



The limits of identification in discrete choice [☆]

Christopher P. Chambers ^a, Christopher Turansick ^{b,*}

^a Department of Economics, Georgetown University, United States of America

^b Department of Decision Sciences and IGIER, Università Bocconi, Italy



ARTICLE INFO

Keywords:

Random utility
Stochastic choice
Identification

ABSTRACT

This paper uncovers tight bounds on the number of preferences permissible in identified random utility models. We show that as the number of alternatives in a discrete choice model becomes large, the fraction of preferences admissible in an identified model rapidly tends to zero. We propose a novel sufficient condition ensuring identification, which is strictly weaker than some of those existing in the literature. While this sufficient condition reaches our upper bound, an example demonstrates that this condition is not necessary for identification. Using our new condition, we show that the classic “Latin Square” example from social choice theory is identified from stochastic choice data.

1. Introduction

The random utility model is the basis for much of modern discrete choice analysis. Random utility moves beyond the classic rational model and supposes that there is some distribution over preferences inducing choice. This distribution over preferences is typically interpreted as heterogeneity within a population or variation of a single agent’s preference across time. A common desideratum for a model in empirical practice is identification. Identification allows for proper counterfactual and welfare analysis. The random utility model is well-known to be unidentified, however there are many refinements of the random utility model which recover identification. We study those models which recover identification by limiting the set of allowable preferences in the model.

Identifying assumptions by nature put restrictions on a model in order to recover identification. We study identifying restrictions which assert that preferences must belong to some prespecified class. Our main result shows that the maximum number of preferences allowed for a set of preferences leading to an identified random utility model is equal to $(n - 2)2^{n-1} + 2$, where n is the number of alternatives in our environment. It follows that the ratio of preferences in such a model to the total number of preferences given a set of alternatives of size n scales at a rate of $\mathcal{O}\left(\frac{2^{n-1}}{(n-1)!}\right)$. This tells us that, as n grows large, the number of preferences in a model identified through preference restrictions is vanishingly small when compared to the total number of possible preferences. When $n = 5$, any identified model must rule out at least half of the possible preferences. When $n = 9$, more than 99.5% of the preferences must be discarded in order to write an identified model. Thus identifying assumptions of this nature must assume away almost all types of behavior.

Our second goal in this paper is to offer a new identifying restriction for random utility models. Our new identifying restriction is called *edge decomposability*. In our setup a random utility model is simply a collection of feasible preferences. Edge decomposability

[☆] We are thankful to two anonymous referees, the advisory editor, Roy Allen, and Yusufcan Masatlioglu for helpful comments and discussions during the course of this project.

* Corresponding author.

E-mail addresses: cc1950@georgetown.edu (C.P. Chambers), christopher.turansick@unibocconi.it (C. Turansick).

<https://doi.org/10.1016/j.geb.2025.02.006>

Received 30 April 2024

asks that, for a given model, for every non-empty subset of preferences of that model, there exists some preference in that subset with an upper contour set unique to that preference within that subset. More specifically it asks that for every non-empty subset of preferences, there exists some pair (x, A) , where x is an alternative in menu A , such that there is a single preference in this subset that chooses x from A but fails to choose x from any strict superset of A . Edge decomposability is a sufficient condition for identification of a model. Among identified models, edge decomposability has large explanatory power. This is due to the fact that the largest edge decomposable model is the same size as the largest identified model. However, we are able to show through an example that there are identified models which are not edge decomposable.

We then compare edge decomposability with the identifying restrictions of Apestegua et al. (2017) and Turansick (2022). We show that every set of preferences which can be the support of some distribution over preferences satisfying the conditions of either Apestegua et al. (2017) or Turansick (2022) must also be edge decomposable. Further, we show that there are sets of preferences that are edge decomposable yet cannot be the support of a distribution satisfying either the conditions of Apestegua et al. (2017) or Turansick (2022). Finally, we introduce a new model of random utility which captures the classic Latin square example from social choice. The Latin square example is a collection of preferences that lead to Condorcet cycles when aggregating preferences via majority rule. We use our edge decomposability condition to show that this model is identified from stochastic choice data.

The rest of this paper proceeds as follows. In Section 2 we introduce our setup and some preliminary constructions. In Section 3 we provide our main result and discuss its proof. In Section 4 we introduce our identifying restriction and compare it to the restrictions of Apestegua et al. (2017), Turansick (2022), as well as the Latin square example. Finally, we conclude with a discussion of the related literature in Section 5.

2. Model and preliminaries

We work with a finite set of alternatives $X_n = \{1, \dots, n\}$ where each alternative is indexed by a number between 1 and n . Further, we use x and y to denote arbitrary elements of X_n . We are interested in strict preferences over this environment. Denote the set of linear orders over X_n by $\mathcal{L}(X_n)$.¹ A typical linear order is denoted by \succ .

Definition 2.1. A model on X_n (or simply a model) is a nonempty set of linear orders $M: \emptyset \subsetneq M \subseteq \mathcal{L}(X_n)$.

We denote the collection of all models on X_n as \mathcal{M}_n . Our notion of model does not capture every possible type of restriction. In our setup, a model restricts which preferences an agent may have. Other random utility models may be constructed through parametric restrictions.

Example 2.1 (Social Choice and Latin Squares). Suppose there is some exogenous ordering \triangleright over alternatives X_n . Without loss we can let it be $1 \triangleright 2 \triangleright \dots \triangleright n$. We can consider the model M which forms the Latin square given this ordering. In this model M , there is one preference \succ_m for each $m \in \{1, \dots, n\}$. \succ_m is given by $m \succ m + 1 \succ \dots \succ n - 1 \succ n \succ 1 \succ \dots \succ m - 1$. We then consider distributions over this model M . This example considers a restriction on the set of available preferences, and thus is a model in our setup.²

Example 2.2 (Luce Random Choice). Consider the following class of random choice rules, described by Luce (1959). Each $x \in X$ is given a weight $w(x) \in (0, \infty)$. The probability that x is chosen from a menu A is then given by $\frac{w(x)}{\sum_{y \in A} w(y)}$. This class of choice rules is often called the “Luce model,” though it is not a model in our sense. Rather, instead of being parametrized by a distribution over preferences, the parameters in this model correspond to $w(x)$. These parameters are not uniquely identified, but the normalized versions $\frac{w(x)}{\sum_{y \in X} w(y)}$ are.

Using the standard Δ notation to represent the set of probability measures over a set, let $\nu \in \Delta(M)$ denote a typical probability distribution over the preferences in model M . Further, let 2^X denote the power set of X and let A denote a typical element of 2^X .

Definition 2.2. A function $p : X_n \times 2^{X_n} \setminus \{\emptyset\} \rightarrow \mathbb{R}$ is a **random choice rule** if it satisfies the following.

- $p(x, A) \geq 0$ for all $x \in A$
- $\sum_{x \in A} p(x, A) = 1$

Random choice rules correspond to our data primitive. They denote how frequently an alternative x is chosen from a menu A . Given $\nu \in \Delta(\mathcal{L}(X_n))$, let p_ν denote the random choice rule induced by ν . Specifically,³

$$p_\nu(x, A) = \sum_{\succ \in \mathcal{L}(X_n)} \nu(\succ) \mathbf{1}\{x \succ y \forall y \in A \setminus \{x\}\}. \tag{1}$$

¹ A linear order is complete, antisymmetric, and transitive.
² Note that the Latin Square example is sometimes known as a cyclic domain which is commonly used to show the presence of Condorcet cycles when aggregating preferences.
³ Here, $\mathbf{1}\{\cdot\}$ corresponds to the standard indicator function. It returns 1 if the logical statement in the brackets is true and returns 0 otherwise.




Fig. 1. The probability flow diagram for the set $X_n = \{x, y, z\}$.

Definition 2.3. A model M is **identified** if for any $v, v' \in \Delta(M)$, $p_v = p_{v'}$ implies that $v = v'$.

Given X_n , we denote the set of identified models as I_n . Our goal is to provide a tight upper bound on the size of models in I_n . Specifically, our main question is what is the largest number of preferences that admit an identified model. In order to answer this question, we rely on the Möbius inverse of a random choice rule.⁴

Definition 2.4. The **Möbius inverse** of $p(x, A)$ is a function $q : X \times 2^X \setminus \{\emptyset\} \rightarrow \mathbb{R}$ defined as follows.

$$\begin{aligned}
 q(x, A) &= p(x, A) - \sum_{A \subsetneq B} q(x, B) \\
 &= \sum_{A \subseteq B} (-1)^{|B \setminus A|} p(x, B)
 \end{aligned}
 \tag{2}$$

The Möbius inverse $q(x, A)$ simply keeps track of how much choice probability is being added to or removed from the choice of x at menu A . Further, the Möbius inverse function has a special relationship with random choice rules induced by models.

Definition 2.5. A random choice rule p is **stochastically rational** if there exists some $v \in \Delta(\mathcal{L}(X))$ such that $p = p_v$.

Consider the set of preferences which choose x from A but do not choose x from any superset of A . This set is given by $L(x, A) = \{> | z > x > y \forall y \in A \setminus \{x\}, \forall z \in X \setminus A\}$.

Theorem 2.1 (Falmagne (1978)). Let q be the Möbius inverse of p_v . Then for all $A \subseteq X$ and all $x \in A$, $q(x, A) = v(L(x, A))$.

This result tells us that $q(x, A)$ is the probability under p_v of $L(x, A)$. Further, our main result relies on a graphical representation of random choice rules. This directed graph is defined as follows. There is one node of this graph for each $A \in 2^X$. We index these nodes via 2^X and use the set A to refer to the node indexed by A . There is a directed edge connecting A to B if $B = A \setminus \{x\}$ for some $x \in A$. We sometimes use e to refer to an arbitrary edge and $e(A, B)$ to refer to the edge connecting node A to node B . Finally, we assign capacities to each edge. The edge capacity of the edge connecting A and $A \setminus \{x\}$ is given by $q(x, A)$. We call this graphical representation the *probability flow diagram*. See Fig. 1 for an example in the three alternative case.⁵

Definition 2.6. A path ρ is a finite sequence of sets $\{A_i\}_{i=0}^n$ such that $A_{i+1} \subsetneq A_i$ for all i , $A_0 = X$, and $A_{|X|} = \emptyset$.

⁴ The Möbius inverse is a fairly general method of inverting cumulative sums on certain partially ordered sets. In our case, the partially ordered set under consideration is (X_n, \subseteq) , and the inversion formula is referred to as the inclusion-exclusion principle. A classical reference is Rota (1964).

⁵ The probability flow diagram is originally due to Fiorini (2004).

Observe that for every path, $|A_i \setminus A_{i+1}| = 1$. In the probability flow diagram, each path from X_n to \emptyset bijectively corresponds to a preference. Each edge of the path corresponds to a unique input (x, A) for the Möbius inverse. Using Theorem 2.1, a path that goes from A_0 to A_n corresponds to the preference $A_0 \setminus A_1 > A_1 \setminus A_2 > \dots > A_{n-1} \setminus A_n$. To further describe this connection, if a random choice rule p is induced by choice according to $>$, then each $q(x, A)$ on the path corresponding to $>$ is equal to one and each $q(y, B)$ not on the path corresponding to $>$ is equal to zero.

3. The dimension of preferences

In this section we present our main result and proof.

Theorem 3.1. For any n , $\max_{M \in \mathcal{I}_n} |M| = (n - 2)2^{n-1} + 2$.

Theorem 3.1 tells us that, if we have an identified model, then there are at most $(n - 2)2^{n-1} + 2$ preferences in that model. There are models which are identified and have fewer preferences, unidentified and have fewer preferences, and there are models which face intuitive restrictions but are not identified.

Example 3.1 (Social Choice and Latin Squares Revisited). Recall the setup of Example 2.1. If our model forms a Latin square then there are $n \leq (n - 2)2^{n-1} + 2$ preferences in the model, with a strict inequality when $n \geq 3$. We will show later that models that form Latin squares are identified.

Example 3.2 (Single-Crossing Random Utility). Suppose there is some exogenous ordering \triangleright over alternatives X_n for $n \geq 2$. Without loss we can let it be $1 \triangleright 2 \triangleright \dots \triangleright n$. An ordered set of preferences $\{>_i\}_{i=1}^m$ satisfies the single-crossing property if $x \triangleright y$, $x >_i y$, and $j \geq i$ imply $x >_j y$. Apestegua et al. (2017) consider random utility models with single-crossing supports. A consequence of their results is that a model satisfying the single crossing property is identified. The maximal size of a model satisfying the single-crossing property is $\binom{n}{2} + 1 \leq (n - 2)2^{n-1} + 2$ with a strict inequality when $n \geq 3$.

Proposition 3.1. Given a set of alternatives X_n for $n \geq 2$, the maximum size of a set of preferences satisfying the single-crossing property is $\binom{n}{2} + 1$.

Example 3.3 (Fishburn and Identification Failure). The classic counterexample to identification of the random utility model is due to Fishburn (1998). Let $X_n = \{a, b, c, d\}$. Consider the following probability distributions over linear orders on X .

$$v_1(>) = \begin{cases} \frac{1}{2} & \text{if } > \in \{a > b > c > d, b > a > d > c\} \\ 0 & \text{otherwise} \end{cases}$$

$$v_2(>) = \begin{cases} \frac{1}{2} & \text{if } > \in \{a > b > d > c, b > a > c > d\} \\ 0 & \text{otherwise} \end{cases}$$

These two probability distributions induce the same random choice rule. If we consider the model given by the four preferences in this example, we have a model which is not identified. The size of this model is $4 < (4 - 2)2^{4-1} + 2 = 18$.

Example 3.4 (Minimal Mutual Agreement). Consider a model M . We say that a model satisfies *minimal mutual agreement* if for each $>, >' \in M$, there are $x, y \in X_n$ such that $x > y$ and $x >' y$. Minimal mutual agreement may be assumed in cases where an analyst is trying to aggregate preferences. When minimal mutual agreement fails, there is a pair $>, >'$ of preferences that are directly opposed, and then any gain by $>$ must be accompanied by a loss in $>'$. Given $>$, define $>^-$ by $x > y \implies y >' x$. It is easy to see that M satisfies minimal mutual agreement if and only if $(> \in M \implies >^- \notin M)$ holds. It then follows that the maximal size of a model satisfying minimal mutual agreement is $\frac{n!}{2} > (n - 2)2^{n-1} + 2$ (for large enough n).

Before presenting our proof of Theorem 3.1, we first discuss some of its implications. Recall that for a given X_n , there are $n!$ strict preferences. We now compare the number of possible preferences with the maximal number of preferences in an identified model. Since preferences here are linear orders, the number of possible preferences is $n!$. According to Theorem 3.1, for reasonably large n , the size of a maximal identified model roughly doubles when adding a new alternative, while the size of the largest possible model, identified or otherwise, increases by a factor of n . In terms of ratios,

$$\begin{aligned}
 \frac{(n-2)2^{n-1} + 2}{n!} &= \frac{2^{n-1}}{n(n-1)(n-3)!} + \frac{2}{n!} \\
 &\leq \frac{2^{n-1}}{(n-1)!} + \frac{2}{n!} \\
 &\leq \frac{2^{n-1}}{(n-1)!} + \frac{2}{(n-1)!} \\
 &= \frac{2^{n-1} + 2}{(n-1)!}.
 \end{aligned}
 \tag{3}$$

This tells us that the ratio of the maximal number of preferences in an identified model and the total number of preferences grows at a rate $\mathcal{O}(\frac{2^{n-1}}{(n-1)!})$. Clearly, as n gets large, this ratio rapidly goes to zero. Thus, any identified model must rule out a large mass of preferences.

We now discuss the proof of Theorem 3.1. First, observe that there is an important connection between identified models and linear independence.

Definition 3.1. A set of vectors $\{v_1, \dots, v_n\}$ is **linearly independent** if any set of scalars $\{c_1, \dots, c_n\}$ with $\sum_{i=1}^n c_i v_i = 0$ implies that $c_i = 0$ for all i .

Let $p_{>}$ denote the real-valued vector, indexed by elements of the form (x, A) with $x \in A$, which encodes the random choice rule induced by choice according to $>$. In other words, $p_{>}(x, A) = 1$ if $x > y$ for all $y \in A \setminus \{x\}$ and $p_{>}(x, A) = 0$ otherwise.

The following result is straightforward and known, but we give a proof of it in the appendix for completeness.

Lemma 3.1. A model M is identified if and only if the set $\{p_{>}\}_{> \in M}$ is linearly independent.

Now let $q_{>}$ denote the vector, indexed by elements of the form (x, A) with $x \in A$, which encodes the Möbius inverse of the random choice rule induced by choice according to $>$. That is, $q_{>}(x, A) = 1$ if $> \in L(x, A)$ and $q_{>}(x, A) = 0$ otherwise. It is less obvious that a model M is identified if and only if the set $\{q_{>}\}_{> \in M}$ is linearly independent. It essentially follows from Lemma 3.1 as the map carrying $p_{>}$ to $q_{>}$ is known to be linear and invertible. We provide a formal proof of the following result in the appendix.

Lemma 3.2. A model M is identified if and only if $\{q_{>}\}_{> \in M}$ is linearly independent.

Our proof of Theorem 3.1 relies on the probability flow diagram with one small alteration. Specifically, we take the probability flow diagram and add an edge connecting \emptyset to X_n (and assign it an edge capacity equal to one). This motivates the following proposition.

Proposition 3.2. Let $\{x^i\}_{i=1}^I$ be a collection of vectors for which for all i, j , $\sum_k x_k^i = \sum_k x_k^j \neq 0$. Then $\{x^i\}_{i=1}^I$ is linearly independent if and only if $\{(x^i, 1)\}_{i=1}^I$ is linearly independent, where $(x^i, 1)$ refers to the concatenation of the vector x^i with a 1.

In the proof of Theorem 3.1, we make use of two concepts from graph theory. First, a circuit is a sequence of directed edges where one edge’s starting node is the previous edge’s ending node with the last edge’s ending node being the same as the first edge’s starting node. In simpler terms, a circuit is a directed cycle. Second, the cyclomatic number of a directed graph corresponds to the minimal number of edges in the graph needed to be removed in order to remove all cycles from the graph. The cyclomatic number also corresponds to the circuit rank of the graph or the maximum number of independent circuits in the graph. We are now ready to proceed with the proof of Theorem 3.1.

Proof. Terminology in this proof is as in Berge (2001). By Lemma 3.2 and Proposition 3.2, a model M is identified if and only if the set of vectors $\{(q_{>}, 1)\}_{> \in M}$ is linearly independent.⁶

Now, consider the directed graph whose nodes are 2^{X_n} and has a directed edge from A to $A \setminus \{x\}$ for any $x \in A$, and also has a directed edge from \emptyset to X_n . We call this the *appended probability flow diagram*. The vector $(q_{>}, 1)$ then represents an indicator function of a subset of the appended probability flow diagram. The discussion following Definition 2.6 explains that $(q_{>}, 1)$ represents an indicator function of a circuit on the appended probability flow diagram, or an oriented loop (the 1 reflects the edge going from \emptyset to X_n).

By Theorem 3 of Chapter 4 of Berge (2001), the cyclomatic number of the appended probability flow diagram is equal to the dimension of the space spanned by indicator functions of circuits of the appended probability flow diagram. Now, the indicator function of any circuit of this diagram must be a linear combination of indicator functions of circuits corresponding to preferences. The reasoning is straightforward: by construction, every circuit must pass through the edge connecting \emptyset to X_n . A circuit which passes through this edge only once clearly corresponds to a preference. A circuit passing through it k times corresponds to a concatenation

⁶ Observe that for each $>$, $\sum_{(x,A):x \in A} q_{>}(x, A) = n$, so that the condition in Lemma 3.2 is satisfied.

of k circuits passing through this edge only once. Hence, the indicator function of the circuit passing through this edge k times is the sum of the indicator functions of the k circuits which pass through the edge only once, each of which correspond to a preference. Hence, the dimension of the circuit space of the appended probability flow diagram is the same as the dimension of the space spanned by $\{(q_{>}, 1)\}$ for all preferences $>$.

The cyclomatic number of a strongly connected graph, as our appended probability flow diagram is, according to Berge (2001), is defined by $E - N + 1$, where E is the number of edges and N is the number of nodes. There are 2^n nodes in our appended probability flow diagram. There is one edge in our probability flow diagram for each (x, A) with $x \in A$. This is given by $\sum_{i=1}^n i \binom{n}{i}$, which is equal to $n2^{n-1}$.⁷ This means that our appended probability flow diagram has $n2^{n-1} + 1$ edges, accounting for the edge going from \emptyset to X_n . It then follows that the cyclomatic number of our appended probability flow diagram is given by the following:

$$n2^{n-1} + 1 - 2^n + 1 = (n - 2)2^{n-1} + 2. \tag{4}$$

Thus, the space spanned by $(q_{>}, 1)$ has dimension $(n - 2)2^{n-1} + 2$. This tells us that for any model M , the dimension of the space spanned by $(q_{>}, 1)$ as $> \in M$ is at most $(n - 2)2^{n-1} + 2$, and that this bound is achieved for some model. In particular, any M for which $(q_{>}, 1)$ is linearly independent must have $|M| \leq (n - 2)2^{n-1} + 2$ and this bound is achieved for some M . Again by Proposition 3.2, any M for which $q_{>}$ is linearly independent must have $|M| \leq (n - 2)2^{n-1} + 2$ and this bound is achieved. Conclude by Lemma 3.2 that any identified model must have $|M| \leq (n - 2)2^{n-1} + 2$ and that this bound is achieved. \square

The dimension of all random choice rules is the same as the dimension of the set of RUM random choice rules; see e.g. Corollary 6(ii) of Saito (2018) or Theorem 2 of Dogan and Yildiz (2022), where it is shown that the rational choice functions span the span of all stochastic choice functions. This immediately gives us the following corollary to Theorem 3.1, which is not so difficult to prove directly by simple combinatorial arguments.

Corollary 3.1. *The set of random choice rules on X_n has dimension $(n - 2)2^{n-1} + 2$.*

Corollary 3.1 obtains bounds on identified models even absent the RUM restriction.

4. A sufficient condition for linear independence

In this section, using the graphical intuition developed in the last section, we propose a new condition which guarantees that a model is identified. While this condition is sufficient for identification and reaches the upper bound from Theorem 3.1, we show by example that it fails to be necessary. We then compare this condition to two other uniqueness conditions in the random utility literature and the Latin square construction from Example 2.1. We now introduce our new condition.

Definition 4.1. We say that a model M is **edge decomposable** if for every nonempty submodel $N \subseteq M$ there exists a preference $> \in N$ and a tuple (x, A) with $x \in A$ such that $N \cap L(x, A) = \{>\}$.

Our arguments often leverage a form of edge decomposability which takes the form of an induction argument:

Definition 4.2. We say that a model M is **sequentially edge decomposable** if it can be enumerated as $\{>_1, \dots, >_k\}$ so that for each $i \in \{1, \dots, k\}$, there is a tuple (x, A) with $x \in A$ such that $\{>_i, \dots, >_k\} \cap L(x, A) = \{>_i\}$.

A simple induction argument guarantees the following, which we will often use without mention.

Proposition 4.1. *A model M satisfies edge decomposability iff it satisfies sequential edge decomposability.*

Keeping in mind Theorem 2.1, when a model is edge decomposable, we are able to do the following to recover the unique distribution over preferences in our model which induces the data. First, as M is a (weak) submodel of M , there exists some preference $>$ and pair (x, A) such that $\{>\} = M \cap L(x, A)$. By Theorem 2.1, this tells us that $q(x, A) = v(L(x, A)) = v(>)$. Now consider the submodel which is given by $N = M \setminus \{>\}$. By edge decomposability, we know that there is some $>' \in N$ such that $\{>'\} = N \cap L(y, B)$ for some $y \in B$. It then follows that $v(>') = q(y, B)$ if $> \notin L(y, B)$ and $v(>') = q(y, B) - v(>)$ if $> \in L(y, B)$. Since M is finite, we can repeat this argument to get full recovery of our distribution v over preferences in M , and so every edge decomposable model M is identified. The following theorem formally restates this observation and collects the identification properties of edge decomposability.

Theorem 4.1.

⁷ The argument is standard, for each element of a set of size n , there are 2^{n-1} sets containing it. So $\sum_{i=1}^n i \binom{n}{i}$, which counts the sum of cardinalities of all subsets, must be given by summing the number 2^{n-1} across each element of the set of size n , resulting in $n2^{n-1}$. See Feller (1968), p.63.

1. If a model M is edge decomposable, it is identified.
2. The maximum size of an edge decomposable model is $(n - 2)2^{n-1} + 2$.
3. There are identified models which are not edge decomposable.

We now interpret our definition of edge decomposability in terms of the probability flow diagram. Recall that every pair (x, A) with $x \in A$ is associated with a specific edge of the probability flow diagram. Specifically, $q(x, A)$ is assigned as the edge capacity of the edge associated with (x, A) . Edge decomposability then says, given a model M , for each submodel N there exists some path associated with $\succ \in N$ such that this path has some edge which is unique to it among paths associated with preferences $\succ' \in N$. Following a similar process to what was described in the last paragraph, this means that when a graph is written as a combination of flows along paths associated with preferences in \succ in an edge decomposable model M , we can always uniquely decompose this graph into flow weights on paths associated with preferences in M .⁸ While edge decomposability may be sufficient for identification of a model M , it fails to be necessary. We show this through Example 4.1 and Fig. 2.

Example 4.1. Consider the environment given by $X_n = \{a, b, c, d, e, f, g, h\}$. We now provide a model with eight preferences which is not edge decomposable and yet is identified. These preferences are represented on the probability flow diagram of Fig. 2. The paths and preferences are as follows.

1. Black path - $\succ_1 - f > g > d > h > c > e > a > b$
2. Red path - $\succ_2 - h > g > e > f > b > d > a > c$
3. Pink path - $\succ_3 - f > g > h > e > d > c > b > a$
4. Blue path - $\succ_4 - h > g > f > d > c > e > b > a$
5. Aqua path - $\succ_5 - g > f > d > h > e > b > a > c$
6. Green path - $\succ_6 - g > h > f > d > e > b > c > a$
7. Purple path - $\succ_7 - g > f > h > e > b > d > c > a$
8. Orange path - $\succ_8 - g > h > e > f > d > c > b > a$

We now proceed to show that $M = \{\succ_1, \succ_2, \succ_3, \succ_4, \succ_5, \succ_6, \succ_7, \succ_8\}$ is an identified model despite not being edge decomposable. It follows immediately from Fig. 2 that M is not edge decomposable as each edge covered by M is covered by two paths. Now we show identification by directly calculating the values of v . To begin, we start by identifying the probability weight on the preferences \succ_2, \succ_4 , and \succ_8 .

$$\begin{aligned}
 v(\succ_2) + v(\succ_4) &= q(h, \{a, b, c, d, e, f, g, h\}) \\
 v(\succ_2) + v(\succ_8) &= q(e, \{a, b, c, d, e, f\}) \\
 v(\succ_4) + v(\succ_8) &= q(b, \{a, b\})
 \end{aligned}
 \tag{5}$$

This gives us the following.

$$\begin{aligned}
 v(\succ_2) &= \frac{q(h, \{a, b, c, d, e, f, g, h\}) + q(e, \{a, b, c, d, e, f\}) - q(b, \{a, b\})}{2} \\
 v(\succ_4) &= \frac{q(h, \{a, b, c, d, e, f, g, h\}) - q(e, \{a, b, c, d, e, f\}) + q(b, \{a, b\})}{2} \\
 v(\succ_8) &= \frac{-q(h, \{a, b, c, d, e, f, g, h\}) + q(e, \{a, b, c, d, e, f\}) + q(b, \{a, b\})}{2}
 \end{aligned}
 \tag{6}$$

The probability weight on each other preference is then given by the following.

$$\begin{aligned}
 v(\succ_1) &= q(c, \{a, b, c, e\}) - v(\succ_4) \\
 v(\succ_6) &= q(f, \{a, b, c, d, e, f\}) - v(\succ_4) \\
 v(\succ_3) &= q(a, \{a, b\}) - v(\succ_1) \\
 v(\succ_5) &= q(e, \{a, b, c, e\}) - v(\succ_6) \\
 v(\succ_7) &= q(c, \{a, c\}) - v(\succ_6)
 \end{aligned}
 \tag{7}$$

Note that each line of Equation (7) is written only using variables that have been identified from Equation (6) or higher lines in Equation (7). Equations (6) and (7) tell us we can recover our distribution over preferences v directly from the data for model M , and so M is identified despite not being edge decomposable.

⁸ Further, we can think of the analogue of edge decomposability for graphs other than the probability flow diagram. In this case, edge decomposability still guarantees that we are able to uniquely decompose our graph into a series of flows into our edge decomposable set of paths.

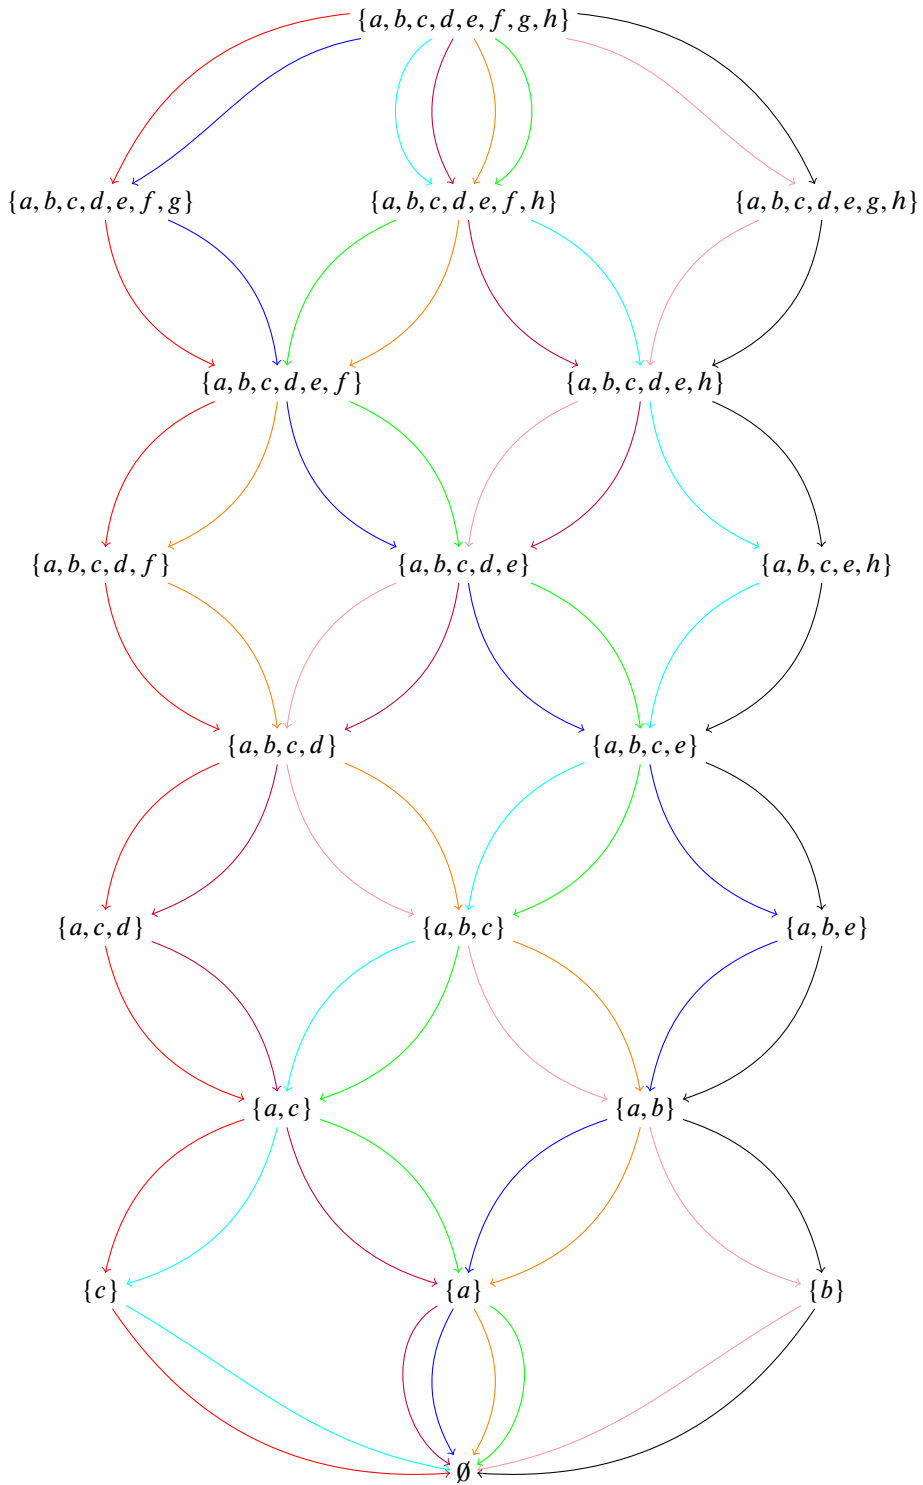


Fig. 2. The graphical representation of Example 4.1. Each differently colored path corresponds to a different preference in our counterexample. (For interpretation of the colors in the figure, the reader is referred to the web version of this article.)

As Example 4.1 demonstrates, edge decomposability is not necessary for identification. However, edge decomposability is a condition that is relatively easy to use. By that we mean that it is easy to constructively create edge decomposable models. Simply start with a preference \succ and consider the model $N = \{\succ\}$. Given N , find a pair (x, A) with $x \in A$ such that $N \cap L(x, A) = \emptyset$. Choose any $\succ' \in L(x, A)$ and create a larger model $N \cup \{\succ'\}$. We can repeat this process until no such pairs (x, A) exist. Once no such pairs exist, we have a maximal edge decomposable model. In fact, as Theorem 4.1 states, maximal edge decomposable models are as large as maximal identified models. However, in order to show this, we need to first introduce two algorithms.

In graph theory, a common question is finding a cycle basis for a given graph. The standard construction is to find a spanning tree for that graph and then construct a single cycle for each edge in the graph which is not in the spanning tree. A spanning tree of a graph is a tree that visits every node of that graph. A tree is a connected undirected graph with no (undirected) cycles. The standard construction develops a cycle basis, but these cycles are not guaranteed to be circuits. Recall that preferences correspond to minimal circuits in the appended probability flow diagram. This means that we are interested in constructing a basis for the circuit space using minimal circuits. As such, we are unable to apply the standard construction right out of the box.

Our first algorithm takes the appended probability flow diagram and constructs a spanning tree which respects the direction of each edge. The fact that each edge in the spanning tree respects the direction of the edges in our original graph is important for our second algorithm. Our second algorithm takes in the spanning tree and creates a cycle basis using minimal circuits. In other words, our second algorithm uses the output of our first algorithm to create our preference basis. The fact that we are able to use minimal circuits in our second algorithm is a result of the specific construction we use in our first algorithm. We now introduce our first algorithm and show that it constructs a spanning tree which respects the direction of edges in the appended probability flow diagram.

Algorithm 4.1.

1. The root node of our spanning tree is chosen as \emptyset .
2. Connect \emptyset with X_n via a directed edge. Set $i = n$.
3. For the sets A with $|A| = i$, enumerate these sets from 1 to J . Set $j = 1$.
4. For each directed (out-)edge, $e(A_j, B)$ leaving A_j , add (the directed edge) $e(A_j, B)$ to our tree if B is not already connected to our tree.
5. Set $j = j + 1$. If $j \leq J$, return to the previous step. If $j > J$, proceed to the next step.
6. Set $i = i - 1$. If $i \leq 0$, terminate the algorithm. If $i \geq 1$, return to step 3 of the algorithm.

Proposition 4.2. *Algorithm 4.1 creates a spanning tree of the probability flow diagram. Further, the edges of the spanning tree respect the direction of the edges in the probability flow diagram.*

Proof. We first prove that this algorithm constructs a spanning tree. Suppose towards a contradiction that we do not. There are three cases.

1. The graph is not spanning.
This means that there is some node A in the probability flow diagram which is not in our tree. This means that every parent node of A in the probability flow diagram is also not in our tree (by step 4 of the algorithm).⁹ Recursively applying this argument eventually takes us to node X_n which is in our tree by construction. This is a contradiction.
2. The graph is not connected.
By the previous case, we know that every node A is connected to some parent node. Repeated application of this logic implies that every node A is connected to X_n via a path which is also connected to \emptyset . Thus the graph is connected.
3. There exists some cycle in the constructed graph.
Since our graph is connected, this means that there are two distinct paths between X_n and some node A . By step 4 of our algorithm, the indegree of A is equal to one.¹⁰ This means that there is a single parent node of A . Call this parent node B . Repeat this argument applied to B and then recursively until we reach X . This means that there is a single path going from A to X_n which goes from child node to parent node.¹¹ Thus if there are two paths from A to X_n , without loss, the second path at some point goes from parent node to child node. However, for every child node, there is also a unique path going from that node to X_n that goes from child to parent nodes. Repeating this argument, we eventually get to some node of the form $\{x\}$. There are no edges on our tree directly connecting $\{x\}$ to \emptyset (by steps 1,2 and 4 in the algorithm), thus there is a unique path from $\{x\}$ to X_n and we are done.

⁹ A parent node of A is any node such that there exists a directed edge from the node to A .
¹⁰ Note that a tree corresponds to the undirected version of the constructed graph. We however included directed edges in our construction as they will be useful in some arguments going forward.
¹¹ A node A is a child node of B if B is a parent node of A .

As none of our three cases can hold, we have a contradiction. Thus our algorithm creates a spanning tree. Further, by the fact that the algorithm initializes going from \emptyset to X_n and then, at each step of the algorithm, starts at A and goes to $A \setminus \{x\}$ for some $x \in A$, the constructed spanning tree respects the direction of edges in the appended probability flow diagram. \square

We now introduce our second algorithm which takes in a spanning tree created from Algorithm 4.1 and creates a preference basis.

Algorithm 4.2.

1. Take as input the spanning (directed) tree from Algorithm 4.1 and the appended probability flow diagram. Initialize with $i = 1$.
2. Enumerate the sets A with $|A| = i$ from 1 to J . Set $j = 1$.
3. Enumerate from 1 to K each directed edge leaving A_j in the probability flow diagram which is not a part of our appended spanning tree. Set $k = 1$. If $K = 0$, proceed to step 6.
4. Add directed edge e_k to the appended spanning tree. By definition of spanning tree, the addition of this edge creates a (new) cycle.

Claim 1. *This cycle can be chosen to be a circuit.*

5. Since this cycle can be chosen to be a circuit, it can be chosen to be a minimal circuit.¹² Minimal circuits in the probability flow diagram correspond to preferences. Add the corresponding preference to the set $\{>_l\}_{l=1}^L$.
6. Set $k = k + 1$. If $k \leq K$, return to step 4. If not, proceed to the next step.
7. Set $j = j + 1$. If $j \leq J$, return to step 3. If not proceed to the next step.
8. Set $i = i + 1$. If $i \leq |X|$, return to step 2. If not, terminate the algorithm.

We now proceed with a proof of Claim 1.

Proof. We proceed by induction on i . Suppose that $i = 1$. In this case, our edge is connecting $\{x\}$ to \emptyset . By construction of our spanning tree in Algorithm 4.1 (step 4), there is a path from $\{x\}$ to X_n in the un-appended spanning tree that goes from children nodes to parent nodes in each step. Since \emptyset is directly connected X_n , by adding the edge connecting $\{x\}$ to \emptyset , we create a circuit.

Now suppose that every edge leaving nodes A where $|A| < i$ has been added to the appended spanning tree. We want to show that adding an edge leaving node B with $|B| = i$ forms some circuit. By a similar argument, there is a directed path from X_n to B in our appended spanning tree. This path goes from parent to children nodes and respects the direction of edges from the probability flow diagram. An edge leaving B connects to a node A with $|A| = i - 1$. Since every edge leaving nodes of size $i - 1$ or less have been added to our appended spanning tree, there exists a path from A to \emptyset which respects the direction of edges from the probability flow diagram. By adding the edge connecting B to A , there is a directed path from X_n to B , from B to A , from A to \emptyset , and from \emptyset to X_n . This is a minimal circuit, and so we are done. \square

Proposition 4.3. *The set of preferences $\{>_l\}_{l=1}^{|M|}$ created by Algorithms 4.1 and 4.2 are (sequentially) edge decomposable. Further, $|M| = (n - 2)2^{n-1} + 2$.*

Proof. Observe that in each iteration of Algorithm 4.2, we create a (minimal) circuit using only edges from the spanning tree, the edge we just added, and edges that were added in prior iterations of the algorithm. This means that the edge we are adding in a specific step of an algorithm is being used for the first time in that step of the algorithm. Each of these edges corresponds to a tuple (x, A) . Let (x_k, A_k) correspond to the edge from the k th iteration of the algorithm. By the prior argument, this means that $\{>_l\}_{l=1}^K \cap L(x_k, A_k) = \{>_k\}$ for $K \leq |M|$. This is exactly the definition of sequential edge decomposability.

Now we show that $|M| = (n - 2)2^{n-1} + 2$. $|M|$ corresponds to the number of edges in the probability flow diagram which are not in the spanning tree. A standard result in graph theory (see Berge (2001) p. 29 Theorem 2) is that the cyclomatic number of a graph is equal to the size of its cycle basis. Further, the standard construction of cycle bases is via constructing a spanning tree and then using each edge not in the spanning tree to create a set of cycles (see again Berge (2001) or more explicitly Bollobás (1998) p. 53 Theorem 9). We have done this in our Algorithms 4.1 and 4.2. We have constructed a set of circuits with one circuit for each edge not in our spanning tree. Thus the number of circuits, and thus preferences in the set $\{>_l\}_{l=1}^{|M|}$, we have created is equal to the cyclomatic number of the probability flow diagram. But this is exactly $(n - 2)2^{n-1} + 2$ by our main theorem, and so we are done. \square

It then follows from the paragraph preceding Theorem 4.1, Example 4.1, and Proposition 4.3 that Theorem 4.1 holds. In effect, Theorem 4.1 tells us that edge decomposability still allows us to form maximal identified models. It simply puts a restriction on which of these models are allowed.

¹² This is because a minimal circuit corresponds to a path through the appended probability flow diagram plus the edge connecting \emptyset with X .

4.1. Relationship to other uniqueness conditions

We now compare edge decomposability to three other uniqueness conditions. Two of these conditions come from the random utility literature and the third comes from Example 2.1. The two conditions from the literature are the uniqueness conditions introduced in Turansick (2022) and Apesteguia et al. (2017). While both of these restrictions recover uniqueness in different ways, they both work by placing restrictions on the preferences in the induced support of their model. Our main goal is to compare these restrictions on supports to our edge decomposability condition on models. We will show that the supports induced by Turansick (2022), Apesteguia et al. (2017), and the model from Example 2.1 are subsets of the possible supports/models induced by edge decomposability.

To begin, we review the setup and characterization of Turansick (2022). The goal of Turansick (2022) is to characterize when the full model $M = \mathcal{L}(X_n)$ admits a unique representation. Unlike our focus, which is developing ex ante conditions that imply identification, the focus of Turansick (2022) is to consider the unrestricted random utility model and characterize which data sets have a unique rationalizing distribution. In the proof of Theorem 1 from Turansick (2022), it is shown that a random choice rule p has a unique representation if and only if every supported path in the probability flow diagram has some edge unique to that path among supported paths. Here, by supported paths, we mean paths where every edge on that path has a strictly positive edge capacity. Restating this in terms of preferences, this means that p has a unique support if and only if one of the following conditions holds for each preference.

1. $\succ \in L(x, A)$ and $q(x, A) = 0$
2. $\exists(x, A)$ with $x \in A$ such that $L(x, A) \cap \{\succ' \mid \succ' \in L(y, B) \implies q(y, B) > 0\} = \{\succ\}$

Relating this back to our definition of edge decomposability, let $M = \{\succ' \mid \succ' \in L(y, B) \implies q(y, B) > 0\}$. Then the result of Turansick (2022) is asking that for every $\succ \in M$, there exists some (x, A) with $x \in A$ such that $M \cap L(x, A) = \{\succ\}$. The condition of Turansick (2022) implies edge decomposability. We now show through an example, related to the counterexample of Fishburn (1998), that it is strictly stronger than edge decomposability.

Example 4.2. Consider the model M given by the following three preferences.

1. $a \succ b \succ c \succ d = \succ_1$
2. $b \succ a \succ d \succ c = \succ_2$
3. $a \succ b \succ d \succ c = \succ_3$

Observe the following.

1. $M \cap L(a, \{a, b, c, d\}) = \{\succ_1, \succ_3\}$
2. $M \cap L(b, \{b, c, d\}) = \{\succ_1, \succ_3\}$
3. $M \cap L(c, \{c, d\}) = \{\succ_2, \succ_3\}$
4. $M \cap L(d, \{d\}) = \{\succ_2, \succ_3\}$

This means that \succ_3 never has a set $L(x, A)$ with $\succ_3 \in L(x, A)$ that is unique to \succ_3 among preferences in M . However, as $M \cap L(b, \{a, b, c, d\}) = \succ_2$ and $M \cap L(d, \{c, d\}) = \{\succ_1\}$, it is easy to see that M is edge decomposable.

We now consider the single crossing random utility model (SCRUM) of Apesteguia et al. (2017). SCRUM puts further structure on X_n in that SCRUM assumes that X_n is endowed with some exogenous linear order \triangleright . We say that a random choice rule p is rationalizable by SCRUM if there exists a distribution over preferences ν such that the support of ν can be ordered so that it satisfies the single-crossing property with respect to \triangleright . Recall from Example 3.2 the single-crossing property.

Definition 4.3. We say that a distribution over preferences ν satisfies the single-crossing property if the support of ν can be ordered in such a way that for all $x \triangleright y$, $x \succ_j y$ implies $x \succ_i y$ for all $j \geq i$.

Unlike previously, it is not immediately obvious that the supports of SCRUM representations are edge decomposable. In order to see that they are, note the following.

Proposition 4.4. Suppose that ν is a distribution over preferences whose support is ordered $(\succ_1, \dots, \succ_n)$ so that it satisfies the single-crossing property with respect to \triangleright . Then there exists a pair (x, A) with $x \in A$ such that $L(x, A) \cap \{\succ_1, \dots, \succ_n\} = \{\succ_1\}$.

It follows from Proposition 4.4 that the supports of SCRUM representations satisfy edge decomposability. Specifically, we know that every subset of a SCRUM support is itself the support of some SCRUM representation. It then follows that the lowest ranked preference in the support is the unique element of some $L(x, A)$ among preferences in the support. In other words, once the support of a SCRUM representation is pinned down, we can recursively find the probability weights on the preferences by looking at the lowest and then the next lowest ranked preferences in the support. We now show through an example that SCRUM supports fail to capture every edge decomposable model.

Example 4.3. Let $X_n = \{a, b, c, d, e, f, g, h\}$ be endowed with the linear order \triangleright and suppose that M is given by the following.

1. $a \triangleright b \triangleright c \triangleright d \triangleright f \triangleright e \Rightarrow \triangleright_1$
2. $a \triangleright b \triangleright d \triangleright c \triangleright e \triangleright f \Rightarrow \triangleright_2$
3. $b \triangleright a \triangleright c \triangleright d \triangleright e \triangleright f \Rightarrow \triangleright_3$

We now proceed with some case work.

1. Suppose $a \triangleright b$. Then it is the case that, abusing notation, if M can satisfy the single-crossing property, $\triangleright_1 \triangleright \triangleright_3$ and $\triangleright_2 \triangleright \triangleright_3$.
 - (a) Suppose $c \triangleright d$. Then it is the case that $\triangleright_1 \triangleright \triangleright_2$ and $\triangleright_3 \triangleright \triangleright_2$. However, given that $a \triangleright b$ and $\triangleright_2 \triangleright \triangleright_3$, $c \triangleright d$ cannot be the case.
 - (b) Suppose $d \triangleright c$. This means that $\triangleright_2 \triangleright \triangleright_1$ which further means that $e \triangleright f$. However, as $a \triangleright b$ and $\triangleright_1 \triangleright \triangleright_3$, we also have that $f \triangleright e$, and so $d \triangleright c$ cannot be the case.
2. Suppose $b \triangleright a$. This means that $\triangleright_3 \triangleright \triangleright_1$ and $\triangleright_3 \triangleright \triangleright_2$. This gives us that $a \triangleright b$, $c \triangleright d$, and $e \triangleright f$. As $c \triangleright_1 d$, $d \triangleright_2 c$, $f \triangleright_1 e$, and $e \triangleright_2 f$, we have that $\triangleright_1 \triangleright \triangleright_2 \triangleright \triangleright_1$, which is a cycle. Thus $b \triangleright a$ cannot be the case.

The above case work tells us that it can neither be the case that $a \triangleright b$ nor $b \triangleright a$, and thus $\{\triangleright_1, \triangleright_2, \triangleright_3\}$ cannot be a SCRUM support. However, as \triangleright_1 is the only preference to rank $f > e$, \triangleright_2 is the only preference to rank $d > c$, and \triangleright_3 is the only preference to rank $b > a$, the model $M = \{\triangleright_1, \triangleright_2, \triangleright_3\}$ is edge decomposable.

We note that while edge decomposability is a strictly weaker condition on supports than single-crossing, part of the novelty of SCRUM is that there is an endogenous mapping from \triangleright and p to the support of v . Alternatively, we could have relied on Theorem 4.1 and Proposition 3.1 to show that edge decomposability is a strictly weaker condition than the single-crossing condition on models. Overall, our discussion in this section shows that edge decomposability is a weaker identifying restriction on supports than two recent identifying restrictions. Further, edge decomposability is an easy to check and constructive criterion for identification.

4.1.1. Preference aggregation and Latin squares

In this section we formalize and discuss the model described in Example 2.1. Importantly, we show that this model ends up being edge decomposable and thus identified. Let \triangleright be a linear order on X_n . Without loss we can have $1 \triangleright 2 \triangleright \dots \triangleright n$.

Definition 4.4. We say that a preference \succ respects \triangleright if there exists $m \in \{1, \dots, n\}$ such that $m \succ m + 1 \succ \dots \succ n - 1 \succ n \succ 1 \succ \dots \succ m - 1$.

The respects relationship is an equivalence relation, whereby any two linear orders inducing the same Latin square are equivalent. Thus, given an order \triangleright , we can think of a full support distribution on preferences which respect \triangleright as the classic example in social choice theory that leads to Condorcet cycles when aggregating preferences.¹³ To extend this further, given an ordering $\triangleright = 1 \triangleright 2 \triangleright \dots \triangleright n$, we can think of appending this ordering to make the cycle $1 \triangleright 2 \triangleright \dots \triangleright n \triangleright 1$. This cycle corresponds to the Condorcet cycle induced when aggregating preferences according to an $n - 1$ majority rule, which respects \triangleright .

Definition 4.5. We say that a model M is a Condorcet Aggregate Random Utility Model (CARUM) if the following conditions hold for some \triangleright :

- Each $\succ \in M$ respects \triangleright .
- Each \succ which respects \triangleright satisfies $\succ \in M$.

For a given \triangleright , denote the corresponding CARUM by $M(\triangleright)$.

The set of CARUM models enjoys a particularly interesting identification property. Assume that the analyst has knowledge that a given random choice rule p is consistent with CARUM, but lacks knowledge about the ordering \triangleright of alternatives. It turns out that, up to orders that induce the same Latin square, the actual \triangleright can be identified. This is the content of the following result.

Proposition 4.5. For each \triangleright , $M(\triangleright)$ is edge decomposable. Further, for any \triangleright and \triangleright' , any random choice rule p , and any $v \in \Delta(M(\triangleright))$ and $v' \in \Delta(M(\triangleright'))$ for which $p = p_v = p_{v'}$, we have $M(\triangleright) = M(\triangleright')$ and $v = v'$.

¹³ We note that there has been recent interest in connecting social choice theory, preference aggregation, and stochastic choice. Brandl et al. (2016) studies aggregation of distributions over preferences and characterizes which aggregation rules satisfy a consistency condition when combining populations and a consistency condition when varying the available menu. Sprumont (2024) studies aggregation via randomized tournaments and gives conditions under which there is a unique extension to a full random choice rule.

The first part of this Proposition 4.5 is easy to see. In a given Latin square $M(\triangleright)$, there is exactly one preference \succ which ranks x first for each $x \in X_n$. Thus $L(x, X_n) \cap M(\triangleright) = \{\succ\}$ for each x and we have edge decomposability. The intuition of the second part of Proposition 4.5 is as follows. Fix a given Latin square $M(\triangleright)$. For each $A \neq X_n$, there is at most a single alternative $x \in A$ such that $q(x, A) > 0$. Consider some edge connecting X_n to $X_n \setminus \{x\}$ in the probability flow diagram. If $q(x, X_n) > 0$, the prior observation implies that there is a unique path from $X_n \setminus \{x\}$ to \emptyset with strictly positive edge capacities along each edge of the path. This allows us to say that the preference \succ which corresponds to this path is in the support of any rationalizing distribution. Since respects is an equivalence relation, we can choose \triangleright to be \succ . From here we simply construct $M(\triangleright)$ and apply edge decomposability to recover identification of v .

5. Related literature

Our paper is primarily related to two strands of literature. The first strand is the one which studies uniqueness and identification in random utility models. Study of the random utility model goes back to Block and Marschak (1959) and Falmagne (1978). Barberá and Pattanaik (1986) and Fishburn (1998) note that the random utility model is in general not identified when there are at least four alternatives available. Much of the literature studying identification aims to recover uniqueness by refining the random utility model. The random expected utility model of Gul and Pesendorfer (2006) recovers uniqueness by restricting to expected utility functions while studying choice over lotteries. As discussed earlier, the single-crossing random utility model of Apesteguia et al. (2017) is able to recover uniqueness by asking that the support of their representation satisfies the single-crossing property with respect to some exogenous order. Yildiz (2023) extends this line of thought and characterizes exactly which models of choice can be identified through a similar progressivity condition. Another approach to identification is through the use of stronger data. Lu (2019) is able to recover both beliefs and preferences when stochastic choice data and information sources are observed. Dardanoni et al. (2020) is able to recover preferences as well as cognitive heterogeneity with data that connects agents’ choices across menus. An alternative approach is taken by Turansick (2022). Turansick (2022) simply asks which realizations of standard random choice rules admit a unique random utility representation. Azrieli and Rehbeck (2022) studies stochastic choice and random utility when we are unable to condition choice on the menu of available alternatives. They find that in this setting that the random utility model is also unidentified. Kashaev et al. (2024) studies separable stochastic choice and finds that separability can be characterized by each agent having well defined marginal choices if and only if the set of feasible choice functions (i.e. preferences) are linearly independent.

The second strand of literature our paper is related to is the strand which uses graph theoretic tools to study the random utility model and more generally stochastic choice. To our knowledge, this strand of literature was started by Fiorini (2004) who used the observation that preferences can be represented as flows on the probability flow diagram in order to provide an alternative proof of the characterization of Falmagne (1978). Turansick (2022) studies conditions on the probability flow diagram which characterize when random choice rules have a unique random utility representation. Chang et al. (2022) studies which preferences have adjacent paths in the probability flow diagram in order to say when random-coefficient models are good approximations of the random utility model. Chambers et al. (2024) extends the probability flow diagram to multiple dimension in order to study choice across multiple people or time periods. Finally, Kono et al. (2023) use graph theoretic tools to study the random utility model when choice probabilities of some alternatives are unobserved.

Declaration of competing interest

Authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

Appendix A. Omitted proofs

A.1. Proof of Lemma 3.1

Proof. If $\{p_\succ\}_{\succ \in M}$ is not linearly independent, then for each $\succ \in M$, there is $c_\succ \in \mathbb{R}$ for which $\sum_{\succ \in M} c_\succ p_\succ = 0$, where at least one $c_\succ \neq 0$. As each $p_\succ \geq 0$ componentwise, it follows that $\{\succ \in M : c_\succ > 0\} \neq \emptyset$ and $\{\succ \in M : c_\succ < 0\} \neq \emptyset$. Observe then that $\sum_{\{\succ \in M : c_\succ > 0\}} c_\succ p_\succ = \sum_{\{\succ \in M : c_\succ < 0\}} -c_\succ p_\succ$. Now, define $v \in \Delta(M)$ as $v(\succ') = \frac{c_{\succ'}}{\sum_{\{\succ \in M : c_\succ > 0\}} c_\succ}$ when $c_{\succ'} > 0$, and $v(\succ') = 0$ otherwise. Similarly define $v' \in \Delta(M)$ as $v(\succ') = \frac{-c_{\succ'}}{\sum_{\{\succ \in M : c_\succ > 0\}} -c_\succ}$ when $c_{\succ'} < 0$ and $v(\succ') = 0$ otherwise. Observe that v and v' have disjoint supports, but by construction $p_v = p_{v'}$, so that M is not identified. If $\{p_\succ\}_{\succ \in M}$ is linearly independent, it is easy to see that M is identified: $p_v = p_{v'}$ implies $\sum_{\succ \in M} (v(\succ) - v'(\succ)) p_\succ = 0$, so that $v(\succ) - v'(\succ) = 0$ for all $\succ \in M$ by the definition of linear independence. \square

A.2. Proof of Lemma 3.2

Proof. Rota (1964), Propositions 1 and 2, show that there is a bijective and linear relationship between a Möbius inverse q and its generating function p . This bijection is through Equation (2). Hence, there is an invertible linear map carrying generating functions p to Möbius inverses q . The result follows as linear independence is preserved under invertible linear transformations. \square

A.3. Proof of Proposition 3.2

Proof. One direction is obvious (the linear independence of $\{x^i\}_{i=1}^I$ entails the linear independence of $\{(x^i, 1)\}_{i=1}^I$). For the other direction, suppose by means of contradiction that $\{(x^i, 1)\}_{i=1}^K$ is linearly independent but that there are c^i not all 0 for which $\sum_i c^i x^i = 0$. Let us suppose without loss that $\sum_k x_k^i = 1$ for all i . Observe then that $\sum_k \sum_i c^i x_k^i = 0$, and that $\sum_k \sum_i c^i x_k^i = \sum_i c^i \sum_k x_k^i = \sum_i c^i$. Consequently $\sum_i c^i = 0$ and hence $\sum_i c^i (x^i, 1) = 0$, contradicting linear independence of $\{(x^i, 1)\}_{i=1}^I$. \square

A.4. Proof of Proposition 3.1

We first consider the number of binary comparisons given a preference over X_n . A binary comparison simply takes two alternatives $x, y \in X_n$ and ascribes either $x > y$ or $y > x$. As such, each preference has $\binom{n}{2}$ binary comparisons. By the definition of single-crossing, if $x > y$ and $x >_i y$ then $x >_j y$ for all $j \geq i$. Let \triangleright^- denote the preference given by $x > y \implies y \triangleright^- x$. Thus the largest a set of single-crossing preferences could possibly be is the size of following construction. Let \triangleright^- be the first preference in our (ordered) set and let \triangleright be the last preference in our ordered set. Then given preference $>_i$, we simply change one binary comparison that agrees with \triangleright^- to one which agrees with \triangleright . In terms of size, we start with a single preference and then iteratively change the ordering of each binary comparison. This corresponds to $\binom{n}{2} + 1$ preferences.

We have just constructed an upper bound for the size of a set of single-crossing preferences. We now give an algorithm that reaches that upper bound. Enumerate each alternative x_i, x_j such that $i > j \implies i \triangleright j$.

1. Initialize at $i = k = 1$ and add $\triangleright^- \Rightarrow_k$ to the set \mathcal{M} .
2. Set $j = i + 1$.
3. Take $>_k$ and construct $>_{k+1}$ by swapping the ranking of x_i and x_j . Add $>_k$ to \mathcal{M} . Set $k = k + 1$.
4. Set $j = j + 1$. If $j > |X|$, set $i = i + 1$ and proceed to the next step. If not, return to the prior step.
5. If $i = |X|$, terminate the algorithm. If not, return to step 2.

Our algorithm effectively takes \triangleright^- as input, takes the lowest ranked alternative and iteratively moves it up one rank at each iteration until it is ranked in the same location as it is in \triangleright . Since at each step of the algorithm, we are only changing the binary comparison of two alternatives which are ranked next to each other, in each step of the algorithm, we construct a preference. We start with \triangleright^- and then make $\sum_{i=1}^n \sum_{j=i+1}^n 1$ swaps of binary comparisons. This corresponds to $\binom{n}{2} + 1$ preferences. Thus we reach our upper bound from the first half of the proof and we are done.

A.5. Proof of Proposition 4.4

Proof. If ν has a single preference in its support, we are done. Suppose otherwise that $n \geq 2$, so that $>_1$ and $>_2$ are in the support of ν . Since $>_1$ and $>_2$ differ, this means that there exists some pair (x, y) such that $x >_1 y$ and $y >_2 x$. Since ν satisfies the single crossing property, for all $i \geq 2$, we have that $y >_i x$. Let A denote the set such that $>_i \in L(x, A)$. Since $x >_1 y, y \in A$ and thus $L(x, A) \cap \{>_1, \dots, >_n\} = \{>_1\}$. The result now continues by induction. \square

A.6. Proof of Proposition 4.5

Proof. The proof of edge decomposability of $M(\triangleright)$ is discussed following Proposition 4.5. Now we consider the rest of the proposition. Suppose that there exists some (\triangleright, ν) with $\nu \in M(\triangleright)$ such that random choice rule p satisfies $p = p_\nu$.

Claim 2. For each $\emptyset \neq A \neq X_n$, there is at most one $x \in A$ such that $q(x, A) > 0$.

Proof. Consider two preferences $>$ and $>'$ with $> \in L(x, A)$ and $>' \in L(y, A)$ for $x \neq y$ and $A \neq X_n$. Without loss, we can take $>$ to be \triangleright . In a Latin square, since $A \neq \emptyset$, there is a single preference which ranks $X \setminus A$ as the first $|X \setminus A|$ alternatives. In our case, this corresponds to $\triangleright \Rightarrow$. This means that $>' \notin M(\triangleright)$ and so we are done. \square

There is some ν such that $p = p_\nu$. This means that each $q(x, A) \geq 0$ (by Falmagne (1978)) and that there is some $x \in X_n$ such that $q(x, X_n) > 0$. Now consider $X_n \setminus \{x\}$. By Falmagne (1978), we have $\sum_{x \in A} q(x, A) = \sum_{y \in X_n \setminus A} q(y, A \cup \{y\})$ for each $A \neq X_n$. This means that there is some $y \in X_n \setminus \{x\}$ such that $q(y, X_n \setminus \{x\}) > 0$. By Claim 2, this y is unique. We can repeat this process, repeatedly going from A to $A \setminus \{z\}$ where z satisfies $q(z, A) > 0$. This process induces the unique path from X_n to \emptyset that passes through $X_n \setminus \{x\}$ and satisfies $q(y, A) > 0$ along every edge of the path. This means that the preference corresponding to this path is in the support of ν . Call this preference $>$. Since respects is an equivalence relation, any preference in the Latin square of $>$, the set of preferences which respect $>$, can be chosen to be \triangleright . In this case, we can choose $>$ as \triangleright , apply edge decomposability, and we are done. \square

Data availability

No data was used for the research described in the article.

References

- Apestequia, J., Ballester, M.A., Lu, J., 2017. Single-crossing random utility models. *Econometrica* 85, 661–674.
- Azrieli, Y., Rehbeck, J., 2022. Marginal stochastic choice. arXiv preprint. arXiv:2208.08492.
- Barberá, S., Pattanaik, P.K., 1986. Falmagne and the rationalizability of stochastic choices in terms of random orderings. *Econometrica* 54, 707–715.
- Berge, C., 2001. *The Theory of Graphs*. Dover Publications.
- Block, H.D., Marschak, J., 1959. *Random Orderings and Stochastic Theories of Response*. Tech. Rep. Cowles Foundation for Research in Economics, Yale University.
- Bollobás, B., 1998. *Modern Graph Theory*, vol. 184. Springer Science & Business Media.
- Brandl, F., Brandt, F., Seedig, H.G., 2016. Consistent probabilistic social choice. *Econometrica* 84, 1839–1880.
- Chambers, C.P., Masatlioglu, Y., Turansick, C., 2024. Correlated choice. *Theor. Econ.* 19, 1087–1117.
- Chang, H., Narita, Y., Saito, K., 2022. Approximating choice data by discrete choice models. arXiv preprint. arXiv:2205.01882.
- Dardanoni, V., Manzini, P., Mariotti, M., Tyson, C.J., 2020. Inferring cognitive heterogeneity from aggregate choices. *Econometrica* 88, 1269–1296.
- Dogan, S., Yildiz, K., 2022. Every choice function is pro-con rationalizable. *Oper. Res.*
- Falmagne, J.-C., 1978. A representation theorem for finite random scale systems. *J. Math. Psychol.* 18, 52–72.
- Feller, W., 1968. *An Introduction to Probability Theory and Its Applications*, vol. 1. Wiley.
- Fiorini, S., 2004. A short proof of a theorem of falmagne. *J. Math. Psychol.* 48, 80–82.
- Fishburn, P.C., 1998. Stochastic utility. In: Barbera, S., Hammond, P., Seidl, C. (Eds.), *Handbook of Utility Theory*. Kluwer, Dordrecht, pp. 273–318.
- Gul, F., Pesendorfer, W., 2006. Random expected utility. *Econometrica* 74, 121–146.
- Kashaev, N., Plávala, M., Aguiar, V.H., 2024. Entangled vs. Separable choice. arXiv preprint. arXiv:2403.09045.
- Kono, H., Saito, K., Sandroni, A., 2023. Axiomatization of random utility model with unobservable alternatives. arXiv preprint. arXiv:2302.03913.
- Lu, J., 2019. Bayesian identification: a theory for state-dependent utilities. *Am. Econ. Rev.* 109, 3192–3228.
- Luce, R.D., 1959. *Individual Choice Behavior*. John Wiley.
- Rota, G.-C., 1964. On the foundations of combinatorial theory I. theory of Möbius functions. *Z. Wahrscheinlichkeitstheor. Verw. Geb.* 2, 340–368.
- Saito, K., 2018. *Axiomatization of the Mixed Logit Model*.
- Sprumont, Y., 2024. Randomized collective choices based on a fractional tournament. *Theor. Econ.*
- Turansick, C., 2022. Identification in the random utility model. *J. Econ. Theory* 203, 105489.
- Yildiz, K., 2023. Foundations of self-progressive choice models.