A graph-based model to discover preference structure from choice data

Cristobal De La Maza (cdelamaz@andrew.cmu.edu)[1]
Alex Davis (alexdavis@cmu.edu)[1]
Cleotilde Gonzalez (coty@cmu.edu)[2]
Ines Azevedo (iazevedo@cmu.edu)[1]

[1]Department of Engineering and Public Policy
[2]Department of Social and Decision Sciences
Carnegie Mellon University, Pittsburgh, PA 15213 USA

Abstract

Theories of preferences are concerned with axioms that describe how people decide or the structure of preference. Although the dominant paradigm has been the set of axioms consistent with utility maximization, several descriptive theories have been proposed leading to alternative preference structures, with better results predicting observed choices. Researchers have developed precise tests of their proposed models, with that testing limited to a priori defined patterns. While this approach is promising and theory-driven, it potentially misses structures not previously considered. As a result, patterns of choices are classified as either fitting a known model or not, with little analysis of the latter. What is needed is an approach that can determine preference structure from choice data even when those data are inconsistent with prior models, suggesting new structures to cognitive researchers or confirming old ones. In this paper we demonstrate how to use graph matching to uncover heterogeneity in the structure of preference across a population of decision-makers. We propose a novel non-parametric approach to formally capture the concept of preference structure using preference graphs, thereafter clustering decision-makers based on graph embedding methods. We explore the approach with simulated choice data from the most common classes of economic and psychological models. We also apply the approach to new empirical implementations of classic experiments in decisions between risky prospects and other stated preferences tasks. The approach uncovers heterogeneity in preference structure across a variety of dimensions, without requiring any prior knowledge of those structures.

Keywords: Heuristics; Preference Structure; Graph Matching; Clustering; Transitivity

Introduction

The study of preferences and the concept of rational choice have been relevant for cognitive science and artificial intelligence research from the early conception of the fields (H. A. Simon, 1956). Theories of preferences are often concerned with the invariant axioms that describe how people decide, or the *structure* of preference. For over a century, the dominant paradigm 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 (Bentham, 1879). This paradigm requires decision-makers to be able to consistently rank any set of alternatives that they come across (Pareto, 1906). This well-behaved description of preference may work well in simple environments, but it is not clear how accurately represents preferences in complex, naturalistic settings. Humans are largely heterogeneous, with preferences that vary over time, and are often inconsistent with their own preferences (Tsetsos, Chater, & Usher, 2012). Large amounts of data are currently available that document choices people make in naturalistic settings. For example, there exists information about purchase decisions, movie selections, and transportation patterns. The rise of automation technologies also poses a challenge to decision researchers on how to mimic accurately human performance. This trend calls for new ways to determine insights from human preferences in the presence of large heterogeneity of naturalistic choices (McFadden, 2001).

The axioms that define well-behaved preferences are both simple and quite powerful (Von Neumann & Morgenstern, 1944; Arrow, 1951). The most relevant ones state that first, all alternatives must be comparable, making the preference relation *complete*. Second, preferences must be *transitive*. With these conditions 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 that corresponds to that ranking. Over the years, this dominant paradigm has not gone without challenge. Researchers in the decision sciences have found that, in many circumstances, preferences are not always well-behaved. Many descriptive theories have proliferated to explain deviations of human behavior from utility maximization. This includes ground-breaking work on bounded rationality, where decision-makers use short-cuts to deal with the limits of human cognitive capacities (H. A. Simon, 1972). For example, the cognitive burden of selecting the best alternative, considering all potential costs and benefits of each alternative, is at best psychologically implausible (Fischhoff, 2005).

Instead, humans use simple rules or heuristics to cope with the cognitive burden of selecting the best alternative (Payne, Bettman, & Johnson, 1993; H. A. Simon, 1972; Gigerenzer, Todd, ABC Research Group, et al., 1999). For example, one psychologically plausible way to deal with complex choices is to simplify the task by choosing based on the attribute that is most important to the decision-maker, only examining other attributes if alternatives are sufficiently close on that attribute to be psychologically "tied". Tversky's lexicographic semiorder is such a process (Tversky, 1969) and can lead to intransitive behavior. Heuristics are designed to reduce effort during the choice process (Shah & Oppenheimer, 2008).

While sub-optimal, heuristics can work surprisingly well under time constraints (Gigerenzer et al., 1999). Human heuristic ability has been signaled as the cornerstone of our superior problem-solving skills (H. Simon & Chase, 1988; Anderson, 1996). As it has been long proposed, recognizing the atomic components of this heuristics can potentially uncover ways of programming efficient machines with human-like performance in natural environments (Anderson & Lebiere, 2014; Slagle, 1971). Although recently it has been proposed that current computer resources allow for the design of systems that use rational strategies in real time (Gershman, Horvitz, & Tenenbaum, 2015), heuristics can still provide significant savings in terms of computation costs and reaction time.

However, there is a blind spot in choice modelling research that limits our possibilities to discover heuristic structures (Maturana & Varela, 1987; Ramachandran, 1992). A blind spot is a region in the retina insensitive to light (Dictionary, 1989). A simple experiment to demonstrate the presence of a blind spot involves looking closely to an image with a black cross on left and a black circle on right, while covering the left eye with the left hand. As we focus our attention on the black cross and get closer to the image, the black circle disappears (Maturana & Varela, 1987). The blind spot in choice modelling is that descriptive theories are apriori defined. Researchers develop precise tests of their proposed models, with that testing limited to a priori defined patterns. As a result, patterns of choices are classified as either fitting a known model or not. For example, thus far structural tests of preference have been limited to specific patterns known a-priori, such as weak stochastic transitivity or the triangular condition (Regenwetter, Dana, & Davis-Stober, 2010). While this approach is promising and theory-driven, it potentially misses structures not previously considered. Some decision process are clearly identifiable a-priori, others might not. We are looking at the choice process too closely and at the same time partially blocking our sight by using tools that are not general enough. What is needed is an approach that can determine preference structure from choice data even when those data are inconsistent with prior models, suggesting at the same time new structures to psychological researchers or confirming old ones, and lending strength to welfare analysis or undermining it.

Next we describe our approach. The current research proposes a novel non-parametric model to formally capture the concept of preference structure using preference graphs, clusters decision-makers based on that structure, and can represent types of preferences currently not possible in existing frameworks (e.g. incomparability (Von Neumann & Morgenstern, 1944)). Because the approach clusters decision-makers with the same structural pattern of preferences, we provide unified method that may account for disparate preference patterns. The paper is structured as follows: we first present our graph-based model; then we test the method in simulations and in new empirical implementations of classic experiments in decisions between risky prospects and other stated preferences tasks; finally we discuss our results and present limitations of the method.

Discovering preference structure heterogeneity

Preference representation as graphs

Both classical utility models and newer descriptive theories imply specific patterns of choices, or *preference structures*. In this work, we exploit the idea that preference structures can be represented as *preference graphs* (Bouyssou & Vincke, 2010). For example, classical utility maximization can be represented as a *completely connected chain* (Varian, 1983; Afriat, 1972). This is, of course, not the only preference structure. For example, a lexicographic semiorder results in cyclic preferences when decision-makers change the weights they apply to attributes of alternatives (Tversky, 1969).

Graphs are a general way to represent binary relations among elements of discrete sets, including preference relations (Diestel, 2000; Varian, 1983; Aiolli & Sperduti, 2010; Bouyssou & Vincke, 2010). Consider a graph $G = (V, E)$ with vertex set V and edge set E. In a *preference graph* the vertices are interpreted as alternatives and edges as binary relations between alternatives where, for all pairs of alternatives, one and only one of the following three edges exists between them (Bouyssou & Vincke, 2010): i) if $a \succ b$, the decision maker strictly prefers a over b, then $a \rightarrow b$ and not $b \rightarrow a$ (*strict preference* or $a\ P\ b$). If instead, $b \succ a$, the decision maker strictly prefers b over a (bPa), then $b \rightarrow a$ and not $a \rightarrow b$; ii) If $a \sim b$, the decision maker is indifferent between a and b, then $a - b$ are connected by an undirected edge (*indifference* or $a\ I\ b$). This can also be represented as a is preferred to b and b to a or $a \leftrightarrow b$; and iii) If a is incomparable with b, then no edge between a and b exists (*incomparability* or $a\ J\ b$). Figure 1 describes a graph representation of preferences.
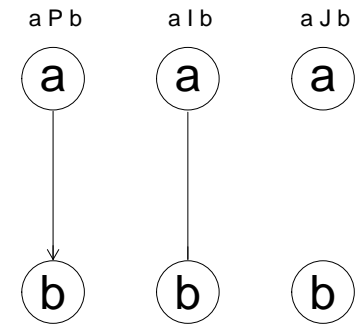

Figure 1: Preference relations in binary choice

An equivalent representation is an adjacency matrix $A = (a_{ij}) \in \{0,1\}^{n \times n}$ where $a_{ij} = 1$ if $(i,j) \in E_A$ and $B = (b_{ij}) \in \{0,1\}^{n \times n}$ where $b_{ij} = 1$ if $(i,j) \in E_B$, indicating preference from i to j. Reflexive loops are usually omitted, meaning the main diagonal of the adjacency matrix has only zeros. In this work we focus on a particular type of preference graphs, namely tournaments, where every alternative is compared and only strict preference is allowed giving a complete di-

rected graph (Gross & Yellen, 2004). The number of vertices in a tournament indicates the order. In the simplest case we find transitive *tournaments* (Moon, 2015), where all relations are strict preferences and there are no cycles. Following (Bouyssou & Vincke, 2010), consider a total order giving a tournament with an adjacency matrix that will show only zeros below the diagonal. A weak order instead, will allow indifference between alternatives and hence giving a tournament with a stepped shape adjacency matrix below the diagonal. For an irrational decision maker, cycles will be observed, giving a tournament with an adjacency matrix that will show elements above and below the diagonal.
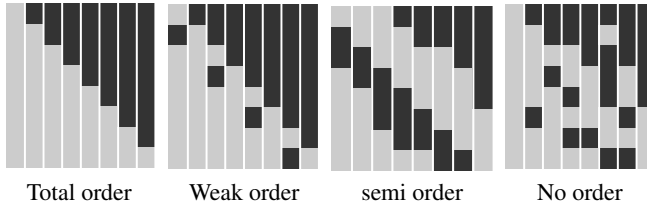

Figure 2: Tournaments with different order structure. Adjacency matrices are colored to ease interpretation with ones in black and zeros in grey

Another representation is a format used by Moon (Moon, 2015), where graphs are drawn based on their *score vector*, which is the number of times each alternative is preferred over other alternatives. For example, 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*. It is drawn by sorting the score vector from highest score at the top to lowest score at the bottom, then adding down arrows from top to bottom. If arrows are omitted (to avoid clutter), this means that the upper alternative is preferred to the lower alternative. Inconsistencies are denoted by upward arrows, where an alternative with a lower score is strictly preferred to an alternative with a higher score. As shown in Figure 3, there are exactly 4 non-isomorphic structures for tournaments of four alternatives (Davis, 1954): a chain, a cycle among the top 3 alternatives, a cycle among the bottom 3 alternatives, and a single long cycle.

Notice that these structures have very different implications for decision-analysis. Given a choice between any subset of four alternatives, a decision-maker with a chain provides a ranking 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.
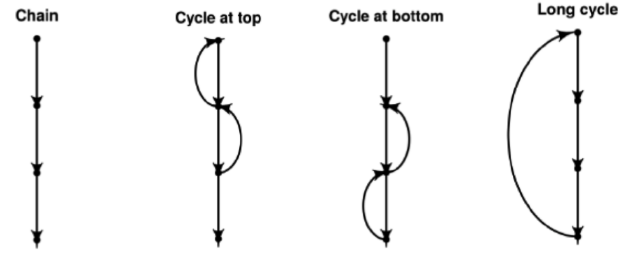

Figure 3: Tournaments on four alternatives. The score vectors are: chain $s = [3, 2, 1, 0]$, cycle at top $s = [2, 2, 2, 0]$, cycle at bottom $s = [3, 1, 1, 1]$, long cycle $s = [2, 2, 1, 1]$.

Preference graph similarity

Our primary analytical tool is a method of calculating the *distance* between graphs. Formally, 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 isomorphic, known as the Hamming distance $d_H(G_1, G_2) = ||vec(G_1) - vec(G_2)||_1$ (Hamming, 1950). 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, we can store the Hamming distance between all pairs of decision-makers in a symmetric $n \times n$ dissimilarity matrix D. From D, a weighted dissimilarity kernel K can be constructed, with values between zero and one (Kevin, 2012). We use standard graph similarity tools to identify clusters of graphs with similar content. This approach is formally equivalent to Coombs' multidimensional unfolding (Coombs & Kao, 1960).

Preference structure cannot be obtained from these Hamming distance computations. For example, two chain preference graphs of equal size with opposing preference content will have a Hamming distance equal to the total number of unordered pairs of vertices $\binom{n}{x}$. Even though they are both chains (identical structures), the Hamming distance indicates that they are as dissimilar as possible. Thus, we need a metric that indicates that these graphs have the same structure and hence that there is a *structural distance* of zero between them. Two graphs have a structural distance of zero if they are isomorphic (Aflalo, Bronstein, & Kimmel, 2015), 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). An *automorphism* of a graph G is a graph that is isomorphic to G, and the *automorphism group* $Aut(G)$ is all of the graphs that are isomorphic to G (Babai & Luks, 1983). We can test whether two graphs are isomorphic by checking whether any of their automorphisms are isomorphic. This is a well studied problem in computer science, called the *graph isomorphism problem* (Babai & Luks, 1983). The minimum Hamming distance between two graphs across all combinations of their auto-

morphisms gives their *structural distance* d_S (Butts & Carley, 2005): $d_S(G_1, G_2) = min(d_H(Aut(G_1), Aut(G_2)))$. If two graphs are similar (but not isomorphic), their structural distance should be small. Clusters of decision-makers with small distances between each other then indicate a common preference structure in a population of decision-makers, partially masked by noise.

Inexact graph matching

With a few alternatives the structural distance between graphs can be quickly calculated using exhaustive search. As the number of alternatives grows, exhaustive search becomes unfeasible. In general, the problem of calculating structural distance is NP-hard (Aflalo et al., 2015), requiring approximation techniques for large graphs with more than 8 alternatives. To make this approximation feasible, we recast the structural distance calculation as an *inexact graph matching problem* (Livi & Rizzi, 2013), where the objective is to find the permutation matrix P* over the space of permutations that makes two adjacency matrices *A* and *B* as similar as possible. The objective function is (Aflalo et al., 2015; Livi & Rizzi, 2013; Vogelstein et al., 2011):

$$P^* = \underset{P \in \mathcal{P}}{\operatorname{argmin}} \ f(P) = \operatorname{dis}_{A \to B}(P) = ||A - P^T B P|| \quad (1)$$

where A, B are the adjacency matrices for the preference graphs of two decision-makers, and $P \in \mathcal{P}$ is in the set of permutation matrices $\mathcal{P}$. If the chosen norm is the Frobenius $L2$ norm squared the problem is know as quadratic assignment (QAP) with non-deterministic polynomial time complexity (Koopmans & Beckmann, 1957). Given that solution set $\mathcal{P}$ is not convex, a common approach is to relax the non-convex restriction $\Pi \in P$, replacing P by its convex hull $\mathcal{D}$, where $\mathcal{D}$ is the set of doubly stochastic matrices (all entries greater than equal to zero and each row and column sums to 1). After some algebra, this relaxation leads to a quadratic program (QCV) solvable in polynomial time (Liu, Qiao, Jia, & Xu, 2014; Aflalo et al., 2015). Nevertheless, this method can lead to inaccurate results (Aflalo et al., 2015). Instead, we followed Vogelstein's approach (rGM) (Vogelstein et al., 2011). We replaced the objective function $f(P)$ by the identity $-tr(APB^T P^T)$ which leads to a non-convex problem where $\nabla^2 f(P) = B \otimes A + B^T \otimes A^T$ is not positive definite (Vogelstein et al., 2011). (Vogelstein et al., 2015) proposed to solve this problem sequentially with Frank-Wolfe algorithm (Frank & Wolfe, 1956). We initialized the optimization with QCV solution (Lyzinski et al., 2016).

Clustering

Once content and structural distances d_s are determined for preference graphs of each pair of decision-makers, the matrix of pairwise structural (or hamming) distances between the graphs of decision-makers can be analyzed using traditional clustering techniques to classify decision-makers into groups with similar preference content and structure. Nonetheless,

nothing ensures that clusters from content and structural dissimilarities will not overlap. Therefore, we need to account for both structural and content dissimilarities simultaneously in the clustering stage. To do so we first embed each dissimilarity matrix in a lower dimensional space and hereafter we bound columns of the resulting embeddings in an $n \times d$ matrix with information about content and structure for each decision-maker, with d the sum of dimensions of the embeddings of both dissimilarity matrices or *embedding fusion*.

We begin by using classical multidimensional scaling to project each distance matrix onto a lower dimensional space (Torgerson, 1952), but based on its superior performance we finally used an autoencoder initialized with weights found using a Restricted Boltzmann Machines (Hinton & Salakhutdinov, 2006; Wang, Yao, & Zhao, 2016). An autoencoder is a neural network model that maps or encodes input space x into a lower dimensional space $h(x)$ at its output layer and then reconstructs or decodes the original input space as $\hat{x}(h)$ (Goodfellow, Bengio, & Courville, 2016).

Therefore, we first convert dissimilarities to probabilities with a radial basis function kernel with fixed σ as the median of the statistic D_{ij} (Kevin, 2012; Karatzoglou, Smola, Hornik, & Karatzoglou, 2016). Next we trained an autoencoder to embed each $n \times n$ kernel dissimilarity matrix in a lower dimensional space $n \times d$. We used a non-linear activation function with a sigmoid transformation for both the encoder $h(x) = Wx + b$ and the decoder $\hat{x} = Wh + c$, a 0.1 learning rate and 1,000 epochs. An autoencoder with a non-linear activation function can be understood as a non-linear version of principal component analysis (Goodfellow et al., 2016). An autoencoder seeks to minimize reconstruction error as (Goodfellow et al., 2016):

$$\underset{W,b,c}{\min} \ l(x) = -\sum_j x_j log(\hat{x}_j) + (1 - x_j) log(1 - \hat{x}_j) \quad (2)$$

Figure 4 presents an schema for the autoencoder network. We selected the number of dimensions in the embedding d that gave the lower training reconstruction error using the elbow method. We pretrained the model with a Restricted Boltzmann Machine (Hinton & Salakhutdinov, 2006). A Restricted Boltzmann Machine is an undirected energy-based graphical model where the visible inputs x are matched with hidden units in a lower dimensional space h. A Restricted Boltzmann Machine is fitted minimizing log-likelihood function using contrastive divergence algorithm as (Hinton, 2002):

$$\underset{\theta = \{W,b,c\}}{\min} l(\theta) = -\sum log(e^{-\mathcal{F}(x)}/Z) \quad (3)$$

$\mathcal{F}(x)$ is the free energy as $\mathcal{F}(x) = -log(\sum_h e^{-E(x,h;\theta)})$, with energy function $E(x,h;\theta) = b^T x + c^T h + x^T Wh$, Z a normalizing constant $Z = \sum_x e^{-\mathcal{F}(x)}$ and W, b and c are model weights. We tested our method against other linear and non linear embedding methods with data sets commonly used for testing clustering methods with superior performance in
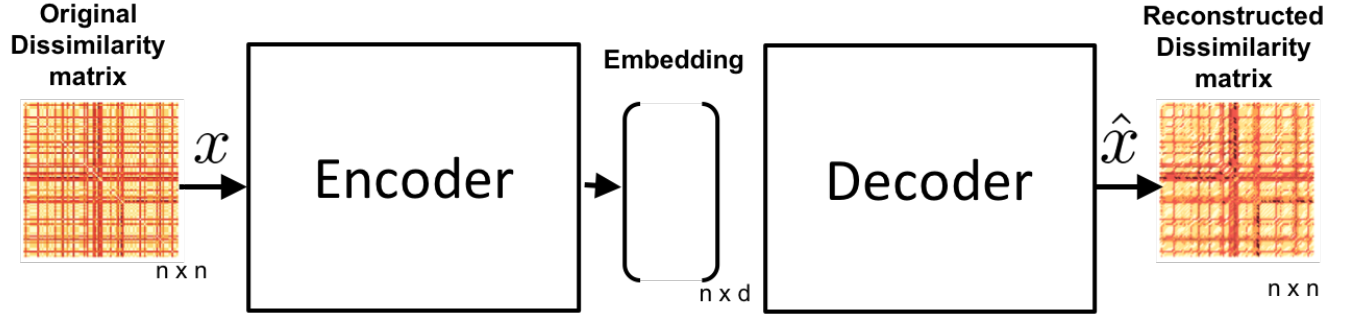
Figure 4: The following autoencoder schema describes our method. The original dissimilarity matrix is encoded into a lower dimensional space minimizing reconstruction error (Wang et al., 2016).

all cases. Finally, to achieve a robust solution, we used k-medians algorithm to determine clusters allocation (Singh, Yadav, & Rana, 2013). Therefore, to determine the clusters we solve the following optimization problem:

$$\min J(\gamma, \mu) = \sum_i^n \sum_j^k \gamma_{ij} ||x_i - \mu_j||_1 \qquad (4)$$

We initialized the algorithm with centroids from a prior hierarchical k-means solution (Hartigan & Wong, 1979; Lucas, 2014; Arai & Barakbah, 2007). Here γ is a binary allocation matrix, k is the apriori defined number of clusters, C is the clusters allocation and μ the vector with medians for each group. We used the gap-statistic to determine the number of clusters k (Tibshirani, Walther, & Hastie, 2001). If necessary, clusters are merged to provide a more general solution.

Preference structure in simulation

We first describe the results of simulations designed to illustrate the method. In our simulation we evaluate our model's ability to separate a popular psychological model, the lexicographic semiorder (Tversky, 1969), from the more traditional expected utility maximization (Von Neumann & Morgenstern, 1944). As an example, consider choosing between pairs of gambles shown in Table 1 from Tversky's classic paper on intransitive preferences (Tversky, 1969), along with three additional gambles (f-h) added to increase graph matching difficulty.

Subjects that chose based on expected value should prefer $a \succ b \succ c \succ d \succ e$, and should have a complete transitive order. Tversky hypothesized that someone following a lexicographic semiorder decision rule would first choose based on differences in gambles probabilities. If the difference in probabilities is small enough, the decision maker would switch to the next attribute and choose based on differences in payoffs. This would result in an intransitive sequence $a \succ b \succ c \succ d \succ e$ and $e \succ a$.

To demonstrate that our approach can reliably cluster decision-makers into groups based on the structure of their preferences, we generate graphs for 100 decision-makers, 37 with lexicographic preferences, 33 with risk neutral expected

Table 1: Gambles from Tversky's (Tversky, 1969) experiment 1 (a-e) plus three added for the simulation (f-h)

Gamble	Probability	Payoff	Expected Value ($)
a	7/24	5.00	1.46
b	8/24	4.75	1.58
c	9/24	4.50	1.69
d	10/24	4.25	1.77
e	11/24	4.00	1.83
f	12/24	3.75	1.88
g	13/24	3.50	1.894
h	14/24	3.25	1.895

value maximizer preferences, and 30 that would choose at random. We first mapped the simulated choices in an adjacency matrix, then computed dissimilarity matrices between adjacency matrices and finally we identified clusters of graphs with similar preference content and structure. Figure 5 summarizes our method. As exposed in Figure 5 k-means clustering (Hartigan & Wong, 1979) on a two-dimensional embedding separates successfully lexicographic semiorders from those who are expected value maximizers.

We first assumed subjects will choose deterministically. A deterministic decision rule will provide structures that are quite easy to distinguish from others because, under all circumstances, the same graph structure will emerge. Nonetheless, noise in the decision process can make preference structures harder to distinguish. For example, in Figure 6 even though a decision rule such as expected value maximization is used, intransitive behavior is still observable if alternatives are harder to compare, confounding the later decision rule with random choice. In Figure 6, the first row presents structural distance dissimilarity matrix for 100 simulated decision-makers with increasing noise for expected utility maximization as $EV + N(0, \sigma^2)$. The second row, presents the sorted dissimilarity matrix by final clustering. The third row, presents a two-dimensional embedding of the dissimilarity matrix.
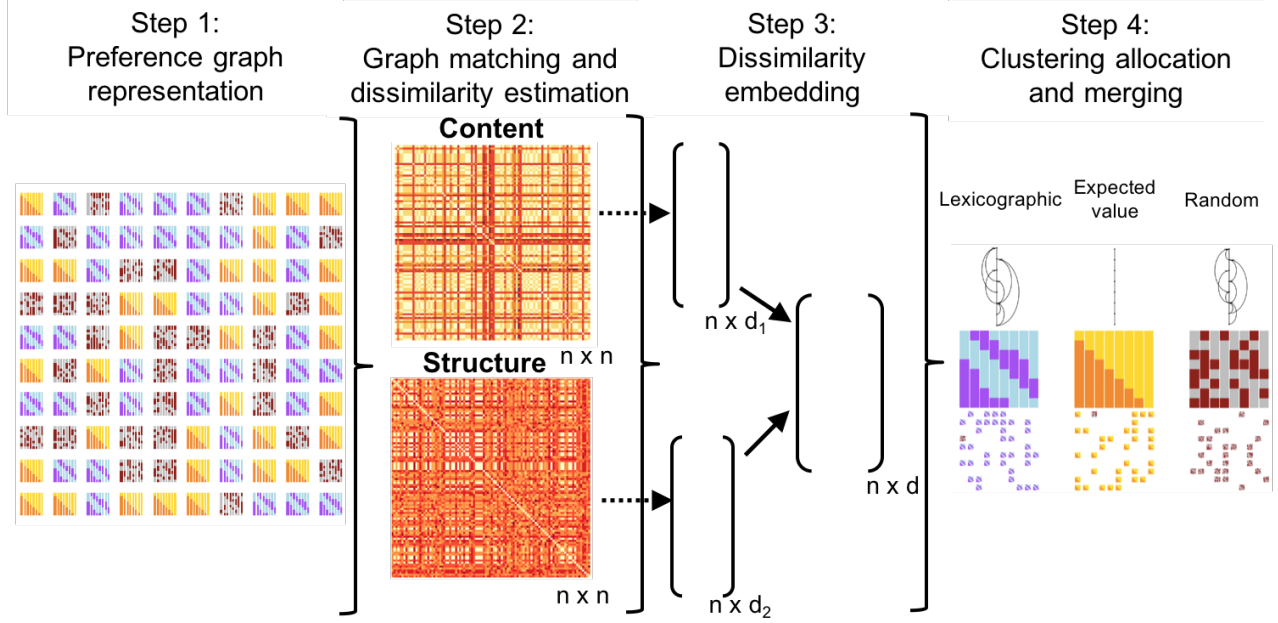
.

Figure 5: 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.
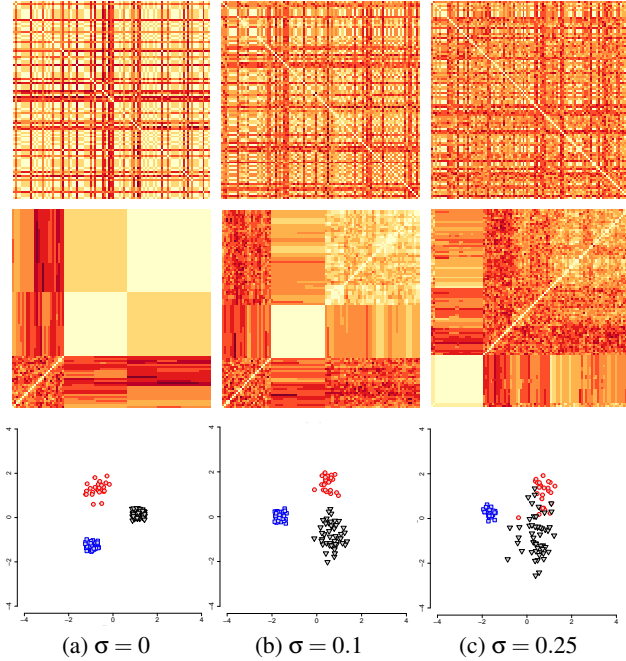

Figure 6: First row, structure dissimilarity matrix with increasing noise for expected utility maximization. Second row, structure dissimilarity matrix sorted by final clustering. Third row, structure dissimilarity matrix two-dimensional embedding (lexicographic purple squares, random choice red circles and expected value grey triangles).

An empirical test of the model

We tested our method in three most relevant tasks: an extension of (Tversky, 1969) transitivity task, a selection of classic anomalies observed in choices between risky prospects and a policy-focused multi-attribute task to study preferences for CO_2 mitigation (Sergi, Davis, & Azevedo, 2017). For each task, we recruited 200 participants using Amazon Mechanical Turk (Mturk). Inclusion criteria were the following: 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. 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. We detected six, seven and two clusters in each of tasks. In Figure 7 we present a two dimensional visualization for dissimilarity embedding fusion results using t-Distributed Stochastic Neighbor embedding (t-SNE) (Maaten & Hinton, 2008). Clusters are indicated in different colors and shapes. Voronoi polygons are plotted to show cluster separation. Next we present a detailed analyzes of our results in each task.

Transitivity task

We first extended Tversky's classic experiment examining lexicographic semiorders (Tversky, 1969). Participants choose between the pairs of gambles shown in Table 1 from Tversky's classic paper on intransitive preferences (Tversky, 1969), along with the three gambles (f-h) considered in the simulation experiment and two additional gambles (i-j) where a higher probability is negatively correlated with a higher ex-
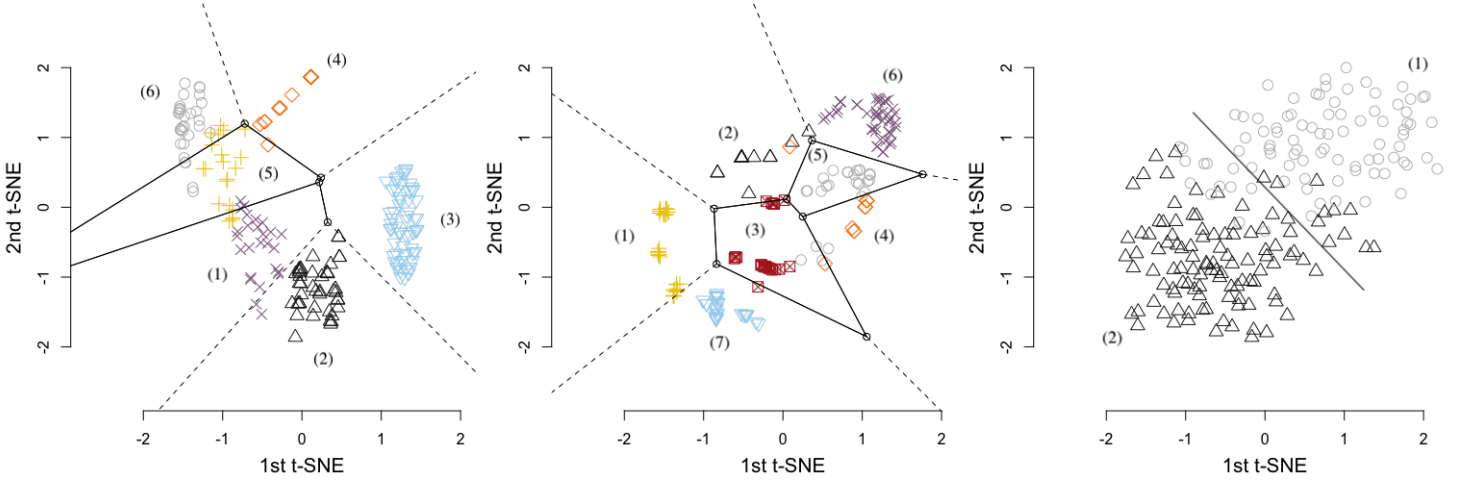
Figure 7: Visualization of dissimilarity embedding fusion results in two dimensions with t-Distributed Stochastic Neighbor Embedding (t-SNE) (Maaten & Hinton, 2008). Clusters are indicated in different colors and shapes. Voronoi polygons are plotted to show cluster separation.

pected value (in i there is 15/24 chance of winning $3, and in j a 16/24 chance of winning $2.75). Probabilities were presented as pie charts without numeric information. An example of the choice task is presented in Figure **??**.

We presented participants with all pair combinations (45 pairs) in three repetitions with the order randomized. 95% of the 200 participants was paid the bonus. Clustering by content and structure six clusters emerged: five groups with chains and one group with multiple cycles. Structural and content heterogeneity, rather than homogeneity, is the primary takeaway. Chain graphs, are the most popular pattern. Figure 8 shows the expected adjacency matrices for the preference graphs in each cluster. Alternatives were pre-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.
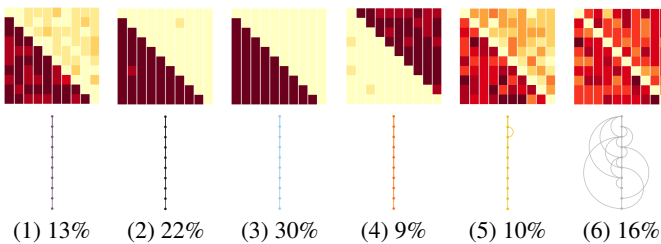

| (1) 13% | (2) 22% | (3) 30% | (4) 9% | (5) 10% | (6) 16% |

Figure 8: 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 in the last row.

To further analyze the obtained clusters we used multinomial logit models to analyze choices (MNL) (McFadden, 1973). An MNL model assumes the probability that an individual chooses alternative $i \in J$ depends linearly on a utility function V_i as $P_i = e^{V_i} / \sum_{j \in J} e^{V_j}$ (McFadden, 1973). We proposed three potential decision rules for the task: maximize expected value $V_i = \beta$ EV, maximize probability of winning $V_i = \beta P(\text{winning})$ and maximize payoffs $V_i = \beta$ Payoff. In Table 3, we fitted a binary logit model with the decision rules as linear utilities in each of the discovered clusters (McFadden, 1973). Figure 9 shows logit probabilities in each case. For all clusters a decision rule based on a single attribute (either probabilities or payoffs) is more likely 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%. It is possible that this clusters respond to the same decision rule with differences in discriminant ability. We must highlight that given that probabilities are not numerically stated, recognizing the alternative with a higher probability in all problems as in cluster 3 requires a superior classification skill. We decided to merge this clusters in one group. Decision-makers in cluster 4 consistently chose the alternative with a higher payoff 93% of times, indicating a single attribute decision rule based on payoffs is a good representation of the choice process. Clusters 5 and 6 seem to respond to a different decision process where decisions are based on probabilities, but other elements are weighted in the decision. More exploration is needed to determine if participants in the later two clusters are following some sort of lexicographic strategy or we observed a new undocumented structure. In cluster 6 multiple cycles are observed. The proportion of choices in cluster 6 favoring the option with the higher probability is significantly distinct from 50% ruling out random

choice. Although details of the different choice rules remain uncovered, we observed a clear tendency to choose based on probabilities (Lichtenstein & Slovic, 1971). It seems the data is more consistent with a lexicographic order (up to noise) in the sense of (Fishburn, 1971) than a lexicographic semiorder as proposed by (Tversky, 1969). (M. H. Birnbaum & Gutierrez, 2007) documented similar findings.

Table 2: Linear utility models per cluster. 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.

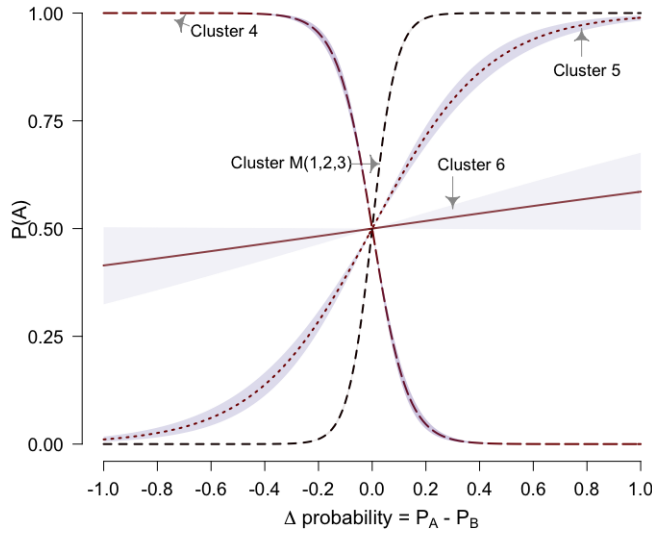| Cluster	Content	$\hat{\beta}$	l(s)	l(EV)	P(p)	N (%)
1	Probs	10***	-542	-621	87%***	26 (13%)
2	Probs	23***	-473	-773	96%***	44 (22%)
3	Probs	111***	-30	-717	100%***	59 (30%)
4	Payoff	3***	-284	-382	7%***	18 (9%)
5	Probs	5***	-558	-586	71%***	20 (10%)
6	Probs	0.3*	-1,029	-1,029	54%***	33 (16%)

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

Figure 9: Logit probabilities P(A) of choosing the alternative with a higher probability of winning (A) per cluster. Clusters 1, 2 and 3 respond to the same decision rule up to noise, so we merged them. While decision-makers in clusters 1, 2 and 3 chose based on the probability of winning, decision-makers in cluster 4 chose maximizing the payoff. Clusters 5 and 6, respond to a different decision process.

CO2 task

Policy-focused researchers have used multi-attribute discrete choice models to estimate policy-relevant quantities for their specific problems, such as the market share of existing and new products (Herriges & Kling, 1999; Greene, 2012), substitution patterns (Hensher, Rose, & Greene, 2015), implicit discount rates (Min, Azevedo, Michalek, & Bruine de Bruin, 2014), willingness-to-pay (McFadden, 1999; Helveston et al., 2015), and consumer's surplus (Small & Rosen, 1981; Williams, 1977). To study the benefits of our method, we collected new data based on work done at Carnegie Mellon done by (Sergi et al., 2017). (Sergi et al., 2017) designed an experimental paradigm to estimate willingness to pay values for CO_2 emission reductions. In our extension of (Sergi et al., 2017) experiment, participants must trade-off between higher (or lower) impacts of electricity generation on climate change and a higher (or lower) electricity bill among the following alternatives (CO_2, Bill): a (-30%, 20%), b (-30%, 5%), c (-25%, 4%), d (-20%, 3%), e (-15%, 2%), f (-10%, 1%), g (30%, -20%), h (30%, -5%). An example of the choice task is presented in Figure **??**. The levels of each scenario were selected in order to discover lexicographic semiorder behavior with electricity bill as first priority. We presented participants with all pair combinations with no repetitions (28 pairs). Alternatives were prearranged so a lower triangular adjacency matrix indicates choices based strictly on bill and an upper triangular adjacency matrix indicates choices based strictly on CO_2. 97% of the 200 participants pass the attention check.

As shown in Figure 10, seven clusters emerged, five groups with chains structures and two groups with multiple cycles. We classified participants in three groups regarding preference content: 1) Greens (focused mainly on CO_2) and 2) Bills (focused mainly on costs) and 3) Cycles were preferences are unclear. While in cluster 1, decision-makers chose strictly based on a lower electricity bill, in clusters 2 and 3 decision-makers are willing to trade-off a higher bill for improvements in environmental quality. Decision-makers in cluster 4 strictly focused their choices on lower CO_2 emissions. Although decision-makers in clusters 5 and 6 also favored the environmental attribute, it seems there is a threshold in terms of higher cost they are not willing to cross. In exchange, decision-makers in cluster 7 show multiple cycles and are less certain of what they want.

Calculation of willingness to pay presumes that preferences are well behave and a cardinal utility function underlies a decision-maker's choices. As a result, researchers typically assume a structure such as a strong utility function (a cardinal utility function plus noise (Block, 1974; Luce & Suppes, 1965; McFadden, 1975, 1976, 1973)). This includes various twists on the generalized extreme value model (which includes the multinomial logit and nested logit) (McFadden, 1997), and the mixed logit model (Revelt & Train, 1998). As we can observe in Figure 10 a significant proportion of the sample (28%) do not have well-behaved preferences. Furthermore, two groups although with a chain structure (24%) respond to a lexicographic order and cannot be represented by a utility function (Mas-Colell, Whinston, Green, et al., 1995).

In the naive approach, a modeler would fit a single multinomial logit model with the full sample (McFadden, 1997).

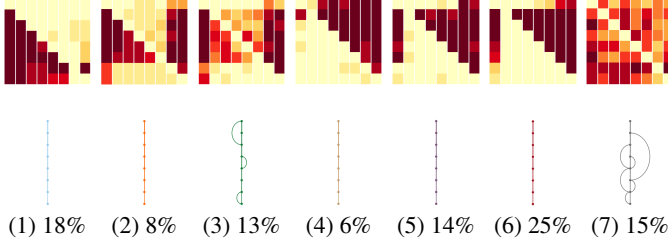(1) 18% (2) 8% (3) 13% (4) 6% (5) 14% (6) 25% (7) 15%

Figure 10: Weighted expected adjacency matrix in each cluster for the CO_2 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 in the last row.

Willingness to pay corresponds to the *marginal rate of substitution* (MRS) between attribute k and the cost of each alternative $MRS_{kc} = \frac{\partial u_i}{\partial x_k} / \frac{\partial u_i}{\partial c_i}$. When a linear in attributes utility function is assumed, it would result in the following specification: $V_i = -8.3 \cdot Bill - 4.5 \cdot CO_2$. Hence, willingness to pay (WTP) corresponds to the ratio of coefficients leading to $WTP = 0.3 \cdot -4.5 / -8.3 = 0.16$ % increment in monthly electricity bill for a 30% percent reduction in CO_2 emissions. The result, at first sight might seem plausible. A more sophisticated modeler could think there is segments in the sample with different preferences. If we presume there are two groups in the sample, a Latent class model (LC) (Greene & Hensher, 2003) would produce two linear utility functions: $V_i^1 = -6.6 \cdot Bill + 0.7 \cdot CO_2$ ($\pi_1 = 33\%$) and $V_i^2 = -16.1 \cdot Bill - 11.8 \cdot CO_2$ ($\pi_2 = 67\%$), with $WTP_1 = -0.03$ and $WTP_2 = 0.22$. Here the second group is willing to pay to hurt the environment, which seems improbable. In the LC model, individuals were assigned to different classes $q \in \{1, 2\}$, with $P_{i/q} = e^{V_i^q} / \sum_{j \in J} e^{V_j^q}$ the probability of choosing alternative $i \in J$ given class q and π_q the probability that an individual belongs to class q. Hence, $P_i = \pi_1 \cdot P_{i/1} + \pi_2 \cdot P_{i/2}$. In Table 3 we present results from different linear models in each of the discovered clusters.

Table 3: Willingness to pay per cluster

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%)

As shown Table 3, Greens in cluster 4 are willing to pay a higher amount to the one accounted in the first LC model. In exchange, Bills in cluster 1 hold a willingness to pay close to zero showing no compromise to protect the environment.

This results suggests a lexicographic order behavior in both groups, with decision-makers focusing the decision only in one attribute (Fishburn, 1971). We observed a compensatory behavior in clusters 2, 3, 5 and 6, where decision-makers are willing to make trade-offs between the two attributes, with a similar willingness to pay in all groups. We observed multiple cycles in both cluster 3 and 7. We cannot confirm if cycles represent some sort of lexicographic semiorder behavior (Tversky, 1969) or a different uncovered structures. Further studies must be undertaken in order to better understand choice strategies in those groups. In Figure 11 we present slopes for both attributes (bill and CO_2) again assuming a weighted additive linear utility model with no intercept as $V_i = \beta_{Bill} \cdot Bill + \beta_{CO_2} \cdot CO_2$. Given their similarities we merged clusters 2 and 3; and clusters 5 and 6. Decision-makers in cluster 1 are insensitive to changes CO_2, whereas decision-makers in cluster 4 are insensitive to changes in electricity bill. Decision-makers in groups with merged clusters M(2,3) and M(5,6), are sensitive to changes in both attributes with different levels of intensity. Decision-makers in cluster 7 are not sensitive to changes in any of the attributes, possibly choosing at random.
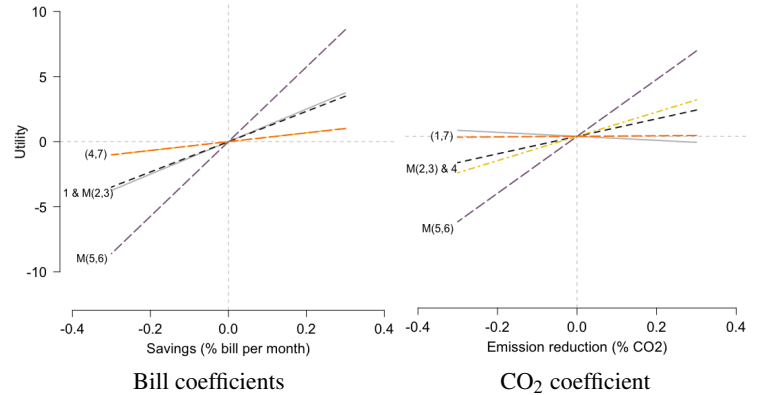

Bill coefficients   CO$_2$ coefficient

Figure 11: Slopes 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.

Anomalies task

Researchers have examined how decision-makers choose among sets of gambles, finding a series of anomalies that challenge the descriptive validity of expected utility theory. The lexicographic semiorder (Tversky, 1969) and Prospect Theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) are two of the most successful in this class. Other researchers have proposed alternative models, including the priority heuristic (Brandstätter, Gigerenzer, & Hertwig, 2006), transfer of attention exchange (M. H. Birnbaum & Chavez, 1997; M. Birnbaum, 2008), decision field theory (Busemeyer & Townsend, 1993; Roe, Busemeyer, & Townsend, 2001), and more recently BEAST model (Erev, Ert, Plonsky, Cohen,

& Cohen, 2017). Each aims to synthesize well-established deviations from rationality into a parsimonious model. A recent choice modeling competition focusing on the BEAST model (Erev et al., 2017), identified a set of gambles to test the presence of the main anomalies detected in descriptive choice. The gambles themselves provide a rich testbed for our approach, as they can expose a variety of patterns that have been previously examined, as well as the possibility of observing new patterns that may have been overlooked. In the experiment, we used 11 gambles with 55 pairs in three repetitions with the order randomized.

The anomalies tested in our experiment comprised five anomalies used to develop Prospect Theory (Kahneman & Tversky, 1979), namely the 1) certainty effect, 2) the reflection effect, 3) overweighting of rare events, 4) loss aversion and 5) risk aversion. In this case, only 55% of the sample pass the attention check. In Figure **??** we present an choice set example. In Table 4 we present the alternatives in the experiments and the expected preference relation for all anomalies. When we clustered participants, two groups emerged. As shown in Figure 12 both groups present a chain structure in expectation. Nonetheless, decision-makers seemed less consistent than in the first two tasks. We must acknowledge that given that gambles in this task were designed to produce choice anomalies it is reasonable to observe irrational behavior. In Table 5 we present the proportion of decision-makers matching the anomalies in each cluster. We could not replicate the certainty effect, neither the reflection effect, nor overweighting of rare events. (Erev et al., 2017). Overall behavior in both groups is very similar. Further inquiries are needed to fully explain the observed choice patterns.

Table 4: Gambles based on (Kahneman & Tversky, 1979; Erev et al., 2017). 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 (Erev et al., 2017).

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

We extended our analysis using latent class models (Greene & Hensher, 2003). We model choices as an ensemble of four decision rules (Erev et al., 2017): R_1) maximize expected value and minimize variance (EV) (Von Neumann & Morgenstern, 1944; Levy & Markowitz, 1979); R_2) maxi-
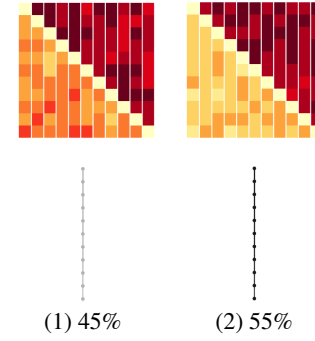

(1) 45%   (2) 55%

Figure 12: 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. The proportion of the sample in each cluster is presented in the last row.

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

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%)

mize the probability of a better outcome (PB) (Venkatraman, Payne, & Huettel, 2014); R_3) maximize a weighted additive function of outcomes (WA) (Payne et al., 1993); and R_4) maximize the probability of winning the high outcome (PW) (Erev et al., 2017). Again we modeled the probability of choosing alternative $i \in J$ with a logit model as $P_i = \sum_{R_q \in R} \pi_{R_q} P_{i/R_q}$, with π_{R_q} the likelihood that an individual uses decision rule $R_q \in R$. In Figure 13 we show results from our ensemble model. We present the proportion of participants assigned to each decision rule defined as $\max_{R_q} P(R_q/Y = y)$ for all 55 problems. There is no noticeable differences between both clusters with probability of a better outcome as most likely decision rule in most problems. The two clusters belong to the same decision process and could be merged.

Discussion and future work

In this paper, we demonstrate how to use graph matching to uncover heterogeneity in the structure of preference across a population of decision-makers and thereafter cluster decision-makers based on graph embedding methods. We explore the approach with simulated choice data from the most common classes of economic and psychological models. We also apply the method to new empirical implementations of
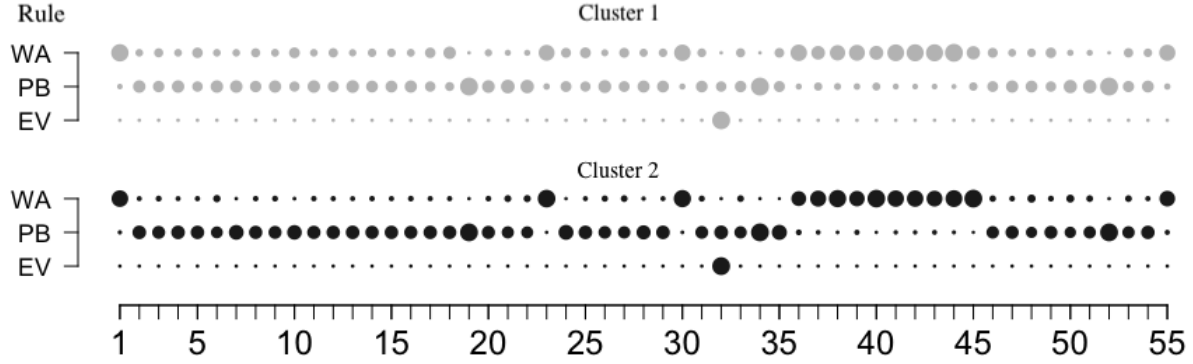
Figure 13: Latent class model results. Proportion per decision rule from posterior distribution $P(R_q/Y = y)$. Decision rules are: WA = weighted additive (outcomes), PB = P(Better outcome) and EV = expected value plus variance.

classic experiments in decisions between risky prospects and other stated preferences tasks. The approach uncovers heterogeneity in preference structure across a variety of dimensions, without requiring any prior knowledge of those structures. Both the proposed method and its results have important implications for researchers in the psychological and economic sciences as it does not depend on apriori defined theories, only the pattern of choices.

As shown in the three empirical tasks, we exploited the regularities in choice patterns to identify individuals using a similar choice rule without prior assumptions. Our model is able to separate patterns consistent with utility theory from those better described by other descriptive theories such as lexicographic order, where decision-makers are choosing based on a single attribute. We also can separate decision-makers showing multiple cycles. The prevailing notion in the decision-sciences is that decision-makers are either rational or irrational. Our results suggest that this is not the case. Some decision-makers seemed to know exactly what they want with a chain structure, while others are less certain. Some decision-makers can order almost every alternative except the bottom few. Others struggle to order the best alternatives. Still others have so many inconsistencies that their choices can hardly be considered to represent preferences at all. Is yet to be confirmed if cyclic preference graphs might be explain by insensitivity to changes in attribute levels or by differences in discriminant ability.

We tested if cluster information can increase predictive accuracy. We used a logistic regression model with linear utility function for all three tasks and we simulated 1,000 bootstrapped samples from our survey data to construct confidence intervals. We split each bootstrapped sample in a train set (66%) and test set (34%). Accuracy is significantly higher if we use a factor variable with the clusters as an interaction term. A likelihood ratio test also shows that the model with a factor variable for clusters is superior to a pooled model (p-value < 0.001). The learned clusters contain relevant information of the choice process.

In the transitivity task the vast majority of the sample (74%) uses a single attribute (up to noise) to make the choice, undermining the plausibility of other more complex rules like expected value calculations. Decisions based on simple rules are quite successful, reducing at the same time the effort required in the task. Although participants with simple rules show more coherent preferences, they can hardly be represented by a utility function. Using classic utility models to approximate lexicographic decision rules can give wildly inaccurate answers, overestimating willingness to compromise. For example, in the CO_2 task we observed both non-compensatory behavior in 24% of the sample. A pooled linear utility model would suggest the population is willing to compromise to protect the environment, whilst a large segment will not be willing to make such trade-offs. Lastly, the anomalies task provided an interesting testbed of the limits of our method. Although we could not replicate the expected anomalous behavior, it shows the common failure of partition algorithms splitting the sample when there is only one cluster.

Our model synthesizes choice models into a general framework for analyzing and discovering preference structures. The approach has the potential to transform current knowledge and approaches to understand preferences, which currently focus on specific structures, and lack a unifying framework for both theoretical and empirical analysis of the preferences of many decision-makers. The results can impact research in the decision-sciences, with applications to health, environmental decision-making, and fundamental studies of human cognition. Practitioners will be able to use this approach to 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. This can inform decision-makers themselves through decision analysis, as well as policy-makers, to help them better understand the welfare impacts of new policies.

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 is defined using the elbow method. Determining the number of dimensions in the op-

timization process can offer a potential improvement (Côté & Larochelle, 2016). Future applications should also developed better ways of determining the number of clusters and herein merging similar clusters. New methods that marry k-means and hierarchical clustering algorithms provide an interesting direction to automate the merging stage (Peterson, Ghosh, & Maitra, 2018). The experimental design also provides some challenges. The number of pairwise comparisons required to complete a tournament grows exponentially with the number of alternatives, increasing the risk of observing mental fatigue through the experiment. A new experimental paradigm needs to be develop in order determine apriori the minimal number of questions required to recover preference structure information and henceforth lower the cognitive burden from human participants. Although our method is valid for any type of pairwise comparison, empirical tests should be extended to other experimental domains.

Acknowledgments

References

Aflalo, Y., Bronstein, A., & Kimmel, R. (2015). On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, *112*(10), 2942–2947.

Afriat, S. (1972). Efficiency estimates of production functions. *International Economic Review*, *13*, 568–598.

Aiolli, F., & Sperduti, A. (2010). A preference optimization based unifying framework for supervised learning problems. In *Preference learning* (pp. 19–42). Springer.

Anderson, J. R. (1996). Act: A simple theory of complex cognition. *American Psychologist*, *51*(4), 355.

Anderson, J. R., & Lebiere, C. J. (2014). *The atomic components of thought*. Psychology Press.

Arai, K., & Barakbah, A. R. (2007). Hierarchical k-means: an algorithm for centroids initialization for k-means. *Reports of the Faculty of Science and Engineering*, *36*(1), 25–31.

Arrow, K. J. (1951). *Social choice and individual values*. John Wiley Sons, Inc: New York.

Babai, L., & Luks, E. M. (1983). Canonical labeling of graphs. In *Proceedings of the fifteenth annual acm symposium on theory of computing* (pp. 171–183).

Bentham, J. (1879). *An introduction to the principles of morals and legislation*. Clarendon Press.

Birnbaum, M. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*(2), 463–501.

Birnbaum, M. H., & Chavez, A. (1997). Tests of theories of decision making: Violations of branch independence and distribution independence. *Organizational Behavior and human decision Processes*, *71*(2), 161–194.

Birnbaum, M. H., & Gutierrez, R. J. (2007). Testing for intransitivity of preferences predicted by a lexicographic semi-order. *Organizational Behavior and Human Decision Processes*, *104*(1), 96–112.

Block, H. (1974). Random orderings and stochastic theories of responses (1960). In *Economic information, decision, and prediction* (pp. 172–217). Springer.

Bouyssou, D., & Vincke, P. (2010). Binary relations and preference modeling. *Decision-making Process: Concepts and Methods*, 49–84.

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: making choices without trade-offs. *Psychological Review*, *113*(2), 409–432.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432–459.

Butts, C. T., & Carley, K. M. (2005). Some simple algorithms for structural comparison. *Computational & Mathematical Organization Theory*, *11*(4), 291–305.

Coombs, C. H., & Kao, R. C. (1960). On a connection between factor analysis and multidimensional unfolding. *Psychometrika*, *25*(3), 219–231.

Côté, M.-A., & Larochelle, H. (2016). An infinite restricted boltzmann machine. *Neural computation*.

Davis, R. L. (1954). Structures of dominance relations. *Bulletin of Mathematical Biology*, *16*(2), 131–140.

Dictionary, O. E. (1989). Oxford english dictionary. *Simpson, JA & Weiner, ESC*.

Diestel, R. (2000). *Graphentheory*. Springer.

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*(4), 369–409.

Fischhoff, B. (2005). Cognitive processes in stated preference methods. *Handbook of environmental economics*, *2*, 937–968.

Fishburn, P. C. (1971). A study of lexicographic expected utility. *Management Science*, *17*(11), 672–678.

Frank, M., & Wolfe, P. (1956). An algorithm for quadratic programming. *Naval Research Logistics (NRL)*, *3*(1-2), 95–110.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.

Gigerenzer, G., Todd, P. M., ABC Research Group, t., et al. (1999). *Simple heuristics that make us smart*. Oxford University Press.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.

Greene, W. H. (2012). *Econometric analysis*. Granite Hill Publishers.

Greene, W. H., & Hensher, D. A. (2003). A latent class model for discrete choice analysis: contrasts with mixed logit. *Transportation Research Part B: Methodological*, *37*(8), 681–698.

Gross, J. L., & Yellen, J. (2004). *Handbook of graph theory*. CRC press.

Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell Labs Technical Journal*, *29*(2), 147–160.

Hartigan, J., & Wong, M. (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, *28*(1), 100–108.

Helveston, J. P., Liu, Y., Feit, E. M., Fuchs, E., Klampfl, E., & Michalek, J. J. (2015). Will subsidies drive electric vehicle adoption? measuring consumer preferences in the us and china. *Transportation Research Part A: Policy and Practice*, *73*, 96–112.

Hensher, D. A., Rose, J. M., & Greene, W. H. (2015). *Applied choice analysis*. Cambridge University Press.

Herriges, J. A., & Kling, C. L. (1999). Nonlinear income effects in random utility models. *Review of Economics and Statistics*, *81*(1), 62–72.

Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural computation*, *14*(8), 1771–1800.

Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, *313*(5786), 504–507.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Karatzoglou, A., Smola, A., Hornik, K., & Karatzoglou, M. A. (2016). Package kernlab. *Google Scholar*.

Kevin, M. (2012). *Machine learning: a probabilistic perspective*. The MIT press.

Koopmans, T. C., & Beckmann, M. (1957). Assignment problems and the location of economic activities. *Econometrica: journal of the Econometric Society*, 53–76.

Levy, H., & Markowitz, H. M. (1979). Approximating expected utility by a function of mean and variance. *The American Economic Review*, *69*(3), 308–317.

Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of experimental psychology*, *89*(1), 46.

Liu, Z.-Y., Qiao, H., Jia, L.-H., & Xu, L. (2014). A graph matching algorithm based on concavely regularized convex relaxation. *Neurocomputing*, *134*, 140–148.

Livi, L., & Rizzi, A. (2013). The graph matching problem. *Pattern Analysis and Applications*, *16*(3), 253–283.

Lucas, A. (2014). amap: Another multidimensional analysis package. *http://CRAN. R-project. org/package= amap*.

Luce, R. D., & Suppes, P. (1965). *Preference, utility, and subjective probability*. Wiley.

Lyzinski, V., Fishkind, D. E., Fiori, M., Vogelstein, J. T.,

Priebe, C. E., & Sapiro, G. (2016). Graph matching: Relax at your own risk. *IEEE transactions on pattern analysis and machine intelligence*, *38*(1), 60–73.

Maaten, L. v. d., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, *9*(Nov), 2579–2605.

Mas-Colell, A., Whinston, M. D., Green, J. R., et al. (1995). *Microeconomic theory* (Vol. 1). Oxford university press New York.

Maturana, H. R., & Varela, F. J. (1987). *The tree of knowledge: The biological roots of human understanding*. New Science Library/Shambhala Publications.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). Academic Press: New York.

McFadden, D. (1975). The revealed preferences of a government bureaucracy: Theory. *The Bell Journal of Economics*, 401–416.

McFadden, D. (1976). The revealed preferences of a government bureaucracy: Empirical evidence. *The Bell Journal of Economics*, 55–72.

McFadden, D. (1997). Modeling the choice of residential location. In J. Quigley (Ed.), *The economics of housing* (Vol. 1, pp. 531–552).

McFadden, D. (1999). Computing willingness-to-pay in random utility models. In J. Moore, R. Reizman, & J. Melvin (Eds.), *Trade: Theory and econometrics* (pp. 253–274). Routledge: London.

McFadden, D. (2001). Economic choices. *The American Economic Review*, *91*(3), 351–378.

Min, J., Azevedo, I., Michalek, J., & Bruine de Bruin, W. (2014). Labeling energy cost on light bulbs lowers implicit discount rates. *Ecological Economics*, *97*, 42–50.

Moon, J. W. (2015). *Topics on tournaments in graph theory*. Courier Dover Publications.

Pareto, V. (1906). *Manuale di economia politica* (Vol. 13). Societa Editrice.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.

Peterson, A. D., Ghosh, A. P., & Maitra, R. (2018). Merging k-means with hierarchical clustering for identifying general-shaped groups. *Stat*, *7*(1).

Ramachandran, V. S. (1992). Blind spots. *Scientific American*, *266*(5), 86–91.

Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2010). Testing transitivity of preferences on two-alternative forced choice data. *Frontiers in psychology*, *1*(148), 1–15.

Revelt, D., & Train, K. (1998). Mixed logit with repeated choices: households' choices of appliance efficiency level. *Review of Economics and Statistics*, *80*(4), 647–657.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*(2), 370–392.

Sergi, B., Davis, A., & Azevedo, I. (2017). Understanding

public perceptions of energy tradeoffs in climate, health, and consumer costs.

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological bulletin*, *134*(2), 207.

Simon, H., & Chase, W. (1988). Skill in chess. In *Computer chess compendium* (pp. 175–188). Springer.

Simon, H. A. (1956). Rational choice and the structure of the environment [Journal Article]. *Psychological Review*, *63*(2), 129-138. doi: 10.1037/h0042769

Simon, H. A. (1972). Theories of bounded rationality. *Decision and Organization*, *1*(1), 161–176.

Singh, A., Yadav, A., & Rana, A. (2013). K-means with three different distance metrics. *International Journal of Computer Applications*, *67*(10).

Slagle, J. R. (1971). *Artificial intelligence: the heuristic programming approach*. McGraw-Hill.

Small, K. A., & Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica*, *49*(1), 105–130.

Tibshirani, R., Walther, G., & Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *63*(2), 411–423.

Torgerson, W. S. (1952). Multidimensional scaling: I. theory and method. *Psychometrika*, *17*(4), 401–419.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal [Journal Article]. *PNAS*, *109*(24), 9659-9664. doi: 10.1073/pnas.1119569109

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*, *76*(1), 31-48.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323.

Varian, H. R. (1983). Non-parametric tests of consumer behaviour. *The review of economic studies*, *50*(1), 99–110.

Venkatraman, V., Payne, J. W., & Huettel, S. A. (2014). An overall probability of winning heuristic for complex risky decisions: Choice and eye fixation evidence. *Organizational Behavior and Human Decision Processes*, *125*(2), 73–87.

Vogelstein, J. T., Conroy, J. M., Lyzinski, V., Podrazik, L. J., Kratzer, S. G., Harley, E. T., ... Priebe, C. E. (2011). Fast approximate quadratic programming for large (brain) graph matching. *arXiv preprint arXiv:1112.5507*.

Vogelstein, J. T., Conroy, J. M., Lyzinski, V., Podrazik, L. J., Kratzer, S. G., Harley, E. T., ... Priebe, C. E. (2015). Fast approximate quadratic programming for graph matching. *PLOS one*, *10*(4), e0121002.

Von Neumann, J., & Morgenstern, O. (1944). Theory of games and economic behavior.

Wang, Y., Yao, H., & Zhao, S. (2016). Auto-encoder based dimensionality reduction. *Neurocomputing*, *184*, 232–242.

Williams, H. C. (1977). On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and planning A*, *9*(3), 285–344.