Discovering preference structure from choice data

Cristóbal De La Maza (cdelamaz@andrew.cmu.edu)[1]
Alex Davis (alexdavis@cmu.edu)[1]
Cleotilde Gonzalez (coty@cmu.edu)[2]
Inês 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

September 21, 2019

Approaches that elicit preference from individual choices often assume that all decision-makers know what they want. That is true if every decision-maker can consistently order the available alternatives, yielding transitive preferences. We propose that decision-makers may be heterogeneous in their preference structure, then cluster decision-makers based on structural heterogeneity using recent advances in graph matching and non-linear embeddings. We characterize heterogeneity of both the content and structure of preferences using two pairwise comparison experiments: the first a classic study of risky choice and the second a two-attribute study about CO_2 mitigation. Decision-makers frequently choose in a way consistent with transitivity, yet some decision-makers make choices consistent with heuristic rules. Furthermore, some participants appear to be uncertain about their preferences, exhibiting violations of transitivity. As a generalization of traditional preference analysis, our approach can be used to make recommendations for those with consistent preferences, uncover complex choice rules, and suggest paths towards clarification for those who are uncertain.

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

1 Introduction

The recent emergence of large 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 understand the determinants of individual choice (Anderson & Lebiere, 2014). For example, many recommender systems use the similarity of individuals' past choices to make suggestions about products (Goldberg, Roeder, Gupta, & Perkins, 2001), estimates of the value of a statistical life are modeled using econometric estimates of the compensation individuals require in exchange for doing a job that has a higher risk of death (Viscusi & Aldy, 2003; Alberini, Hunt, & Markandya, 2006; Hammitt, 2000), and votes are tallied under the assumption that each vote contains a well-defined expression of the voter's preference (Roberts, 1980).

In these important choice situations there could be as many decision rules as there are decision-makers. Such heterogeneity cannot be captured with existing modeling approaches. For example, the mixed logit model assumes that decision-makers all use the same decision rule, varying on the degree to which individuals weight different attributes (McFadden & Train, 2000). Latent class models, on the other hand, assume that decision-makers use different decision rules, but require that those decision rules are known ahead of time, potentially missing patterns not previously considered (Greene & Hensher, 2003). New methods are required to uncover heterogeneity in decision rules across large populations of decision-makers (McFadden, 2001).

There are two ways to capture preference heterogeneity. The first examines whether decision-makers choose the *same alternatives*, or whether they have the same *preference content*. A decision-maker who chooses political candidate A over B and B over C has the same preference content as another that chooses A over B and B over C. The second type of heterogeneity is about the *relationships* between the things that are chosen, or the preference *structure*. A decision-maker that chooses A over B, B over C, and A over C has the same structure as a decision-maker that chooses C over B, B over A, and C over A, because both decision-makers have *transitive* preferences, even though their preference content differs. Prior work typically restricts preference heterogeneity analysis to differences in preference content.

Individual choice behavior reveals preferences that are consistent with utility maximization only if decision-makers can order all the available alternatives (Von Neumann & Morgenstern, 1944; Arrow, 1951), and are not susceptible to subtle but inconsequential changes in how the alternatives are described or presented (framing effects, context effects, reference dependence) (Tversky & Kahneman, 1981; Bhatia, 2013). 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 (Tversky & Kahneman, 1981; Bhatia, 2013). One reason for these deviations is that the burden of selecting the best alternative among a large set, where considering the potential costs and benefits of each alternative, is too difficult (Fischhoff, 2005), forcing individuals to use short-cuts or heuristics to make their choices (Simon, 1972; Gigerenzer & Todd, 1999; Payne, Bettman, & Johnson, 1993). One 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, but there is both theoretical and empirical evidence that choice by lexicographic semiorder can be intransitive (Tversky, 1969).

In simple decisions, for example between two alternatives, each with two 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 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 (Marschak, 1959; Davis-Stober, 2009). If inconsistency arises only in the short-run, then giving decision-makers more time or more opportunities to choose will lead to stable preferences (Busemeyer & Townsend, 1993). Yet not all choice inconsistencies can be characterized as random deviations from well-ordered preferences or failures to reach long-run stability (Tversky & Kahneman, 1981; Bhatia, 2013).

We introduce a new approach that classifies decision-makers into discrete groups based on the pattern of their choices, without having to specify the decision rules used by those groups ahead of time. This approach allows clusters of decision-makers with similar preference patterns to emerge, even when those patterns are inconsistent with prior behavioral theories, potentially improving choice prediction, and suggesting new theories to researchers. Tests of structural properties of preference, such as weak stochastic transitivity, can then be conducted within each cluster (Regenwetter et al., 2014). The approach extends prior work on non-parametric preference clustering that focuses on preference content (Brown, Park, Steinley, & Davis-Stober, 2018), by clustering based on both the content and structure of decision-maker preferences. In what follows we describe the method and apply the approach across two experiments to classify decision-makers into homogeneous preference groups. The next section describes the method in detail.

2 Discovering preference structure

To cluster decision-makers based on their preferences, 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 cluster decision-makers based on these embeddings to propose a decision rule for each cluster. Figure 1 summarizes the four steps for a simulated sample of 100 decision-makers.

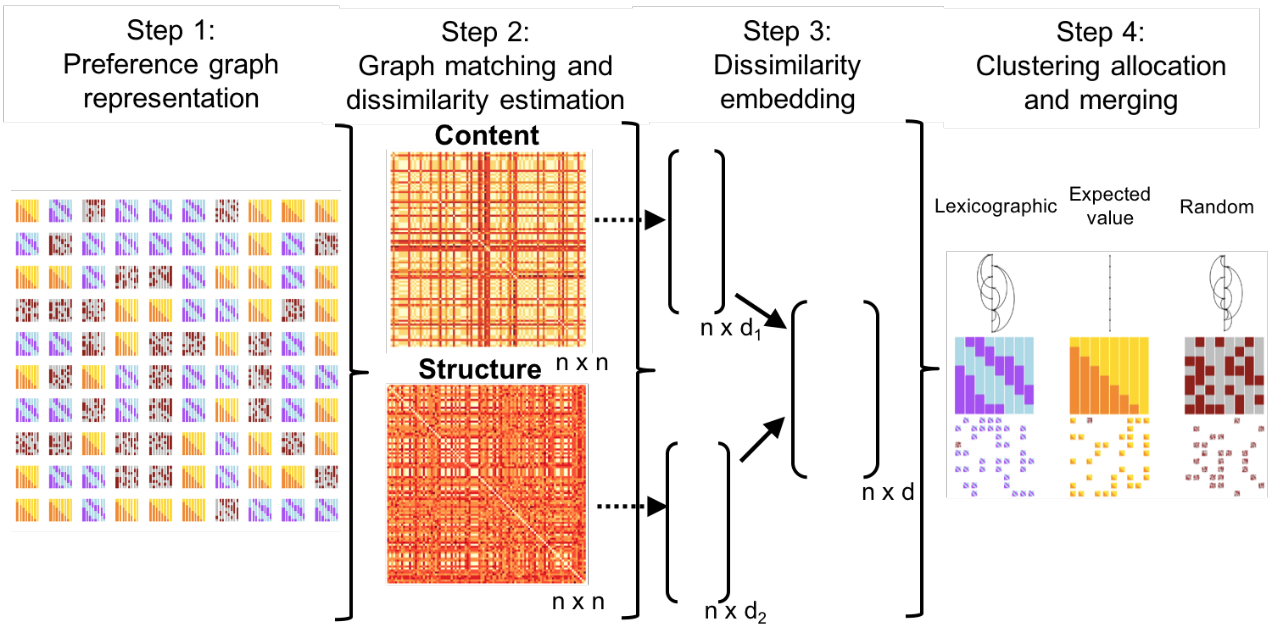

Figure 1: Method summary. The scheme summarizes our method four steps for a simulated sample of 100 decision-makers. First, we represent choices as preference graphs. Next, we compute dissimilarities on both content and structure. Further, we estimate a lower dimensional embedding for each dissimilarity matrices. Finally, we find clustering allocations.

2.1 Preference representation as graphs

First, we represent sequences of individual choices for a single decision-maker as a *preference graph* $G = (V, E)$ (see step 1 in Figure 1), which consists of a set of vertices V and edges E where vertices represent *alternatives* and edges represent binary *preference relations* between alternatives, as illustrated in Figure 2. For all pairs of alternatives a and b, one and only one of the following three preference relations holds (Bouyssou & Vincke, 2010): 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* (Moon, 2015), 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 2a, 2b, 2c, and 2d (Davis, 1954). 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 2e, 2f, 2g, and 2h, with reflexive preferences (along the main diagonal) omitted.

To examine their structure, preference graphs can be unlabeled and 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 (Moon, 2015). 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 2i, and is consistent with classical utility maximization (Varian, 1983; Afriat, 1972). In contrast, the lexicographic semiorder is a preference graph that can contain cycles (Tversky, 1969), such as those shown in Figures 2j, 2k, and 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 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

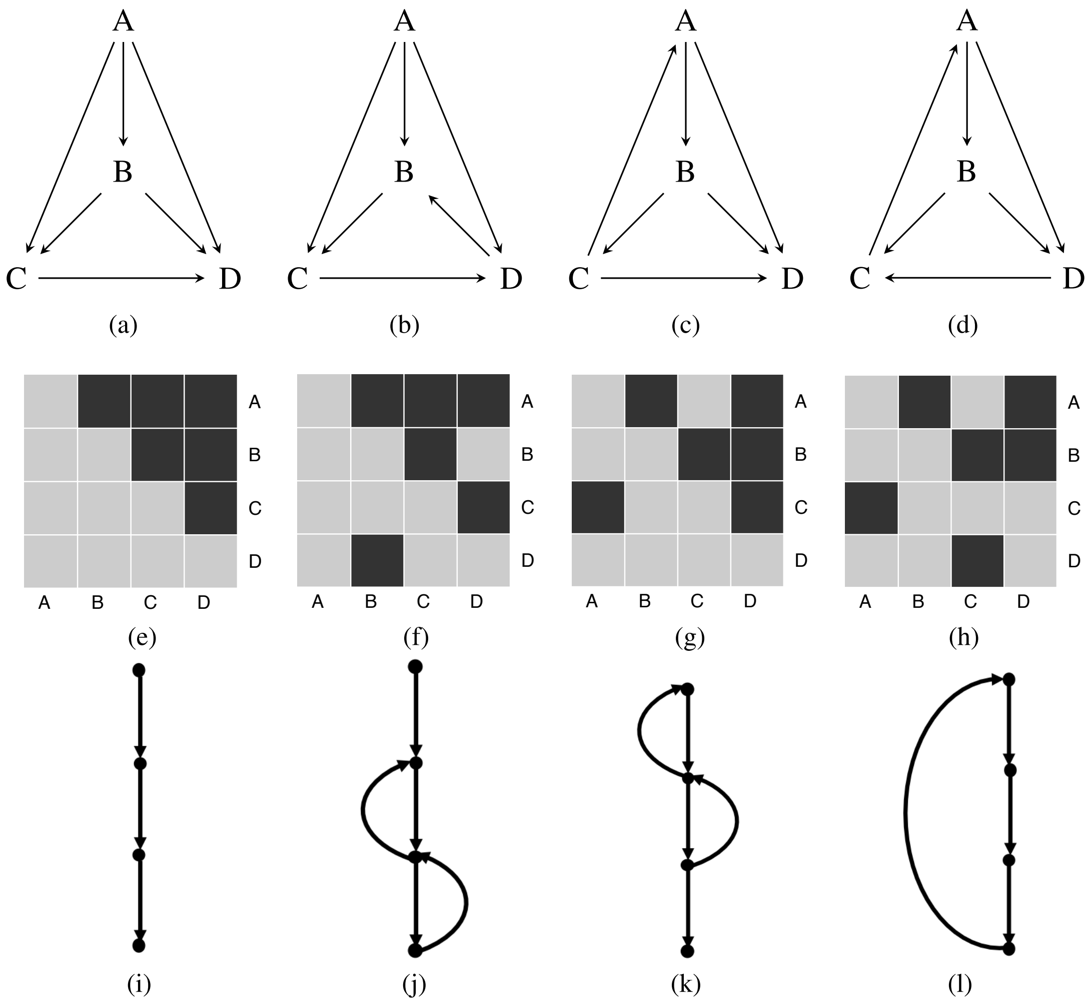

Figure 2: Tournament graphs for four alternatives. First row, preference graphs. Second row, adjacency matrices. Third row, unlabeled tournament structures over four alternatives can be defined by their score vectors: chain ($s = [3,2,1,0]$) (i), cycle at bottom ($s = [3,1,1,1]$) (j), cycle at top ($s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

consistently rank only the best alternative. A decision-maker with the long cycle has a consistent ranking over any subset of three alternatives, but no global ranking over all four.

The preference graph representation uses binary relations (0's and 1's) to represent preference between two alternatives. The representation can be extended to weighted edges (in the unit interval [0, 1]) if choice probabilities are available. Those binary choice probabilities could be elicited from an individual (through repeated choices), or by aggregating binary preference graphs from many individuals who have the same preference content and structure. We use the latter weighted graph approach to represent clusters of decision-makers.

2.2 Graph matching and dissimilarity estimation

Our primary analytical tool, depicted in Step 2 of Figure 1, is a method of calculating the *distance* between graphs. As described in the next two sections, we calculate this distance for both the labeled graph (representing the content of preferences), and unlabeled graph (representing the preference structure).

2.2.1 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* $d_H(G_1, G_2) = ||vec(G_1) - vec(G_2)||_1$ (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, 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 (Coombs & Kao, 1960) and closely related to the approach proposed by Brown *et al.* (Brown et al., 2018) who focus on preference content.

2.2.2 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 (Aflalo, Bronstein, & Kimmel, 2015), meaning there is a bijection $f : V_1 \to V_2$ 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 (Babai & Luks, 1983), 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* (Babai & Luks, 1983). The minimum Hamming distance between two graphs across their automorphism groups then 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. Without this isomorphism step, two decision-makers who make similar choices, but one is intransitive and the other is not, have a chance of being classified similarly. This is because in many cases one only needs to reverse a few edges in an intransitive graph to make it transitive. Including the distance between two graphs over their isomorphisms emphasizes the structural component of preference.

With a few alternatives the structural distance between graphs can be quickly calculated using exhaustive search. However, the problem is NP-hard (Aflalo et al., 2015), requiring approximation techniques for large graphs with more than 8 alternatives. 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 set of permutations that makes two adjacency matrices A_1 and A_2 as similar as possible. The objective function is (Aflalo et al., 2015; Livi & Rizzi, 2013; Vogelstein et al., 2015):

$$\mathrm{P}^* = \underset{\mathrm{P} \in \mathcal{P}}{\operatorname{argmin}} \; f(P) = \mathrm{dis}_{A_1 \to A_2}(\mathrm{P}) = ||A_1 - P^T A_2 P|| \tag{1}$$

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 know as quadratic assignment (QAP) with non-deterministic polynomial time complexity (Koopmans & Beckmann, 1957). 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 (Liu, Qiao, Jia, & Xu, 2014; Aflalo et al., 2015). Because this relaxation can lead to inaccurate results (Aflalo et al., 2015), we instead use Vogelstein's approach (rGM) (Vogelstein et al., 2015) that replaces the objective function $f(P)$ by the identity $-tr(A_1 \mathrm{P} A_2^T \mathrm{P}^T)$, leading to a non-convex problem, where $\nabla^2 f(P) = B \otimes A_1 + A_2^T \otimes A_1^T$ is not positive definite (Vogelstein et al., 2015). Vogelstein *et al.* proposed to solve this problem sequentially with Frank-Wolfe algorithm (Frank & Wolfe, 1956; Vogelstein et al., 2015). We initialized the optimization with the QCV solution (Lyzinski et al., 2016).

2.3 Lower dimensional dissimilarity embedding

In step 3 in Figure 1, we embed preference graph dissimilarities into a lower dimensional space, taking a very high dimensional similarity matrix across all decision-makers (e.g., a 200×200 matrix for 200 decision-makers) and summarizing that matrix with only a few dimensions. The embedding enables subsequent clustering techniques, helping avoid too many clusters and thus overestimation of heterogeneity. 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 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 (Kevin, 2012; Karatzoglou, Smola, Hornik, & Zeileis, 2004). 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 (Goodfellow, Bengio, & Courville, 2016), 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 $\hat{x}(h)$ (Goodfellow et al., 2016). 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, we show the autoencoder network. We selected the number of dimensions in the embedding d using at elbow point of the loss function, and we pretrained the autoencoder with a Restricted Boltzmann Machine (Hinton & Salakhutdinov, 2006; Wang, Yao, & Zhao, 2016).

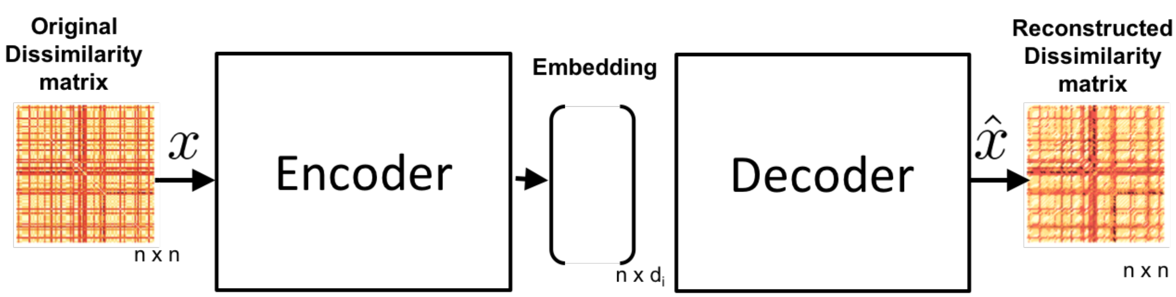

Figure 3: Autoencoder summary. Each original dissimilarity matrix is encoded into a lower dimensional space minimizing reconstruction error (Wang et al., 2016).

2.4 Clustering allocation and merging

Next, we cluster decision-makers based on dissimilarity embeddings (see step 4 in Figure 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 (Singh, Yadav, & Rana, 2013), 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 (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 cluster 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.

2.5 Within-cluster modeling and prediction

Lastly, also in step 4 in Figure 1, 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 (McFadden, 1973), 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 use the logit model as a simple way to summarize the differences between the groups rather than an accurate description of their decision process, which would require complicated (and possibly unknown) parametric functions.

Another reason we use the logit model is to highlight potential bias from aggregating across clusters, where coefficients that differ strongly across clusters would change conclusions if inappropriate aggregation is used across clusters.

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 Empirical analysis

To test our method, we collected choice data from Amazon Mechanical Turk (MTurk) workers for two stated preference tasks: 1) choices between two risky options based on a classic study by Tversky (Tversky, 1969) (*transitivity task*), and 2) choices between electricity generation options for one's state that trade-off CO_2 emissions and electricity bill impacts (Sergi, Davis, & Azevedo, 2018) (*CO_2 task*). Prior work has found that data from the MTurk platform is comparable to laboratory experiments (Crump, McDonnell, & Gureckis, 2013). For both tasks, attributes were presented graphically. Figure 4 shows an example of each choice task.

For each task, we recruited 200 MTurk participants with inclusion criteria: age of at least 18

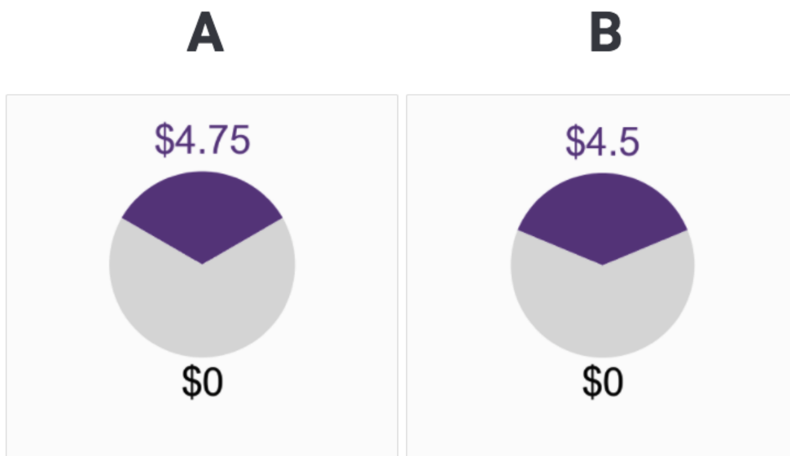

(a) Transitivity task

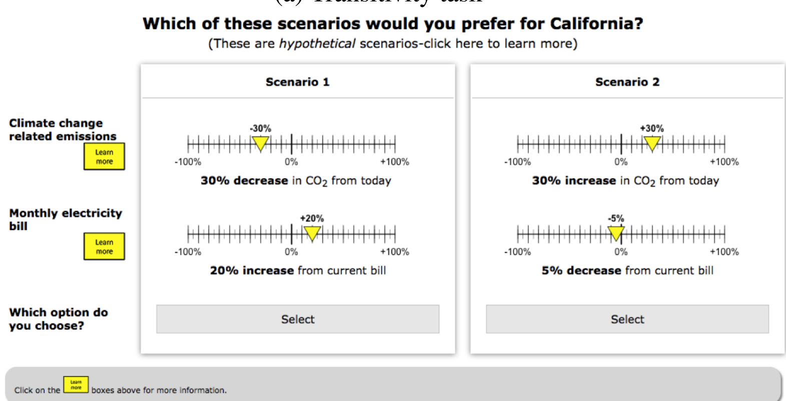

(b) CO_2 task

Figure 4: Choice set examples in each task. Transitivity task: Choice set example alternative b vs. c in Table 1(a). CO_2 task: Choice set example a vs. b in Table 1(a).

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. Table 1(a), presents the selected alternatives for each experiment. Problems for the transitivity task were based on gambles from (Tversky, 1969). Problems for the CO_2 task were based on alternatives from (Sergi et al., 2018). Participants were presented with all pairs of alternatives. Table 1(b) and Table 1(c) presents each pair of problems represented in adjacency matrix form.

Table 1: Alternatives selected for the experimental tasks and problem ID in adjacency matrix. Problems for the transitivity task based on gambles from (Tversky, 1969). Problems for the CO_2 task based on alternatives from (Sergi et al., 2018).

(a) Alternatives selected

	Transitivity task			CO_2 task	
	Probability	Payoff ($)	Expected Value ($)	CO_2	Bill
a	7/24	5.00	1.458	-30%	20%
b	8/24	4.75	1.583	-30%	5%
c	9/24	4.50	1.688	-25%	4%
d	10/24	4.25	1.771	-20%	3%
e	11/24	4.00	1.833	-15%	2%
f	12/24	3.75	1.875	-10%	1%
g	13/24	3.50	1.894	30%	-20%
h	14/24	3.25	1.895	30%	-5%
i	15/24	3.00	1.875		
j	16/24	2.75	1.833		

(b) Transitivity task

$$
\begin{array}{c}
\begin{array}{cccccccccc} a & b & c & d & e & f & g & h & i & j \end{array} \\
\left(\begin{array}{cccccccccc}
 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\
 & & 10 & 11 & 12 & 13 & 14 & 15 & 16 & 17 \\
 & & & 18 & 19 & 20 & 21 & 22 & 23 & 24 \\
 & & & & 25 & 26 & 27 & 28 & 29 & 30 \\
 & & & & & 31 & 32 & 33 & 34 & 35 \\
 & & & & & & 36 & 37 & 38 & 39 \\
 & & & & & & & 40 & 41 & 42 \\
 & & & & & & & & 43 & 44 \\
 & & & & & & & & & 45 \\
 & & & & & & & & &
\end{array}\right)
\begin{array}{c} a \\ b \\ c \\ d \\ e \\ f \\ g \\ h \\ i \\ j \end{array}
\end{array}
$$

(c) CO_2 task

$$
\begin{array}{c}
\begin{array}{cccccccc} a & b & c & d & e & f & g & h \end{array} \\
\left(\begin{array}{cccccccc}
 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\
 & & 8 & 9 & 10 & 11 & 12 & 13 \\
 & & & 14 & 15 & 16 & 17 & 18 \\
 & & & & 19 & 20 & 21 & 22 \\
 & & & & & 23 & 24 & 25 \\
 & & & & & & 26 & 27 \\
 & & & & & & & 28 \\
 & & & & & & &
\end{array}\right)
\begin{array}{c} a \\ b \\ c \\ d \\ e \\ f \\ g \\ h \end{array}
\end{array}
$$

3.1 Transitivity in risky choices

In the *transitivity task*, participants chose between the pairs of gambles included in Table 1(a), from Tversky's classic paper on intransitive preferences (a-e) (Tversky, 1969), along with five additional gambles (f-j), where last two gambles a higher probability is negatively correlated with a higher expected value. As shown in Table 1(b), participants were presented all pairs from 10 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%) passed the attention check questions and were provided with the bonus payment .

Using our approach, we identify six clusters: four with chain structures, one with a small cycle, and one with multiple cycles. Figure 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 5b shows the predicted probabilities from the logistic regressions that fit the data the best in each cluster. For details refer to Table 2(a) in the appendix section, where we test three decision rules to explain decision-maker choice behavior in each cluster: 1) maximize expected value $V_{i1} = \beta_1 \times \mathrm{EV}_i$, 2) maximize probability of winning $V_{i2} = \beta_2 \times P(\text{winning})_i$, 3) and maximize payoffs $V_{i3} = \beta_3 \times \mathrm{Payoff}_i$.

We find 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 choose the alternative with the higher payoff (93% of the time). Although decision-makers in cluster 5 and 6 show a cyclic structure, the

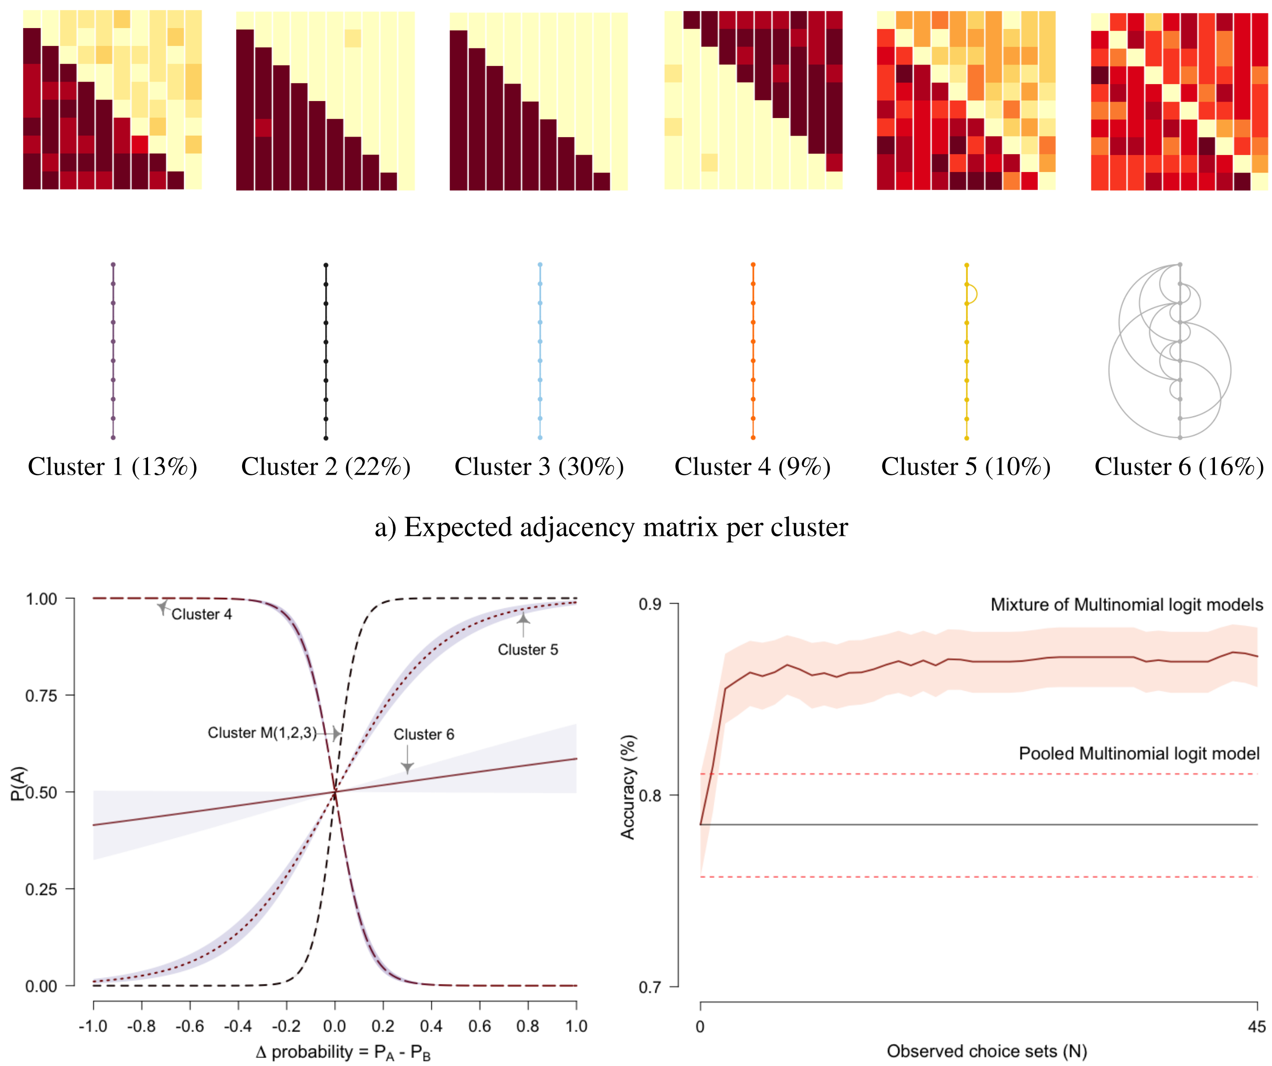

a) Expected adjacency matrix per cluster

b) Logit probabilities

c) Model accuracy

Figure 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.

proportion of choices favoring the option with the higher probability of winning is significantly different from 50%, suggesting their choices were not entirely random.

Figure 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 use a single-attribute choice rule (clusters 1-4), simplifying the task, and leading to transitive preferences within-cluster. Decision-makers whose choices can not be easily explained by a single attribute are also more likely to have intransitive preferences (clusters 5 and 6).

3.2 Preferences for CO_2 emission reductions

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 (Herriges & Kling, 1999; Greene, 2012), substitution patterns (Hensher, Rose, & Greene, 2015), implicit discount rates (Min, Azevedo, Michalek, & de Bruin, 2014), willingness-to-pay (McFadden, 1999; Helveston et al., 2015), and consumer's surplus (Small & Rosen, 1981; Williams, 1977). In the CO_2 *task*, we collected data based on a recent paper by Sergi *et al.* (Sergi et al., 2018), who elicit preferences for CO_2 emission reductions. In our extension of their experiment, participants were 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 1(c), 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.

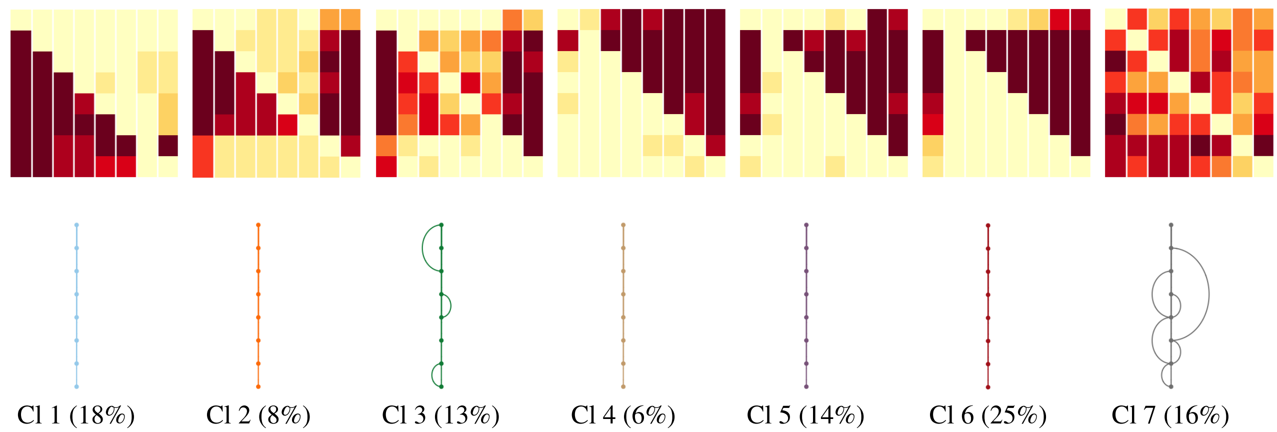

a) Expected adjacency matrix per cluster (Cl: Cluster)

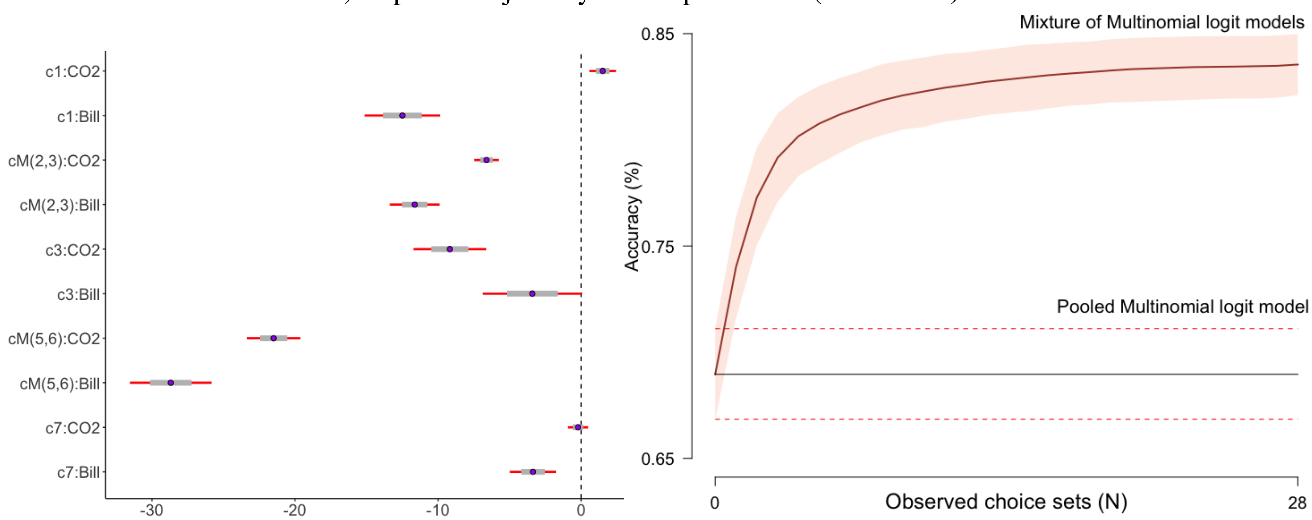

b) Coefficients

c) Model accuracy

Figure 6: 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 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, 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.

Figure 6a 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 6a, 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. To model behavior in each cluster, we used a weighted additive linear utility model with no intercept over both attributes (bill and CO_2) as $V_i = \beta_{Bill} \cdot Bill_i + \beta_{CO_2} \cdot CO_{2i}$ (see Table 2(b) in the appendix section). In Figure 6b we present coefficient values for the weighted additive linear utility model in 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 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 Classification of decision rules

To better understand the relationship between the clusters for each task, we use a hierarchical clustering approach (Kemp, Tenenbaum, Griffiths, Yamada, & Ueda, 2006). We applied hierarchical clustering using the Hamming distance between expected adjacency matrices as the dissimilarity metric. As shown in Figure 7a 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 (Birnbaum, 2011), 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.

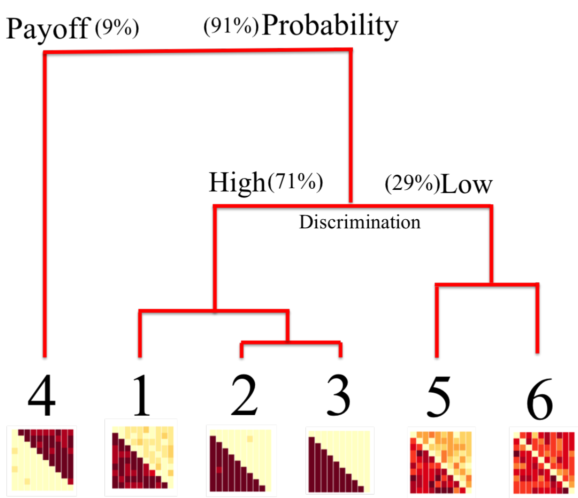

(a) Transitivity task

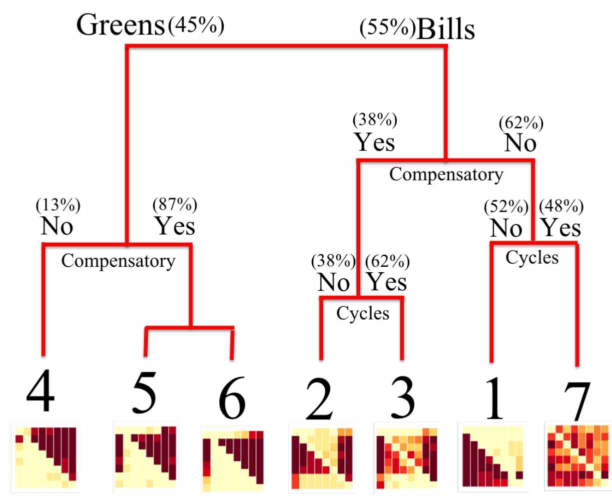

(b) CO_2 task

Figure 7: Hierarchical clustering on the expected adjacency matrices for each cluster in the transitivity and CO_2 tasks.

For the CO_2 task (see Figure 7b), 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). Lastly, some individuals were uncertain about their preferences, as in cluster 7, who had intransitive cycles.

4 Discussion

Policy decisions that rely on understanding the preferences of individuals, groups, and society may benefit from preference models that allow heterogeneity in individual decision rules. Existing approaches require that those rules are known and specified ahead of time. In this work, we use a novel strategy to discovery heterogeneity in decision rules automatically. We use graph matching to develop a method of clustering decision-makers based on the content and structure of their decision rules. By clustering decision-makers into groups with homogeneous preferences, researchers can better test their models against choice patterns within each group, for example by comparing utility maximization to a lexicographic semiorder. Knowledge of these groups also improves prediction of the behavior of new decision-makers when only a few choices from those decision-makers are available. The approach can aid policy analysis, by allowing subgroups with heterogeneous preference content and structure to express the trade-offs that they are willing (or not willing) to make. Further, the approach is not limited to preference analysis, and can be used in any task where decision-makers must choose between two alternatives.

We explore the approach with a classic experiment of risky decisions, and a new policy-relevant stated preference task. In the classic experimental design by Tversky, *the transitivity task*, we find that the the vast majority of the sample uses a single attribute (up to noise) to choose, where 65% of the sample chose only based on the probability of winning a gamble, and 9% chose only based on the gamble's payoff (Birnbaum & Gutierrez, 2007; Lichtenstein & Slovic, 1971; Birnbaum & Gutierrez, 2007; Birnbaum, 2011), undermining the plausibility of other more complex rules like expected value calculations (Brandstätter, Gigerenzer, & Hertwig, 2006; Gigerenzer & Goldstein, 1999; Slovic, 1975; Fishburn, 1971). The data are consistent with a lexicographic order, where decision-makers use only one attribute unless there are exact ties (Fishburn, 1971), rather than a lexicographic semiorder which allows for inexact ties within a just-noticeable-difference (Tversky, 1969). Surprisingly, even though decision-makers used a simplifying choice rule, almost all of them had transitive preferences (Regenwetter, Dana, & Davis-Stober, 2011), likely because a simplifying single-attribute lexicographic decision rule makes consistency (and transitivity) a foregone conclusion. This is a surprising finding given most of the literature using these stimuli finds intransitive decision-makers. What we find is that there are small pockets of intransitive decision-makers (they are not choosing randomly according to attention checks). The conclusion seems to be more balanced than previous literature suggests: some people know what they want, many people disagree with each other, and a few are unsure about what they want.

We also apply our method to a policy-relevant choice task, *the CO_2 task*, where we 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. A naive approach that does not take non-compensatory behavior into account, for example by fitting a single multinomial logit model (MNL) with two attributes on the full sample, would lead to significantly distorted policy conclusions (McFadden, 1973, 1997). 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 $V_i = -8.3 \cdot Bill - 4.5 \cdot CO_2$ 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. 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.

Moreover, our method can improve choice prediction. In both experiments, predictive accuracy increased around 10% when only a few choices were available to estimate mixing probabilities. In Figures 5 and 6, 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.

There are also several important 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 point of the loss function, a useful heuristic (Goodfellow et al., 2016), but allowing for the number of dimensions to be determined automatically by the data in the optimization process would be an important improvement (Côté & Larochelle, 2016). 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.

5 Conclusions

Under the premise that policy decisions ought to be based on trade-offs decision-makers would make between different private or public goods (Arrow et al., 1996), 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.

Acknowledgments

We 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.

Appendix: Modeling results for each experiment

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, *WTP*: Willingness to pay. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

(a) Transitivity task

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%)

(b) 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%)

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

Alberini, A., Hunt, A., & Markandya, A. (2006). Willingness to pay to reduce mortality risks: evidence from a three-country contingent valuation study. *Environmental and Resource Economics*, *33*(2), 251–264.

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.

Arrow, K. J., Cropper, M. L., Eads, G. C., Hahn, R. W., Lave, L. B., Noll, R. G., ... others (1996). Is there a role for benefit-cost analysis in environmental, health, and safety regulation? *Science*, *272*(5259), 221–222.

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).

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, *120*(3), 522–543.

Birnbaum, M. H. (2011). Testing mixture models of transitive preference: Comment on regenwetter, dana, and davis-stober (2011). *Psychological Review*, *118*(4), 675–683.

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.

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.

Brown, N., Park, S., Steinley, D., & Davis-Stober, C. P. (2018). Modeling between-subject variability in decision strategies via statistical clustering: A p-median approach. *Journal of Behavioral Decision Making*, *31*(2), 250–264.

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*, *28*(7), 1265–1288.

Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating amazon's mechanical turk as a tool for experimental behavioral research [Journal Article]. *PloS one*, *8*(3), e57410.

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

Davis-Stober, C. P. (2009). Analysis of multinomial models under inequality constraints: Applications to measurement theory. *Journal of Mathematical Psychology*, *53*(1), 1–13.

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.

Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason: Take the best and its relatives. In G. Gigerenzer, P. Todd, & ABC Research Group (Eds.), *Simple heuristics that make us smart.* (pp. 75–95). Oxford University Press, New York.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart* [Book]. Oxford University Press, USA.

Goldberg, K., Roeder, T., Gupta, D., & Perkins, C. (2001). Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, *4*(2), 133–151.

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.

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

Hammitt, J. K. (2000). Valuing mortality risk: Theory and practice. *Environmental Science and Technology*, *34*, 1396–1400.

Hartigan, J. A., & Wong, M. A. (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., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, *313*(5786), 504–507.

Karatzoglou, A., Smola, A., Hornik, K., & Zeileis, A. (2004). kernlab-an s4 package for kernel methods in r. *Journal of statistical software*, *11*(9), 1–20.

Kemp, C., Tenenbaum, J. B., Griffiths, T., Yamada, T., & Ueda, N. (2006). Learning systems

of concepts with an infinite relational model. In *Proceedings of the national conference on artificial intelligence* (Vol. 21, p. 381).
Kevin, M. (2012). *Machine learning: a probabilistic perspective*. The MIT press, Cambridge, Massachusetts.
Koopmans, T. C., & Beckmann, M. (1957). Assignment problems and the location of economic activities. *Econometrica: journal of the Econometric Society*, 53–76.
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*.
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.
Marschak, J. (1959). *Binary choice constraints on random utility indicators* (Tech. Rep.). Cowles Foundation for Research in Economics, Yale University.
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. (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. *American economic review*, *91*(3), 351–378.
McFadden, D., & Train, K. E. (2000). Mixed mnl models for discrete response. *Journal of applied Econometrics*, 447–470.
Min, J., Azevedo, I. L., Michalek, J., & de Bruin, W. B. (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.
Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review*, *118*(1), 42–56.
Regenwetter, M., Davis-Stober, C. P., Lim, S. H., Guo, Y., Popova, A., Zwilling, C., … Messner, W. (2014). Qtest: Quantitative testing of theories of binary choice. *Decision*, *1*(1), 2.
Roberts, K. W. (1980). Interpersonal comparability and social choice theory. *The Review of Economic Studies*, *47*(2), 421–439.

Sergi, B., Davis, A., & Azevedo, I. (2018). The effect of providing climate and health information on support for alternative electricity portfolios. *Environmental Research Letters*, *13*(2), 024026.

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).

Slovic, P. (1975). Choice between equally valued alternatives. *Journal of Experimental Psychology: Human Perception and Performance*, *1*(3), 280.

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.

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

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458.

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

Viscusi, W. K., & Aldy, J. E. (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of risk and uncertainty*, *27*(1), 5–76.

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* (Vol. 60) [Book]. Princeton: Princeton university press.

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

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