”. In: *Applied Economics* 45.14 (2013), pp. 1897–1902.
- [101] Torsten Grothmann and Fritz Reusswig. “People at risk of flooding: why some residents take precautionary action while others do not”. In: *Natural hazards* 38.1-2 (2006), pp. 101–120.

- [102] Richard W Hamming. “Error detecting and error correcting codes”. In: *Bell Labs Technical Journal* 29.2 (1950), pp. 147–160.
- [103] James K Hammitt. “Valuing Mortality Risk: Theory and Practice”. In: *Environmental Science and Technology* 34 (2000), pp. 1396–1400.
- [104] James K Hammitt and Kevin Haninger. “Willingness to pay for food safety: sensitivity to duration and severity of illness”. In: *American Journal of Agricultural Economics* 89.5 (2007), pp. 1170–1175.
- [105] Nick Hanley, Robert E Wright, and Wiktor Adamowicz. “Using choice experiments to value the environment”. In: *Environmental and resource economics* 11.3-4 (1998), pp. 413–428.
- [106] John A. Hartigan and M. Anthony Wong. “Algorithm AS 136: A k-means clustering algorithm”. In: *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28.1 (1979), pp. 100–108.
- [107] Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical learning with sparsity*. CRC press, 2015.
- [108] John P Helveston et al. “Will subsidies drive electric vehicle adoption? Measuring consumer preferences in the US and China”. In: *Transportation Research Part A: Policy and Practice* 73 (2015), pp. 96–112.
- [109] David A Hensher. “Attribute processing, heuristics and preference construction in choice analysis”. In: *Choice Modelling: The State-of-the-art and The State-of-practice: Proceedings from the Inaugural International Choice Modelling Conference*. Emerald Group Publishing Limited. 2010, pp. 35–69.
- [110] David A Hensher. “Stated preference analysis of travel choices: the state of practice”. In: *Transportation* 21.2 (1994), pp. 107–133.
- [111] David A Hensher and William H Greene. “Non-attendance and dual processing of common-metric attributes in choice analysis: a latent class specification”. In: *Empirical economics* 39.2 (2010), pp. 413–426.
- [112] David A Hensher, John M Rose, and William H Greene. *Applied choice analysis*. Cambridge University Press, 2015.
- [113] David A Hensher et al. “Estimating the willingness to pay and value of risk reduction for car occupants in the road environment”. In: *Transportation research part A: policy and practice* 43.7 (2009), pp. 692–707.
- [114] Joseph A Herriges and Catherine L Kling. “Nonlinear income effects in random utility models”. In: *Review of Economics and Statistics* 81.1 (1999), pp. 62–72.
- [115] Ralph Hertwig et al. “Decisions from experience and the effect of rare events in risky choice”. In: *Psychological science* 15.8 (2004), pp. 534–539. ISSN: 0956-7976.

- [116] Stephane Hess, Michel Bierlaire, and John Polak. “Mixed GEV modelling of mode-choice”. In: *Swiss Transportation Research Conference*. TRANSP-OR-CONF-2006-031. 2005.
- [117] Stephane Hess, Amanda Stathopoulos, and Andrew Daly. “Allowing for heterogeneous decision rules in discrete choice models: an approach and four case studies”. In: *Transportation* 39.3 (2012), pp. 565–591.
- [118] Wesley E Highfield, Sarah A Norman, and Samuel D Brody. “Examining the 100-Year Floodplain as a Metric of Risk, Loss, and Household Adjustment”. In: *Risk Analysis* 33.2 (2013), pp. 186–191.
- [119] Geoffrey E Hinton and Ruslan R Salakhutdinov. “Reducing the dimensionality of data with neural networks”. In: *science* 313.5786 (2006), pp. 504–507.
- [120] Ming-Chou Ho et al. “How do disaster characteristics influence risk perception?” In: *Risk Analysis* 28.3 (2008), pp. 635–643.
- [121] Robert Holmes and Karen Dinicola. *100-Year Flood - It's All About Chance*. Government Document. 2010. URL: <http://pubs.usgs.gov/gip/106/pdf/100-year-flood-handout-042610.pdf>.
- [122] Christopher K Hsee et al. “Preference reversals between joint and separate evaluations of options: a review and theoretical analysis.” In: *Psychological bulletin* 125.5 (1999), p. 576.
- [123] Francis KC Hui, Samuel Müller, and AH Welsh. “Joint selection in mixed models using regularized PQL”. In: *Journal of the American Statistical Association* 112.519 (2017), pp. 1323–1333.
- [124] Nicolas Jacquemet et al. “Preference elicitation under oath”. In: *Journal of Environmental Economics and Management* 65.1 (2013), pp. 110–132.
- [125] F Reed Johnson, Melissa Ruby Banzhaf, and William H Desvousges. “Willingness to pay for improved respiratory and cardiovascular health: a multiple-format, stated-preference approach”. In: *Health Economics* 9.4 (2000), pp. 295–317.
- [126] F Reed Johnson et al. “Eliciting stated health preferences: an application to willingness to pay for longevity”. In: *Medical Decision Making* 18.2_suppl (1998), S57–S67.
- [127] Peter Juslin, Marcus Lindskog, and Bastian Mayerhofer. “Is there something special with probabilities?â€”Insight vs. computational ability in multiple risk combination”. In: *Cognition* 136 (2015), pp. 282–303. ISSN: 0010-0277.
- [128] Daniel Kahneman. “Maps of bounded rationality: Psychology for behavioral economics”. In: *The American economic review* 93.5 (2003), pp. 1449–1475.
- [129] Daniel Kahneman and Amos Tversky. “Prospect theory: An analysis of decision under risk”. In: *Econometrica: Journal of the Econometric Society* (1979), pp. 263–291. ISSN: 0012-9682.

- [130] Alexandros Karatzoglou et al. “kernlab-an S4 package for kernel methods in R”. In: *Journal of statistical software* 11.9 (2004), pp. 1–20.
- [131] Roger E Kasperson et al. “The social amplification of risk: A conceptual framework”. In: *Risk analysis* 8.2 (1988), pp. 177–187.
- [132] Wim Kellens, Teun Terpstra, and Philippe De Maeyer. “Perception and communication of flood risks: a systematic review of empirical research”. In: *Risk analysis* 33.1 (2013), pp. 24–49.
- [133] Wim Kellens et al. “An analysis of the public perception of flood risk on the Belgian coast”. In: *Risk analysis* 31.7 (2011), pp. 1055–1068.
- [134] Carmen Keller, Michael Siegrist, and Heinz Gutscher. “The role of the affect and availability heuristics in risk communication”. In: *Risk analysis* 26.3 (2006), pp. 631–639. ISSN: 1539-6924.
- [135] Charles Kemp et al. “Learning systems of concepts with an infinite relational model”. In: *Proceedings of the national conference on artificial intelligence*. Vol. 21. 1. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999. 2006, p. 381.
- [136] Murphy Kevin. *Machine learning: a probabilistic perspective*. The MIT press, Cambridge, Massachusetts., 2012.
- [137] Barbel Knauper et al. “Motivation influences the underestimation of cumulative risk”. In: *Personality and Social Psychology Bulletin* 31.11 (2005), pp. 1511–1523. ISSN: 0146-1672.
- [138] Ethan T Knocke and Korine N Kolivras. “Flash flood awareness in southwest Virginia”. In: *Risk analysis* 27.1 (2007), pp. 155–169.
- [139] Daniela Knuth et al. “Risk Perception, Experience, and Objective Risk: A Cross-National Study with European Emergency Survivors”. In: *Risk analysis* 34.7 (2014), pp. 1286–1298.
- [140] Tjalling C Koopmans and Martin Beckmann. “Assignment problems and the location of economic activities”. In: *Econometrica: journal of the Econometric Society* (1957), pp. 53–76.
- [141] Alan Krupnick. “Mortality-risk valuation and age: stated preference evidence”. In: *Review of Environmental Economics and Policy* 1.2 (2007), pp. 261–282.
- [142] Mylene Lagarde. “Investigating attribute non-attendance and its consequences in choice experiments with latent class models”. In: *Health economics* 22.5 (2013), pp. 554–567.
- [143] Jason D Lee, Yuekai Sun, and Michael A Saunders. “Proximal Newton-type methods for convex optimization.” In: *NIPS*. Vol. 25. 2012, pp. 827–835.
- [144] Haim Levy and Harry M Markowitz. “Approximating expected utility by a function of mean and variance”. In: *The American Economic Review* 69.3 (1979), pp. 308–317.
- [145] Sarah Lichtenstein and Paul Slovic. “Reversals of preference between bids and choices in gambling decisions.” In: *Journal of experimental psychology* 89.1 (1971), p. 46.

- [146] Michael K Lindell and Seong Nam Hwang. “Households’ perceived personal risk and responses in a multihazard environment”. In: *Risk Analysis* 28.2 (2008), pp. 539–556.
- [147] Patricia W Linville, Gregory W Fischer, and Baruch Fischhoff. “AIDS risk perceptions and decision biases”. In: *The social psychology of HIV infection*. Ed. by J. B. Pryor G. D. Reeder. Hillsdale, NJ: Erlbaum, 1993, pp. 5–38. ISBN: 0805809910.
- [148] Jin-Tan Liu et al. “Valuation of the risk of SARS in Taiwan”. In: *Health Economics* 14.1 (2005), pp. 83–91.
- [149] Zhi-Yong Liu et al. “A graph matching algorithm based on concavely regularized convex relaxation”. In: *Neurocomputing* 134 (2014), pp. 140–148.
- [150] Lorenzo Livi and Antonello Rizzi. “The graph matching problem”. In: *Pattern Analysis and Applications* 16.3 (2013), pp. 253–283.
- [151] George F Loewenstein et al. “Risk as feelings”. In: *Psychological bulletin* 127.2 (2001), p. 267. ISSN: 1939-1455.
- [152] Alberto Longo, Anil Markandya, and Marta Petrucci. “The internalization of externalities in the production of electricity: willingness to pay for the attributes of a policy for renewable energy”. In: *Ecological Economics* 67.1 (2008), pp. 140–152.
- [153] Jordan J Louviere. *Analyzing decision making: Metric conjoint analysis*. 67. Sage, 1988.
- [154] Jordan J Louviere, David A Hensher, and Joffre D Swait. *Stated choice methods: analysis and applications*. Cambridge university press, 2000.
- [155] Antoine Lucas. “amap: Another Multidimensional Analysis Package”. In: [http://CRAN.R-project.org/package= amap](http://CRAN.R-project.org/package=amap) (2014).
- [156] Christopher G Lucas et al. “A rational model of function learning”. In: *Psychonomic bulletin & review* 22.5 (2015), pp. 1193–1215.
- [157] Vince Lyzinski et al. “Graph matching: Relax at your own risk”. In: *IEEE transactions on pattern analysis and machine intelligence* 38.1 (2016), pp. 60–73.
- [158] Laurens van der Maaten and Geoffrey Hinton. “Visualizing data using t-SNE”. In: *Journal of machine learning research* 9.Nov (2008), pp. 2579–2605.
- [159] Samer M Madanat, CY Yang, and Ying-Ming Yen. “Analysis of stated route diversion intentions under advanced traveler information systems using latent variable modeling”. In: *Transportation Research Record* 1485 (1995).
- [160] Jacob Marschak. *Binary choice constraints on random utility indicators*. Tech. rep. Cowles Foundation for Research in Economics, Yale University, 1959.
- [161] Peter Martinsson and Fredrik Carlsson. “Do experience and cheap talk influence willingness to pay in an open-ended contingent valuation survey?” In: *rapport nr.: Working Papers in Economics* 190 (2006).

- [162] Rachel McCloy, Ruth MJ Byrne, and Philip N Johnson-Laird. “Understanding cumulative risk”. In: *The Quarterly Journal of Experimental Psychology* 63.3 (2010), pp. 499–515. ISSN: 1747-0218.
- [163] Rachel McCloy et al. “Training conditional and cumulative risk judgements: the role of frequencies, problem structure and Einstellung”. In: *Applied Cognitive Psychology* 21.3 (2007), pp. 325–344. ISSN: 1099-0720.
- [164] Mark A McDaniel and Jerome R Busemeyer. “The conceptual basis of function learning and extrapolation: Comparison of rule-based and associative-based models”. In: *Psychonomic bulletin & review* 12.1 (2005), pp. 24–42.
- [165] Daniel McFadden. “Computing willingness-to-pay in random utility models”. In: *Trade: Theory and Econometrics*. Ed. by J. Moore, R. Reizman, and J. Melvin. Routledge: London, 1999, pp. 253–274.
- [166] Daniel McFadden. “Conditional logit analysis of qualitative choice behavior”. In: *Frontiers in Econometrics*. Ed. by P. Zarembka. Academic Press: New York, 1973, pp. 105–142.
- [167] Daniel McFadden. “Economic choices”. In: *American economic review* 91.3 (2001), pp. 351–378.
- [168] Daniel McFadden. *Foundations of Welfare Economics and Product Market Applications*. Tech. rep. National Bureau of Economic Research, 2017.
- [169] Daniel McFadden. “Modeling the choice of residential location”. In: *The Economics of Housing*. Ed. by J. Quigley. Vol. 1. 1997, pp. 531–552.
- [170] Nicolai Meinshausen. “Relaxed lasso”. In: *Computational Statistics & Data Analysis* 52.1 (2007), pp. 374–393.
- [171] Erwann Michel-Kerjan and Howard Kunreuther. “Redesigning flood insurance”. In: *Science* 333.6041 (2011), pp. 408–409.
- [172] Jihoon Min et al. “Labeling energy cost on light bulbs lowers implicit discount rates”. In: *Ecological Economics* 97 (2014), pp. 42–50.
- [173] Robert Cameron Mitchell and Richard T Carson. *Using surveys to value public goods: the contingent valuation method*. Resources for the Future, 1989.
- [174] John W Moon. *Topics on tournaments in graph theory*. Courier Dover Publications, 2015.
- [175] Yurii Nesterov. “A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$ ”. In: *Doklady an SSSR*. Vol. 269. 3. 1983, pp. 543–547.
- [176] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [177] Vilfredo Pareto. *Manuale di economia politica*. Vol. 13. Societa Editrice, 1906.

- [178] John W Payne, James R Bettman, and Eric J Johnson. *The adaptive decision maker*. Cambridge University Press, 1993.
- [179] Ellen Peters et al. “Numeracy and decision making”. In: *Psychological science* 17.5 (2006), pp. 407–413. ISSN: 0956-7976.
- [180] Wei Qian and Yuhong Yang. “Model selection via standard error adjusted adaptive lasso”. In: *Annals of the Institute of Statistical Mathematics* (2013), pp. 1–24.
- [181] Michel Regenwetter, Jason Dana, and Clinton P Davis-Stober. “Transitivity of preferences.” In: *Psychological Review* 118.1 (2011), pp. 42–56.
- [182] William Revelle and Maintainer William Revelle. “Package "psych"”. In: *The Comprehensive R Archive Network* (2015).
- [183] Luis I Rizzi and Cristobal De La Maza. “The external costs of private versus public road transport in the Metropolitan Area of Santiago, Chile”. In: *Transportation Research Part A: Policy and Practice* 98 (2017), pp. 123–140.
- [184] Luis I Rizzi and Juan de Dios Ortúzar. “Road safety valuation under a stated choice framework”. In: *Journal of Transport Economics and Policy (JTEP)* 40.1 (2006), pp. 69–94.
- [185] Luis I Rizzi and Juan de Dios Ortúzar. “Stated preference in the valuation of interurban road safety”. In: *Accident Analysis & Prevention* 35.1 (2003), pp. 9–22.
- [186] Luis I Rizzi et al. “Valuing air quality impacts using stated choice analysis: Trading off visibility against morbidity effects”. In: *Journal of environmental management* 146 (2014), pp. 470–480.
- [187] Kevin WS Roberts. “Interpersonal comparability and social choice theory”. In: *The Review of Economic Studies* 47.2 (1980), pp. 421–439.
- [188] Brian Roe et al. “US consumers’ willingness to pay for green electricity”. In: *Energy policy* 29.11 (2001), pp. 917–925.
- [189] Robert M Roe, Jermone R Busemeyer, and James T Townsend. “Multialternative decision field theory: A dynamic connectionist model of decision making.” In: *Psychological Review* 108.2 (2001), pp. 370–392.
- [190] John M Rose et al. “Ngene 1.2 stated choice experiment design software”. In: *University of Sydney*: <http://www.choice-metrics.com/NgeneManual120.pdf> (2018).
- [191] Michael Rothschild and Joseph Stiglitz. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information”. In: *The Quarterly Journal of Economics* 90.4 (1976), pp. 629–649.
- [192] Isabelle Ruin, Jean-Christophe Gaillard, and Celine Lutoff. “How to get there? Assessing motorists’ flash flood risk perception on daily itineraries”. In: *Environmental hazards* 7.3 (2007), pp. 235–244.
- [193] Ana M Ruz, Karen Ubilla, and Diego Luna. “Externalities and energy policy”. In: *Escenarios Energéticos-Chile 2030*: <http://escenariosenergeticos.cl> (2015).

- [194] Katharina Sammer and Rolf Wüstenhagen. “The influence of eco-labelling on consumer behaviour—Results of a discrete choice analysis for washing machines”. In: *Business Strategy and the Environment* 15.3 (2006), pp. 185–199.
- [195] Paul A Samuelson. “A note on the pure theory of consumer’s behaviour”. In: *Economica* 5.17 (1938), pp. 61–71.
- [196] Riccardo Scarpa et al. “Modelling attribute non-attendance in choice experiments for rural landscape valuation”. In: *European review of agricultural economics* 36.2 (2009), pp. 151–174.
- [197] Brian Sergi, Alexander Davis, and Inês Azevedo. “The effect of providing climate and health information on support for alternative electricity portfolios”. In: *Environmental Research Letters* 13.2 (2018), p. 024026.
- [198] Anuj K Shah and Daniel M Oppenheimer. “Heuristics made easy: An effort-reduction framework.” In: *Psychological bulletin* 134.2 (2008), p. 207.
- [199] Harriet Shaklee and Baruch Fischhoff. “The Psychology of Contraceptive Surprises: Cumulative Risk and Contraceptive Effectiveness¹”. In: *Journal of Applied Social Psychology* 20.5 (1990), pp. 385–403. ISSN: 1559-1816.
- [200] David F Shanno. “Conditioning of quasi-Newton methods for function minimization”. In: *Mathematics of computation* 24.111 (1970), pp. 647–656.
- [201] Chen Shaobing and David Donoho. “Basis pursuit”. In: *28th Asilomar conf. Signals, Systems Computers*. 1994.
- [202] Paschal Sheeran, Peter R Harris, and Tracy Epton. “Does heightening risk appraisals change people’s intentions and behavior? A meta-analysis of experimental studies”. In: *Psychological bulletin* 140.2 (2014), p. 511. ISSN: 1939-1455.
- [203] Bruce J Sherrick et al. “Farmers’ preferences for crop insurance attributes”. In: *Review of Agricultural Economics* 25.2 (2003), pp. 415–429.
- [204] Michael Siegrist and Heinz Gutscher. “Flooding risks: A comparison of lay people’s perceptions and expert’s assessments in Switzerland”. In: *Risk Analysis* 26.4 (2006), pp. 971–979.
- [205] Michael Siegrist and Heinz Gutscher. “Natural hazards and motivation for mitigation behavior: People cannot predict the affect evoked by a severe flood”. In: *Risk Analysis* 28.3 (2008), pp. 771–778.
- [206] Michael Siegrist, Heinz Gutscher, and Timothy C Earle. “Perception of risk: the influence of general trust, and general confidence”. In: *Journal of Risk Research* 8.2 (2005), pp. 145–156.
- [207] Mauricio Sillano and Juan de Dios Ortúzar. “Willingness-to-pay estimation with mixed logit models: some new evidence”. In: *Environment and Planning A* 37.3 (2005), pp. 525–550.

- [208] Joseph P Simmons, Leif D Nelson, and Uri Simonsohn. “False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant”. In: *Psychological science* 22.11 (2011), pp. 1359–1366.
- [209] Herbert A Simon. “Rational choice and the structure of the environment”. In: *Psychological Review* 63.2 (1956), pp. 129–138. DOI: 10.1037/h0042769.
- [210] Herbert A Simon. “Theories of bounded rationality”. In: *Decision and Organization* 1.1 (1972), pp. 161–176.
- [211] Itamar Simonson and Amos Tversky. “Choice in context: Tradeoff contrast and extremeness aversion.” In: *Journal of Marketing Research* 29.3 (1992), pp. 281–295.
- [212] Archana Singh, Avantika Yadav, and Ajay Rana. “K-means with Three different Distance Metrics”. In: *International Journal of Computer Applications* 67.10 (2013).
- [213] Paul Slovic. “Choice between equally valued alternatives.” In: *Journal of Experimental Psychology: Human Perception and Performance* 1.3 (1975), p. 280.
- [214] Paul Slovic. “Manipulating the attractiveness of a gamble without changing its expected value”. In: *Journal of Experimental Psychology* 79.1p1 (1969), p. 139. ISSN: 0022-1015.
- [215] Paul Slovic. “What does it mean to know a cumulative risk? Adolescents’ perceptions of short-term and long-term consequences of smoking”. In: *Journal of Behavioral Decision Making* 13.2 (2000), pp. 259–266. ISSN: 1099-0771.
- [216] Paul Slovic, Baruch Fischhoff, and Sarah Lichtenstein. “Accident probabilities and seat belt usage: A psychological perspective”. In: *Accident Analysis Prevention* 10.4 (1978), pp. 281–285. ISSN: 0001-4575.
- [217] Paul Slovic et al. “The affect heuristic”. In: *European journal of operational research* 177.3 (2007), pp. 1333–1352. ISSN: 0377-2217.
- [218] Kenneth A Small and Harvey S Rosen. “Applied welfare economics with discrete choice models”. In: *Econometrica* 49.1 (1981), pp. 105–130.
- [219] Adam B Smith and Richard W Katz. “US billion-dollar weather and climate disasters: data sources, trends, accuracy and biases”. In: *Natural hazards* 67.2 (2013), pp. 387–410. ISSN: 0921-030X.
- [220] Johannes Spinnewijn. “Insurance and perceptions: how to screen optimists and pessimists”. In: *The Economic Journal* 123.569 (2013), pp. 606–633. ISSN: 1468-0297.
- [221] David E Sprott, David M Hardesty, and Anthony D Miyazaki. “Disclosure of odds information: An experimental investigation of odds format and numeric complexity”. In: *Journal of Public Policy Marketing* (1998), pp. 11–23. ISSN: 0743-9156.
- [222] Joseph E Stiglitz. “The theory of screening, education, and the distribution of income”. In: *The American economic review* 65.3 (1975), pp. 283–300. ISSN: 0002-8282.
- [223] Joakim Sundh and Peter Juslin. “Compound risk judgment in tasks with both idiosyncratic and systematic risk: The Robust Beauty of additive probability integration”. In: *Cognition* 171 (2018), pp. 25–41. ISSN: 0010-0277.

- [224] Thomas Sundqvist. “Quantifying household preferences over the environmental impacts of hydropower in Sweden: A choice experiment approach”. In: *Power Generation Choice in the Presence of Environmental Externalities, Doctoral Thesis, Lulea University of Technology* (2002).
- [225] Ola Svenson. “Cognitive processes in judging cumulative risk over different periods of time”. In: *Organizational behavior and human performance* 33.1 (1984), pp. 22–41. ISSN: 0030-5073.
- [226] Ola Svenson. “Cognitive strategies in a complex judgment task: Analyses of concurrent verbal reports and judgments of cumulated risk over different exposure times”. In: *Organizational Behavior and Human Decision Processes* 36.1 (1985), pp. 1–15. ISSN: 0749-5978.
- [227] Ola Svenson. “Towards a framework for human judgments of quantitative information: the numerical judgment process, NJP model”. In: *Journal of Cognitive Psychology* (2016), pp. 1–15. ISSN: 2044-5911.
- [228] Kenji Takao et al. “Factors determining residents’ preparedness for floods in modern megalopolises: the case of the Tokai flood disaster in Japan”. In: *Journal of Risk Research* 7.7-8 (2004), pp. 775–787.
- [229] Karl H Teigen. “Overestimation of subjective probabilities”. In: *Scandinavian Journal of Psychology* 15.1 (1974), pp. 56–62. ISSN: 1467-9450.
- [230] Teun Terpstra. “Emotions, trust, and perceived risk: Affective and cognitive routes to flood preparedness behavior”. In: *Risk analysis* 31.10 (2011), pp. 1658–1675.
- [231] Richard H Thaler and Cass Sunstein. “Nudge”. In: *Improving Decisions about Health, Wealth, and Happiness* (2008).
- [232] Annegret H Thielen et al. “Insurability and mitigation of flood losses in private households in Germany”. In: *Risk Analysis* 26.2 (2006), pp. 383–395.
- [233] Robert Tibshirani. “Regression shrinkage and selection via the lasso”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), pp. 267–288.
- [234] Robert Tibshirani, Guenther Walther, and Trevor Hastie. “Estimating the number of clusters in a data set via the gap statistic”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63.2 (2001), pp. 411–423.
- [235] Kenneth E Train. *Discrete choice methods with simulation*. Cambridge university press, 2009. ISBN: 1139480375.
- [236] Efthymios Tsakonas et al. “Sparse conjoint analysis through maximum likelihood estimation”. In: *IEEE Transactions on Signal Processing* 61.22 (2013), pp. 5704–5715.
- [237] Konstantinos Tsetsos, Nick Chater, and Marius Usher. “Salience driven value integration explains decision biases and preference reversal”. In: *PNAS* 109.24 (2012), pp. 9659–9664.

- [238] Amos Tversky. “Choice by elimination”. In: *Journal of Mathematical Psychology* 9.4 (1972), pp. 341–367.
- [239] Amos Tversky. “Intransitivity of preferences.” In: *Psychological Review* 76.1 (1969), pp. 31–48.
- [240] Amos Tversky and Craig R Fox. “Weighing risk and uncertainty”. In: *Psychological review* 102.2 (1995), p. 269. ISSN: 1939-1471.
- [241] Amos Tversky and Daniel Kahneman. “Advances in prospect theory: Cumulative representation of uncertainty”. In: *Journal of Risk and uncertainty* 5.4 (1992), pp. 297–323. ISSN: 0895-5646.
- [242] Amos Tversky and Daniel Kahneman. “Availability: A heuristic for judging frequency and probability”. In: *Cognitive Psychology* 5.2 (1973), pp. 207–232.
- [243] Amos Tversky and Daniel Kahneman. “Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment”. In: *Psychological review* 90.4 (1983), p. 293. ISSN: 1939-1471.
- [244] Amos Tversky and Daniel Kahneman. “Judgment under uncertainty: Heuristics and biases”. In: *science* 185.4157 (1974), pp. 1124–1131. ISSN: 0036-8075.
- [245] Amos Tversky and Daniel Kahneman. “Loss aversion in riskless choice: A reference-dependent model”. In: *The Quarterly Journal of Economics* 106.4 (1991), pp. 1039–1061.
- [246] Amos Tversky and Daniel Kahneman. “The framing of decisions and the psychology of choice”. In: *Science* 211.4481 (1981), pp. 453–458.
- [247] Amos Tversky and Itamar Simonson. “Context-dependent preferences”. In: *Management Science* 39.10 (1993), pp. 1179–1189.
- [248] Amos Tversky and Richard H Thaler. “Anomalies: preference reversals”. In: *Journal of Economic Perspectives* 4.2 (1990), pp. 201–211.
- [249] Hal R Varian. “Non-parametric tests of consumer behaviour”. In: *The review of economic studies* 50.1 (1983), pp. 99–110.
- [250] Hal R Varian. “The nonparametric approach to demand analysis”. In: *Econometrica* 50.4 (1982), pp. 945–973.
- [251] Vinod Venkatraman, John W Payne, and Scott A Huettel. “An overall probability of winning heuristic for complex risky decisions: Choice and eye fixation evidence”. In: *Organizational Behavior and Human Decision Processes* 125.2 (2014), pp. 73–87.
- [252] Bernardita Vial and Felipe Zurita. *Microeconomia intermedia*. Pontificia Universidad Catolica de Chile, Instituto de Economia, Oficina de Publicaciones, 2006.
- [253] W Kip Viscusi and Joseph E Aldy. “The value of a statistical life: a critical review of market estimates throughout the world”. In: *Journal of risk and uncertainty* 27.1 (2003), pp. 5–76.

- [254] Vivianne HM Visschers et al. “Probability information in risk communication: a review of the research literature”. In: *Risk Analysis* 29.2 (2009), pp. 267–287. ISSN: 1539-6924.
- [255] Joshua T Vogelstein et al. “Fast approximate quadratic programming for graph matching”. In: *PLOS one* 10.4 (2015), e0121002.
- [256] John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Vol. 60. Princeton: Princeton university press, 1944.
- [257] Gisela Wachinger et al. “The risk perception paradox—implications for governance and communication of natural hazards”. In: *Risk analysis* 33.6 (2013), pp. 1049–1065.
- [258] Mei Wang and Paul S Fischbeck. “Incorporating framing into prospect theory modeling: a mixture-model approach”. In: *Journal of Risk and Uncertainty* 29.2 (2004), pp. 181–197. ISSN: 0895-5646.
- [259] Yasi Wang, Hongxun Yao, and Sicheng Zhao. “Auto-encoder based dimensionality reduction”. In: *Neurocomputing* 184 (2016), pp. 232–242.
- [260] Larry Wasserman. *All of statistics: a concise course in statistical inference*. Springer Science & Business Media, 2013.
- [261] Larry Wasserman and Ryan Tibshirani. “Lecture notes Statistical Machine Learning 10-702”. In: *Carnegie Mellon University* (2017).
- [262] Elke U Weber, Ann-Renee Blais, and Nancy E Betz. “A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors”. In: *Journal of behavioral decision making* 15.4 (2002), pp. 263–290. ISSN: 1099-0771.
- [263] Neil D Weinstein. “Optimistic biases about personal risks”. In: *Science* 246.4935 (1989), pp. 1232–1234.
- [264] Huw Williams. “On the formation of travel demand models and economic evaluation measures of user benefit”. In: *Environment and planning A* 9.3 (1977), pp. 285–344.
- [265] Ruud Zaalberg et al. “Prevention, adaptation, and threat denial: Flooding experiences in the Netherlands”. In: *Risk Analysis* 29.12 (2009), pp. 1759–1778.
- [266] Tomasz Zaleskiewicz, Zbigniew Piskorz, and Anna Borkowska. “Fear or money? Decisions on insuring oneself against flood”. In: *Risk, Decision and Policy* 7.3 (2002), pp. 221–233.
- [267] Cun-Hui Zhang et al. “Nearly unbiased variable selection under minimax concave penalty”. In: *The Annals of statistics* 38.2 (2010), pp. 894–942.
- [268] Hui Zou. “The adaptive lasso and its oracle properties”. In: *Journal of the American statistical association* 101.476 (2006), pp. 1418–1429.
- [269] Hui Zou, Trevor Hastie, and Robert Tibshirani. “On the “degrees of freedom” of the lasso”. In: *The Annals of Statistics* 35.5 (2007), pp. 2173–2192.

7

Appendix: Cumulative risk perception from judgments and choices




Figure 7.1: Expected loss for each option and risk premium under the assumption of a constant relative risk aversion (CRRA) utility function with $r=0.5$ normalized between 0 and 1 ($U(x) = (X^{1-r}(1-r))/(W_0^{1-r}(1-r))$). Utility function and certainty equivalents for option A (no coverage) with a 4%, 2% and 1% chance of a flood each year (A'', B'' and C''), and expected loss for the same gambles (A', B' and C'). If we assume participants are risk-averse, participants should be willing to pay their expected loss (EL) plus an additional risk premium (RP) to be protected. The red dots represent utility of option A (no coverage) at different risk levels and the white dots represent utility of option B (full coverage) at different insurance premiums assuming they decide based on final wealth levels.




Figure 7.2: Indifference curves for individuals with accurate judgment, additive and mean heuristics for a 2% chance flood each year. Each axis represents welfare in the state of the world where either no flood or flood occurs. The x-axis represents final wealth in the state of the world where there are no floods, or equivalently the agent initial wealth minus the cost of the insurance policy. The y-axis is final wealth level in the state of the world where a flood occurs, or equivalently initial wealth minus monetary loss plus insurance coverage minus the cost of the insurance. When offered a fair insurance premium a risk-averse individual with accurate judgment should prefer full insurance coverage, illustrated at the certainty line where indifference curve intersects actuarially fair insurance line [252]. All contracts to the right of the dotted line should be preferred to no insurance coverage. Thus, contracts represented by points A (\$150,000), B (\$90,000) or C (\$45,000) should be preferred to contract D (no coverage). The same applies for an individual following an additive heuristic. Contracts A, B or C fall left of a mean heuristic individual indifference curve and henceforth should be worst than contract D.

8

Appendix: Learning preference structure with graph matching

Table 8.1: Problems used for the choice experiments. See also problem ID in adjacency matrix in Table 8.2

(a) Transitivity task based on [239]. Gambles from Tversky’s classic paper on intransitive preferences [239] (a-e), along with five additional gambles (f-j), where last two gambles a higher probability is negatively correlated with a higher expected value.

Gamble	Probability	Payoff	Expected Value (\$)
a	7/24	5.00	1.458
b	8/24	4.75	1.583
c	9/24	4.50	1.688
d	10/24	4.25	1.771
e	11/24	4.00	1.833
f	12/24	3.75	1.875
g	13/24	3.50	1.894
h	14/24	3.25	1.895
i	15/24	3.00	1.875
j	16/24	2.75	1.833

(b) Anomalies task Gambles based on [129, 60]. In the column labeled as P, we present the proportion of participants with anomalous behavior. Proportion P, should be above 0.5 in all cases to match previous findings [60].

Alt	p.1	o.1	o.2	Anomalies	P
a	1	3,000	0	Certainty (a \succ b)	0.7
b	0.80	4,000	0	Certainty (a \succ b)	
c	0.25	3,000	0	Certainty (d \succ c)	
d	0.20	4,000	0	Certainty (d \succ c)	0.3
e	1	-3,000	0	Reflection (f \succ e)	
f	0.8	-4,000	0	Reflection (f \succ e)	0.3
g	1	50	0	Overweight (h \succ g)	
h	0.01	5,000	0	Overweight (h \succ g)	0.3
i	1	0	0	Loss aversion (i \succ j)	0.7
j	0.5	1,000	-1,000	Loss aversion (j \succ i)	
k	0.5	6,000	0	Risk aversion (a \succ k)	0.3

(c) Alternatives CO2 task (CO_2 %, Bill %) based on [197]

a (-30%, 20%), b (-30%, 5%), c (-25%, 4%), d (-20%, 3%),
e (-15%, 2%), f (-10%, 1%), g (30%, -20%), h (30%, -5%)

Table 8.2: Problem ID in adjacency matrix.

1	2	3	4	5	6	7	8	1	2	3	4	5	6	7	8	9	10	1	2	3	4	5	6	7	8	9	10																				
1	1	2	3	4	5	6	7	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45				
8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55

Table 8.3: Linear utility models per cluster transitivity task. $l(s)$: log-likelihood model with a single parameter, $l(EV)$: log-likelihood model expected value rule, $P(p)$: proportion choosing the alternative with a higher probability of winning.

Cluster	Content	$\hat{\beta}$	$l(s)$	$l(EV)$	$P(p)$	N (%)
1	Probs	10***	-542	-621	87%***	26 (13%)
2	Probs	23***	-473	-773	96%***	44 (22%)
3	Probs	111***	-30	-717	100%***	59 (30%)
4	Payoff	3***	-284	-382	7%***	18 (9%)
5	Probs	5***	-558	-586	71%***	20 (10%)
6	Probs	0.3*	-1,029	-1,029	54%***	33 (16%)

*p<0.1; **p<0.05; ***p<0.01

Table 8.4: Willingness to pay per cluster CO_2 task

Cluster	Content	$\hat{\beta}_{CO_2}$	$\hat{\beta}_{Bill}$	WTP	N
1	Bills	1.5**	-12.5***	-0.04	37 (18%)
2	Bills	-5.1***	-9.1***	0.17	17 (8%)
3	Bills	-7.9***	-13.8***	0.17	26 (13%)
4	Greens	-9.2***	-3.4	0.81	12 (6%)
5	Greens	-15.6***	-20.4***	0.23	28 (14%)
6	Greens	-27.9***	-37.7***	0.22	50 (25%)
7	Cycles	-0.2	-3.4***	0.02	30 (15%)

Table 8.5: Latent decision rule model anomalies task

<i>Dependent variable: Choice</i>		
Rule 1 (EV)	Expected value	30.1* (17.8)
	Standard deviation	-30.3 (18.7)
Rule 2 (Better)	P(Better outcome)	1.8*** (0.2)
Rule 3 (Weighted additive)	Outcome low	-7.5*** (1.1)
	Outcome high	0.9** (0.2)
Rule 4 (Probability matching)	Probability low	3.1*** (0.9)
Rule 1	($\pi_1 = 15\%$)	"_"
Rule 2	($\pi_2 = 38\%$)	0.9*** (0.2)
Rule 3	($\pi_3 = 36\%$)	0.9*** (0.2)
Rule 4	($\pi_4 = 11\%$)	-0.3 (0.3)
Observations	11,000	
Log Likelihood	-7,052	

Note: *p<0.1; **p<0.05; ***p<0.01

Table 8.6: Proportion matching expected anomalies per cluster. (1) Certainty effect, (2) Reflection effect, (3) Rare event overweighting, (4) Loss aversion, (5) Risk aversion

Cluster	(1)	(2)	(3)	(4)	(5)	N
All	0.2	0.2	0.3	0.7	0.7	200 (100%)
1	0.2	0.3	0.4	0.6	0.6	89 (45%)
2	0.2	0.1	0.3	0.7	0.8	111 (55%)

9

Appendix: Welfare analysis using the sparse multinomial logit model




Figure 9.1: Correlation matrix survey

Table 9.1: Summary statistics

Variable	n	mean	sd	median	mad	min	max	skew
Have seen hydro	488	0.34	0.47	0	0	0	1	0.66
Have seen power	488	0.17	0.38	0	0	0	1	1.75
Have seen renewable	488	0.31	0.46	0	0	0	1	0.80
Signed oath	488	0.49	0.50	0	0	0	1	0.02
Public support	483	0.27	0.26	0.12	0.11	0.05	1	1.20
Gender (female)	486	0.56	0.50	1	0	0	1	-0.23
Age	474	49	16	48	16	17	90	0.13
Ethnicity	485	0.04	0.20	0	0	0	1	4.47
NGO	485	0.01	0.11	0	0	0	1	8.80
Electricity bill (USD/month)	479	47	27	41	20	4	263	2.37
Income (USD/month)	475	1,146	843	940	836	188	3,759	1.02
Others responsibility (Values)	437	0	1	0.05	1.18	-3.10	1.56	-0.54
Individual responsibility (Values)	437	0	1	0.36	0.99	-2.39	1.28	-0.72
Pro-social (Values)	437	0	1	0.19	1.41	-1.61	1.30	-0.24
Altruistic (Values)	437	0	1	0.27	0.82	-3.81	1.33	-1.07
Efficient (Behavior)	455	0	1	0.20	0.97	-2.72	1.32	-0.87
Green (Behavior)	455	0	1	0.04	1.08	-2.06	1.75	-0.14
Activist (Behavior)	455	0	1	-0.27	0.95	-1.48	2.71	0.85
Biocentric (Beliefs)	444	0	1	0.04	1.02	-2.24	2.64	0.13
Economicist (Beliefs)	444	0	1	0.15	1.12	-3.29	1.54	-0.56
Trust in technology (Beliefs)	444	0	1	0.27	0.85	-3.64	1.85	-0.86
Trust in government (integrity)	467	0	1	-0.02	1.16	-1.84	1.71	-0.21
Trust in government (competence)	467	0	1	-0.03	0.97	-2.99	1.94	-0.42
Family	485	0.12	0.33	0	0	0	1	2.28
Children	488	0.64	0.48	1	0	0	1	-0.56
Elder	488	0.34	0.47	0	0	0	1	0.68
Past visits	488	0.20	0.40	0	0	0	1	1.49
Future visits	484	0.33	0.47	0	0	0	1	0.70

Table 9.2: Factor analysis qualitative scales

Statement (Scale 1 to 7)	Factor	α	MR1	MR2	MR3	MR4
Trust in government						
The government has the knowledge and human resources to reduce environmental impacts	Competence	1		0.41		
The government will act without political or industry pressures to regulate electricity generation	Integrity	0.8	0.58	0.23		
The government will take into account the interests of citizens and the environment to regulate electricity generation			0.78	-0.14		
The government has sufficient judgment to make good decisions to regulate electricity generation			0.78			
Social values						
The industry is responsible for reducing environmental impacts of electricity generation and not households like mine	Resp. Other	0.8		0.63		
The government is responsible for reducing environmental impacts of electricity generation and not households like mine				0.94		0.13
All households use electricity and should contribute to reduce the impacts of its generation	Resp. Ind.	0.7	0.77			0.21
Households like mine should not be responsible for reducing the environmental impacts of electricity generation*			0.45	-0.17		
It is my responsibility as a citizen to help reduce the environmental impacts of electricity generation			0.65			
I am not willing to collaborate to reduce the environmental impacts of electricity generation if others are not*			0.48	-0.10		
It is my duty to help others when they can not help themselves			0.49	0.11		-0.22
Each individual is solely responsible for its well-being in life*	Pro-social	0.4	0.13		0.31	
My only responsibility is to provide welfare for my family and I*			-0.31	-0.16	0.99	0.16
Donations rarely improve the lives of others*	Altruistic	0.2	0.10			0.30
My personal actions can improve the life of people I do not know			0.34			-0.46
Environmental beliefs						
Environmental problems can be solved with better technology*	Ecologist	0.6		0.73	-0.61	
Plants and animals have the same right to live than humans				0.67		
Humans are depleting the limited resources available on the planet				0.64		
People worry too much about the impact of economic development	Economicist	0.7	0.78	0.20	0.15	
Humans have the right to modify the natural environment			0.54			
We care too much about the impacts on the environment and not enough to create jobs today			0.51		-0.12	
Nature is strong enough to withstand the impact of our lifestyle*			-0.50		0.19	
The deterioration of the environment is not as bad as they say*	Technology	0.5	-0.22	0.27	0.42	
There are more important things life than the environment*			-0.24	0.21	0.25	
Environmental behavior						
I keep electric appliances unplugged when not using them	Efficient	0.7	0.86			
I take shorter showers to save water			0.67			
I turn off lights when leaving a room			0.65			
I save fuel biking or walking			0.31	0.18	0.20	
I recycle newspapers, glass bottles and other items	Green	0.7	-0.15	0.93		
I prefer environmentally friendly products				0.63		
I participate in protests for the protection of the environment	Activist	0.5		-0.10	0.73	
I donate money to environmental protection					0.53	

Table 9.3: Environmental impacts per technology

Source	Use of Space (ha/MW)	Forest (ha/MW)	CO _{2e} (Ton/GWh)	MP (Kg/GWh)	NO _x (Kg/GWh)	SO ₂ (Kg/MWh)
Biomass	0.1	0	24	170	1,100	80
Coal	0.7	0	1,001	90	610	610
Diesel	0.01	0	779	50	300	20
LNG	0.01	0	436	20	200	10
Wind	0.6	0	11	0	0	0
Geothermic	0.4	0	28	0	0	0
Dam	1.5	1.500	28	0	0	0
Run-of-the-river	0.2	0.200	7.200	0	0	0
Solar	1.4	0	48	0	0	0

Personal communication with experts from escenariosenergeticos.cl (2014) (Based on www.e-seia.cl)

Table 9.4: Private costs per technology

Source	Investment (USD/kW)	Life use (years)	Capacity (%)	Investment (USD/MWh)	O&M (USD/MWh)	Fuel (USD/MWh)	Total (USD/MWh)
Biomass	3,610	40	0.6	70	21	23	114
Coal	2,078	35	0.85	29	3	38	70
Diesel	500	25	0.85	7	15	144	166
LNG	893	25	0.9	13	3	81	97
Wind	1,945	20	0.4	65	12	0	77
Geothermic	3,896	40	0.85	54	13	0	67
Dam	2,000	45	0.55	42	5	0	47
Run-river	2,000	45	0.53	43	5	0	48
Solar	2,110	25	0.24	110	11	0	121

Personal communication with experts from escenariosenergeticos.cl (2014)