Economics Working Paper Series
Faculty of Business, Economics and Law, AUT

Cycle conditions for "Luce rationality"

José A. Rodrigues-Neto, Matthew Ryan and James Taylor

2024/03

Cycle conditions for "Luce rationality"

José A. Rodrigues-Neto, Matthew Ryan and James Taylor*

March 26, 2024

Abstract

We extend and refine conditions for "Luce rationality" (i.e., the existence of a Luce – or logit – model) in the context of stochastic choice. When choice probabilities satisfy *positivity*, we show that the *cyclical independence (CI)* condition of Ahumada and Ülkü (2018) and Echenique and Saito (2019) is necessary and sufficient for Luce rationality, even if choice is only observed for a restricted set of menus. We then adapt results from the *cycles approach* (Rodrigues-Neto, 2009) to the common prior problem (Harsanyi, 1967-1968) to refine the CI condition, by reducing the number of *cycle equations* that need to be checked. A general algorithm is provided to identify a minimal sufficient set of equations (depending on the collection of menus for which choice is observed). Three cases are discussed in detail: (i) when choice is only observed from binary menus, (ii) when all menus contain a common default; and (iii) when all menus contain an element from a common binary default set. Investigation of case (i) leads to a refinement of the famous *product rule*.

*Rodrigues-Neto and Taylor: *Research School of Economics, Australian National University;* Ryan: *Department of Economics and Finance, Auckland University of Technology.*

1 Introduction

The classical Luce – or logit – model (Luce, 1959) assumes *positivity* of choice probabilities: each available option is chosen with strictly positive probability. Echenique and Saito (2019) define a *general Luce model (GLM)* that is compatible with options being chosen with zero probability, and prove that it is characterised by a *cyclical independence* condition. The same result was obtained, independently, by Ahumada and Ülkü (2018).[1] In a GLM there is a strictly-positive-valued utility function that determines, in Luce fashion, the allocation of probability within the (finite) support of the choice probability function on each menu: a supported alternative x is chosen with probability equal to the utility of x as a proportion of the total utility in the support (i.e., the sum of utilities of supported alternatives).

Cyclical independence is a restriction on the probability of choice cycles. A choice cycle is a sequence of choices where each successively chosen element is also available in the immediately preceding menu and the final chosen element is identical to the first; it is a cycle in "revealed (weak) preference". Cyclical independence requires (roughly speaking) that any choice cycle has the same probability as the reverse cycle (with which it is paired) when each cycle in the pair may occur with positive probability. This requirement generates a *cycle equation* associated with the cycle pair.

Ahumada and Ülkü (2018) and Echenique and Saito (2019) show that existence of a GLM is equivalent to cyclical independence under the maintained assumption that choice probabilities are defined for **all** non-empty menus that may be constructed from a given **finite** universal set of alternatives, X. We show (Theorem 3) that the result remains valid without any restriction on X, or on the set of menus, other than the requirement that all menus are finite (i.e., in the set-up of Cerreia-Vioglio *et. al.*, 2021). This added generality is useful for practical applications. Choice is usually only observed from a limited set of menus.[2] Moreover, for the same general setting, we show that cyclical independence, *together with positivity*, characterises the Luce model (Corollary 1).

Cyclical independence is therefore the fundamental empirical signature of *(general) Luce rationality* – the compatibility of choice probabilities with the existence of a (general) Luce model – for an *arbitrary* collection of finite menus.

Section 4 connects our Corollary 1 with results from a related literature. When the

[1] Ahumada and Ülkü use the term *Luce rule with limited consideration* rather than GLM, and provide a different interpretation of the model. They do not name the cyclical independence condition but it appears as their Axiom 1.

[2] See McCausland *et al.* (2020) for a rare exception.

universal set of alternatives is finite, we show that Corollary 1 is implied by Proposition 1 in Rodrigues-Neto (2009), though expressed in a very different idiom. The latter result establishes necessary and sufficient conditions for the existence of a *common prior* in the framework of Harsanyi (1967-1968). In that idiom, X is a state space and menus are conditioning events; choice probability functions become conditional probability functions.

It is well known that the existence of a common prior is mathematically analogous to the existence of a Luce model, at least for finite X. It should therefore come as little surprise that an analogue of cyclical independence also makes an appearance in the common prior literature. Following Rodrigues-Neto (2009), several papers have explored the so-called *cycles approach* to the common prior problem.[3] We import ideas from this literature to refine the cyclical independence condition for Luce rationality. By "refine" we mean identify a subset of cycle equations whose satisfaction guarantees cyclical independence (assuming positivity): a "basis" for the set of cycle equations.

The amount of refinement that can be achieved depends on the structure of the menu set. Two such results are already familiar (at least for finite X). When the menu set comprises all binary menus, it suffices to satisfy the cycle equations for cycles of length three (Luce and Suppes, 1965, Theorem 48): this is equivalent to the well-known *product rule*.[4] When all possible menus are present, it suffices to check cycles of length two (Luce, 1959, Lemmas 2-3 and Theorem 3): this is nothing other than the famous *independence of irrelevant alternatives (IIA)* condition.

The central result of Rodrigues-Neto (2012) may be used to establish the cardinality of a basis for the cycle equations for any given menu set (see Theorem 7). We use this result to further refine the product rule for binary menus (Theorem 8), and to show that a subset of the IIA conditions suffice to ensure cyclical independence when all menus share a common default (Theorems 5 and 9). We also consider menu sets where every menu contains at least one element from a *pair* of defaults. In this case, we show that it suffices to check a subset of the cycles of length four (Lemma 2 and Theorem 6). Finally, we describe an algorithm from graph theory to identify a basis of cycle equations for any given menu set. In on-going work, we hope to develop these tools into an empirical strategy to test for Luce rationality with an arbitrary collection of menus.

[3] See, for example, Rodrigues-Neto (2012, 2014), Hellwig (2013), Hellman and Samet (2012) and Fiorini and Rodrigues-Neto (2017).

[4] The product rule condition appears, unnamed, in Luce (1959, p.16). The *product rule* terminology seems to have been coined by Estes (1960, p.272).

2 The (general) Luce model

We adopt the set-up of Cerreia-Vioglio *et. al.* (2021). Let X be a non-empty set, which is interpreted as the universal domain of alternatives. Let $\mathcal{M}$ be a non-empty collection of non-empty, **finite** subsets of X. An element of $\mathcal{M}$ is a menu from which a single alternative must be chosen. Note that neither X nor $\mathcal{M}$ need be finite.[5]

A *random choice function (RCF)* describes the stochastic choice behavior of some individual. An RCF is a function $p : X \times \mathcal{M} \to [0, 1]$ satisfying $\sum_{x \in E} p(x, E) = 1$ for any $E \in \mathcal{M}$ and $p(x, E) = 0$ for any $E \in \mathcal{M}$ and any $x \in X \setminus E$. We interpret $p(x, E)$ as the probability that the individual chooses x when confronted with menu E.

Unless otherwise stated, we assume that $\mathcal{M}$ includes all singletons – the definition of a random choice function fixes its value on any singleton, so this assumption is without loss of generality. If $\mathcal{M}$ contains all non-empty, finite subsets of X then we say that the menu set is *comprehensive*.

If p is an RCF we define $\Gamma_p : \mathcal{M} \to 2^X$ to be the support function for p:

$$\Gamma_p(E) \;=\; \{x \in E \mid p(x, E) > 0\}$$

for each $E \in \mathcal{M}$. Note that Γ_p satisfies the properties of a choice function: $\emptyset \neq \Gamma_p(E) \subseteq E$ for each $E \in \mathcal{M}$.

Next, we recall some properties of RCFs and a classical result:

Definition 1 *An RCF satisfies* ***positivity*** *if* $p(x, E) > 0$ *when* $x \in E \in \mathcal{M}$.

Definition 2 *An RCF satisfies* ***independence of irrelevant alternatives (IIA)*** *if*

$$p(x, E)\, p(y, F) = p(x, F)\, p(y, E)$$

whenever $E, F \in \mathcal{M}$ *and* $\{x, y\} \subseteq E \cap F$.

If p satisfies positivity, the IIA condition may be expressed in a more familiar form:

$$\frac{p(x, E)}{p(y, E)} \;=\; \frac{p(x, F)}{p(y, F)}$$

(i.e., the relative likelihood of choosing x over y is independent of the menu in which they appear). We adopt the less familiar product form so we can consider contexts where positivity may not be satisfied.[6]

[5] However, we restrict attention to the case of finite X from Section 4 onwards.

[6] Cerreia-Vioglio *et al.* (2021, Lemma 6) study the relationships amongst several variants of the IIA condition and Luce's (1959) *choice axiom* when positivity is relaxed.

Definition 3 *If $p : X \times \mathcal{M} \to [0,1]$ is an RCF, then p has a **Luce model** if there exists some (utility) function $v : X \to \mathbb{R}_{++}$ such that*

$$p(x, E) = \frac{v(x)}{\sum_{y \in E} v(y)}$$

whenever $x \in E \in \mathcal{M}$. We say that v is a Luce model for p.

Note that if v is a Luce model for p then so is $u = kv$ for any $k > 0$. In particular, if p possesses a Luce model and X is finite, then p possesses a Luce model such that $\sum_{x \in X} v(x) = 1$.

Theorem 1 (cf., Luce, 1959) *Let $p : X \times \mathcal{M} \to [0,1]$ be an RCF. Suppose that (i) $\mathcal{M}$ is comprehensive, and (ii) p satisfies positivity. Then p has a Luce model if and only if it satisfies IIA.*

Luce (1959) proves this result for finite X. For general X, the standard proof requires some modification. We prove Theorem 1 in Section 3.

This classical result is a foundation stone for the large literature on the Luce model. However, assumption (i) is a significant constraint for empirical work, while (ii) has proved troubling to theorists.[7]

Positivity is obviously a necessary condition for an RCF to possess a Luce model. Ahumada and Ülkü (2018) and Echenique and Saito (2019) relax the positivity assumption and consider RCFs which possess a model of the following more general form:[8]

Definition 4 *Let $p : X \times \mathcal{M} \to [0,1]$ be an RCF. Then p has a **general Luce model (GLM)** if there exists a (utility) function $v : X \to \mathbb{R}_{++}$ such that*

$$p(x, E) = \frac{v(x)}{\sum_{y \in \Gamma_p(E)} v(y)}$$

whenever $E \in \mathcal{M}$ and $x \in \Gamma_p(E)$.

This is a kind of 'bounded awareness' generalisation of the Luce model. The set $\Gamma_p(E)$ is the collection of alternatives in E which the decision maker is aware of. The model does not suggest that those alternatives not in $\Gamma_p(E)$ are strictly worse in some way than those in $\Gamma_p(E)$, merely that they have not been properly considered.

[7]Empiricists have been less troubled. As McFadden (1973, p.109) famously observed: "Since empirically a zero probability is indistinguishable from one that is extremely small, there is little loss of generality in assuming that the selection probabilities are all positive".

[8]As noted above, we adopt the terminology of Echenique and Saito (2019).

Ahumada and Ülkü (2018) and Echenique and Saito (2019) characterise the GLM when X is finite and $\mathcal{M}$ is comprehensive. They show that an RCF has a general Luce model (under these maintained assumptions) iff it satisfies a condition that generalises IIA. To state this generalised condition we require some additional notation and terminology, which are mostly adopted from Echenique and Saito (2019). A sequence $\{(x_i, x_{i+1}, E_i)\}_{i=1}^m$ with $m \in \{1, 2, ...\}$ and $\{x_i, x_{i+1}\} \subseteq E_i \in \mathcal{M}$ for each i (and repetition allowed)[9] is called a *connected sequence* of *length* m. It is a *cycle* if $x_1 = x_{m+1}$. A cycle of length m will be called an m-cycle. A connected sequence $\{(x_i, x_{i+1}, E_i)\}_{i=1}^m$ will be sometimes be abbreviated as

$$x_1 E_1 x_2 E_2 \cdots E_{m-1} x_m E_m x_{m+1}$$

when notationally convenient.

A connected sequence $\{(x_i, x_{i+1}, E_i)\}_{i=1}^m$ is *positive* (for a given RCF, p) if

$$p(x_i, E_i) p(x_{i+1}, E_i) > 0$$

for each i. We then say that there is a positive connected sequence from x_1 to x_{m+1} and denote this fact by $x_1 \rightarrow_p x_{m+1}$. Of course, all connected sequences are positive when p satisfies positivity. In general, $\rightarrow_p$ is a reflexive,[10] symmetric and transitive binary relation on X. Hence, the equivalence classes of $\rightarrow_p$ partition X. The symmetry of $\rightarrow_p$ reflects the fact that positive connected sequences come in natural pairs. If $\{(x_i, x_{i+1}, E_i)\}_{i=1}^m$ is a positive connected sequence (hence $x_1 \rightarrow_p x_{m+1}$), then so is the "reverse" sequence

$$\{(x_{m+2-i}, x_{m+1-i}, E_{m+1-i})\}_{i=1}^m$$

(which shows that $x_{m+1} \rightarrow_p x_1$).

Definition 5 *An RCF, $p : X \times \mathcal{M} \rightarrow [0, 1]$, satisfies* ***cyclical independence (CI)*** *if*

$$\prod_{i=1}^m p(x_i, E_i) = \prod_{i=1}^m p(x_{i+1}, E_i) \qquad (*)$$

for any positive cycle $\{(x_i, x_{i+1}, E_i)\}_{i=1}^m$.

We refer to $(*)$ as the *cycle equation* for the associated cycle (pair).[11] We say that a cycle is *consistent* if it satisfies $(*)$. The IIA condition requires that every 2-cycle is

[9] A connected sequence is an indexed family, not a set.
[10] Recall that $\mathcal{M}$ includes all singletons.
[11] The cycle equation for the reverse cycle is identical.

consistent. For RCFs satisfying positivity, CI therefore generalises – and strengthens – the IIA condition.

We next state a generalised version of the result in Ahumada and Ülkü (2018) and Echenique and Saito (2019); one that relaxes the assumption of finite X. It is a special case of an even more general result presented later: see Theorem 3 below.

Theorem 2 (cf., Ahumada and Ülkü, 2018; Echenique and Saito, 2019) *Let $\mathcal{M}$ be comprehensive and let $p : X \times \mathcal{M} \to [0,1]$ be an RCF. Then p possesses a general Luce model iff it satisfies CI.*

3 Arbitrary menu sets

Theorems 1 and 2 maintain the restrictive assumption of a comprehensive menu set. It turns out that this assumption is entirely redundant for Theorem 2, and may be dispensed with in Theorem 1 by replacing IIA with the stronger condition, CI.

The proof of the following result follows essentially the same lines as the proofs of Ahumada and Ülkü (2018, Theorem 1) and Echenique and Saito (2019, Theorem 1), but is included for completeness.

Theorem 3 *Suppose $p : X \times \mathcal{M} \to [0,1]$ is an RCF. Then p possesses a general Luce model if and only if it satisfies CI.*

Proof. It is straightforward to verify that if p has a GLM then it satisfies CI.

Conversely, suppose p satisfies CI. We first observe that CI implies the following condition:[12] for any $x, \hat{x} \in X$ with $x \to_p \hat{x}$, and any two positive connected sequences $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ and $\left\{\left(x'_j, x'_{j+1}, E'_j\right)\right\}_{j=1}^{m'}$ from x to $\hat{x}$, we have

$$\prod_{i=1}^{m} \frac{p\left(x_i, E_i\right)}{p\left(x_{i+1}, E_i\right)} = \prod_{j=1}^{m'} \frac{p\left(x'_j, E'_j\right)}{p\left(x'_{j+1}, E'_j\right)} \tag{1}$$

To see why, suppose that $x \to_p x$ via the positive connected sequences $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ and $\left\{\left(x'_j, x'_{j+1}, E'_j\right)\right\}_{j=1}^{m'}$. Concatenating the former with the reverse of the latter gives a positive cycle $(x \to_p \hat{x} \to_p x)$, so CI implies

$$\prod_{i=1}^{m} \frac{p\left(x_i, E_i\right)}{p\left(x_{i+1}, E_i\right)} \prod_{j=1}^{m'} \frac{p\left(x'_{j+1}, E'_j\right)}{p\left(x'_j, E'_j\right)} = 1 \quad \Leftrightarrow \quad \prod_{i=1}^{m} \frac{p\left(x_i, E_i\right)}{p\left(x_{i+1}, E_i\right)} = \prod_{j=1}^{m'} \frac{p\left(x'_j, E'_j\right)}{p\left(x'_{j+1}, E'_j\right)}$$

which gives (1).

[12] In fact, it is easy to see that the two conditions are equivalent.

For each equivalence class, A, of $\to_p$ we assign strictly positive utilities to the elements of A as follows. Fix some $\hat{x} \in A$ and set $v(\hat{x}) = 1$. For each $x \in A \diagdown \{\hat{x}\}$ we have $x \to_p \hat{x}$ so define

$$v(x) = \prod_{i=1}^{m} \frac{p(x_i, E_i)}{p(x_{i+1}, E_i)}$$

for any positive connected path $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ from x to $\hat{x}$. Property (1) ensures that this definition is unambiguous. Moreover, we have $v(x) > 0$; and v is defined on all of X since the equivalence classes of $\to_p$ partition X.

It remains to show that v generates a GLM for p. Suppose $E \in \mathcal{M}$ and $x, y \in \Gamma_p(E)$. It is evident that $\Gamma_p(E)$ is contained in a single equivalence class of $\to_p$ since any two alternatives in $\Gamma_p(E)$ are linked by a positive connected sequence of length 1. Let $\hat{x}$ be the fixed element used to define utilities in this equivalence class and consider the positive cycle $\hat{x} \to_p x \to_p y \to_p \hat{x}$. By concatenating positive connected sequences that make up this cycle and applying CI we have:

$$\frac{1}{v(x)} \frac{p(x, E)}{p(y, E)} v(y) = 1 \quad \Leftrightarrow \quad v(x) = \frac{p(x, E)}{p(y, E)} v(y).$$

Hence:

$$\sum_{x \in \Gamma_p(E)} v(x) = \frac{v(y)}{p(y, E)} \quad \Leftrightarrow \quad p(y, E) = \frac{v(y)}{\sum_{x \in \Gamma_p(E)} v(x)}.$$

This completes the proof. □

When p satisfies positivity, v is a GLM for p iff it is a LM for p. We immediately deduce:

Corollary 1 *Suppose $p : X \times \mathcal{M} \to [0, 1]$ is an RCF that satisfies positivity. Then p has a LM iff it satisfies CI.*

We may now prove Theorem 1.

Proof of Theorem 1. The "only if" part is straightforward. To prove the "if" part we first show that under the conditions of the theorem, IIA implies CI. Let $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ be a positive cycle and define

$$E = \bigcup_{i=1}^{m} E_i.$$

Note that E is finite and $E \in \mathcal{M}$ since $\mathcal{M}$ is comprehensive. Now observe that IIA and positivity imply

$$\frac{p(x_i, E_i)}{p(x_{i+1}, E_i)} = \frac{p(x_i, E)}{p(x_{i+1}, E)}$$

for each i, and CI follows. Hence p has a GLM (Theorem 3). Any GLM is an LM when p satisfies positivity. □

In summary, when $\mathcal{M}$ is an **arbitrary** collection of **finite** menus, CI characterises *general Luce rationality*, and CI together with positivity characterises *Luce rationality*.

4 Luce models and common priors

The special case of Corollary 1 for finite X has a precursor in Proposition 1 of Rodrigues-Neto (2009), which concerns the existence of a common prior for the posteriors of a set of differentially informed agents. In the latter context, X is a state space, menus are reinterpreted as conditioning events, and choice probabilities become posterior beliefs.

More precisely, there is a set $J = \{1, 2, \ldots, n\}$ of agents. The information of agent j is described by a knowledge partition Π^j of (the finite set) X. For each $x \in X$, we use $\pi^j(x)$ to denote the element of Π^j containing x. The posterior beliefs of agent j are described by a function $p^j : X \times \Pi^j \to [0, 1]$ that satisfies $p^j(x, E) = 0$ if $x \notin E$ and $\sum_{x \in X} p^j(x, E) = 1$ for each $E \in \Pi^j$. Here, $p^j(x, \pi^j(x))$ represents the posterior probability that agent j places on state x after learning event $\pi^j(x) \in \Pi^j$. Let us call $(p^1, p^2, \ldots, p^n)$ a *posterior belief system (PBS)*. Rodrigues-Neto (2009) restricts attention to posterior belief systems that satisfy the analogue of positivity: $p^j(x, E) > 0$ whenever $x \in E \in \Pi^j$. We therefore take this to be part of the definition of a PBS.

In this context, a *common prior* for $(p^1, p^2, \ldots, p^n)$ is a function $\mu : X \to (0, 1]$ such that:

$$p^j(x, \pi^j(x)) = \frac{\mu(x)}{\sum_{y \in \pi^j(x)} \mu(y)}$$

for all $j \in J$ and $x \in X$. We say that $(p^1, p^2, \ldots, p^n)$, is *consistent* if $p^j(\cdot, E) \equiv p^\ell(\cdot, E)$ whenever $E \in \Pi^j \cap \Pi^\ell$. Consistency is obviously a necessary condition for the existence of a common prior.

Here is the result of Rodrigues-Neto (2009):

Theorem 4 (Rodrigues-Neto, 2009, Proposition 1) *Let X be a finite state space and let $\mathcal{P} = (p^1, p^j, \ldots, p^n)$ be a PBS. Then $\mathcal{P}$ has a common prior iff*

$$\prod_{i=1}^{m} p^{j_i}(x_i, E_i) = \prod_{i=1}^{m} p^{j_i}(x_{i+1}, E_i) \qquad (**)$$

for any sequences $(j_1, j_2, \ldots, j_m)$, $(x_1, x_2, \ldots, x_m, x_{m+1})$ and $(E_1, E_2, \ldots, E_m)$ such that $x_1 = x_{m+1}$ and $\{x_i, x_{i+1}\} \subseteq E_i = \pi^{j_i}(x_i)$ for each $i \in \{1, 2, \ldots, m\}$.

Note that condition $(**)$ enforces consistency: if

$$x \neq y \in E = \pi^j(x) = \pi^\ell(x)$$

then $(**)$ implies $p^j(x, E)p^\ell(y, E) = p^j(y, E)p^\ell(x, E)$ and hence

$$\frac{p^j(x, E)}{p^j(y, E)} = \frac{p^\ell(x, E)}{p^\ell(y, E)}.$$

Given the finiteness of E, it follows that $p^j(\cdot, E) \equiv p^\ell(\cdot, E)$.

To connect Theorem 4 to our Corollary 1 we make the following observations.

1. Given any consistent PBS, $\mathcal{P} = (p^1, p^2, \ldots, p^n)$, we may set $\mathcal{M} = \bigcup_{j\in J} \Pi^j$ and define an RCF, $p : X \times \mathcal{M} \to [0, 1]$ as follows: $p(x, \pi^j(x)) = p^j(x, \pi^j(x))$. This function is well-defined by consistency and inherits positivity from $\mathcal{P}$. Conversely, if p is a random choice function satisfying positivity, with $\mathcal{M}$ equal to a union of partitions of X, then there is a consistent PBS that induces $\mathcal{M}$ and p in the manner just described. Moreover, recalling the remark immediately following Definition 3, $\mathcal{P}$ has a common prior iff p has a Luce model.

2. If $\mathcal{P} = (p^1, p^2, \ldots, p^n)$ is a consistent PBS, and $\mathcal{M}$ and p are constructed as above, then the condition in Theorem 4 may be expressed as follows:

$$\prod_{i=1}^{m} p(x_i, E_i) = \prod_{i=1}^{m} p(x_{i+1}, E_i)$$

   for any sequences $(x_1, x_2, \ldots, x_m, x_{m+1})$ and $(E_1, E_2, \ldots, E_m)$ such that $x_1 = x_{m+1}$ and $\{x_i, x_{i+1}\} \subseteq E_i$ for each $i \in \{1, 2, \ldots, m\}$. Since p satisfies positivity, this is equivalent to imposing the CI condition on p.

3. Given any menu set, $\mathcal{M}$, that includes all singletons, we may express $\mathcal{M}$ as a union of partitions of X: for each $E \in \mathcal{M}$ define the partition

$$\Pi^E = \{E\} \cup \{\{x\} \mid x \in X \setminus E\}$$

   so that $\mathcal{M} = \bigcup_{E\in\mathcal{M}} \Pi^E$.

Combining these observations we see that Corollary 1 is implied by Theorem 4, under a suitable translation between frameworks. The results are not quite equivalent in that condition $(**)$ does additional work to impose consistency on the PBS; work that is redundant in the (single agent) stochastic choice environment.

The process of translation provides a bridge from the cycles approach to the common prior existence problem on the one side, to the characterisation of Luce models for arbitrary menu sets on the other. Any result on consistent posterior belief systems can be directly translated into to an analogous result in the stochastic choice context (and *vice versa*).

In the next section we exploit developments from the common prior literature to refine the CI condition for Luce rationality, based on the structure of $\mathcal{M}$. Since the common prior literature on which we rely assumes a finite state space, *we assume X is finite for the remainder of the paper.*

5 Independent cycle equations

Depending on the structure of $\mathcal{M}$, it may not be necessary to check all positive cycles to verify that CI holds. There may exist a subset of positive cycles whose associated cycle equations imply the cycle equations of all other positive cycles. We say that the excluded cycles are "spanned" by the subset.[13]

5.1 Bounding the length of cycles

Note that if every (positive) m-cycle is consistent, then so is every (positive) k-cycle for any $k < m$: a (positive) k-cycle $\{(x_i, x_{i+1}, E_i)\}_{i=1}^k$ can be extended to a (positive) m-cycle by appending $m-k$ terms of the form (x_{k+1}, x_{k+1}, E_k), and the extended cycle generates the same cycle equation as the original k-cycle.

In some situations, it may suffice – for the purposes of verifying the CI property – to check the cycle equations for positive cycles of a given length. When p satisfies positivity – the Corollary 1 scenario – there are two familiar situations of this sort. If $\mathcal{M}$ is comprehensive, we know that CI holds iff all 2-cycles are consistent (Theorem 1). If $\mathcal{M}$ comprises all menus of cardinality up to two – the *binary menus* case – then CI holds iff the *product rule* is satisfied (Luce and Suppes, 1965, Theorem 48); the product rule requires that all 3-cycles are consistent:

$$p(x,y)\,p(z,x)\,p(y,z) \;=\; p(y,x)\,p(x,z)\,p(z,y) \tag{PR}$$

for any $\{x, y, z\} \subseteq X$. Indeed, we may leverage the latter result into a stronger one. If $\mathcal{M}$ *includes* all binary menus, then it is straightforward to show that a LM exists (hence

[13]The linear algebra terminology will be motivated shortly.

CI is satisfied) iff p satisfies the product rule and IIA.[14] Note that the product rule only makes reference to binary menus, while the IIA condition applies to all menus. It follows that CI holds iff all 3-cycles are consistent.

For the common prior context (Theorem 4), Hellwig (2013) exhibits a restriction on conditioning events (knowledge partitions) that implies a bound on the length of cycles for which the CI-analogue condition (∗∗) must be checked. The required restriction is that $\pi^j(x) \cap \pi^\ell(x') \neq \emptyset$ for all $x, x' \in X$ and all $j \neq \ell$. Under this restriction, Hellwig shows that a common prior exists iff all 4-cycles satisfy (∗∗).[15] Using the analysis of Section 4, we may translate Hellwig's restriction on knowledge partitions into a corresponding restriction on $\mathcal{M}$. A salient special case of the implied restriction is when every non-singleton menu in $\mathcal{M}$ includes a common "default" option. For this special case, we obtain an even stronger result: it suffices to check cycles of length two.

Theorem 5 *Suppose there is some fixed alternative $d \in X$ such that $d \in E$ for all non-singleton $E \in \mathcal{M}$. Let $p : X \times \mathcal{M} \to [0,1]$ be an RCF that satisfies positivity. Then p satisfies CI if and only if it satisfies IIA.*

Recall that IIA is equivalent to all 2-cycles being consistent, which is equivalent to all positive 2-cycles being consistent when p satisfies positivity. Before proving Theorem 5, we first establish a useful Lemma.

Lemma 1 *Let $p : X \times \mathcal{M} \to [0,1]$ be an RCF that satisfies positivity. Let*

$$c = \{(x_i, x_{i+1}, E_i)\}_{i=1}^m$$

be a cycle with $E_1 \cap \cdots \cap E_m \neq \emptyset$. If p satisfies IIA, then the cycle equation for c is satisfied.

Proof. First, suppose $x_1 \in E_1 \cap \cdots \cap E_m$. Recall that $x_1 = x_{m+1}$. For ease of exposition, let $y = x_1 = x_{m+1}$. In this case, we must show that

$$\frac{p(x_2, E_2)}{p(x_2, E_1)} \frac{p(x_3, E_3)}{p(x_3, E_2)} \cdots \frac{p(x_m, E_m)}{p(x_m, E_{m-1})} \frac{p(y, E_1)}{p(y, E_m)} = 1$$

or, equivalently,

$$\left[\frac{p(x_2, E_2)\, p(y, E_1)}{p(x_2, E_1)\, p(y, E_2)}\right] \left[\frac{p(x_3, E_3)\, p(y, E_2)}{p(x_3, E_2)\, p(y, E_3)}\right] \cdots \left[\frac{p(x_m, E_m)\, p(y, E_{m-1})}{p(x_m, E_{m-1})\, p(y, E_m)}\right] = 1.$$

[14] Kovach and Tserenjigmid (2022, Theorem 1) extend this idea to characterise a generalisation of the LM which they call the *focal Luce model*.

[15] Recall that positivity of p is an implicit assumption in the common prior context.

But this equation holds since each square-bracketed term is equal to 1 by IIA.

Next, suppose $m > 1$ and $x_k \in E_1 \cap \cdots \cap E_m$ for some $k \in \{2, 3, \ldots, m\}$. Let $y = x_k$. From the previous case, we know that the cycles

$$\{(y, x_1, E_1), (x_1, x_2, E_1) \ldots, (x_{k-1}, y, E_{k-1})\}$$

and

$$\{(y, x_{k+1}, E_k), (x_{k+1}, x_{k+2}, E_{k+1}), \ldots, (x_m, x_{m+1}, E_m)(x_1, y, E_1)\}$$

are consistent. The cycle equations for these two cycles are, respectively:

$$\frac{p(y, E_1)}{p(x_1, E_1)} \prod_{i=1}^{k-1} \frac{p(x_i, E_i)}{p(x_{i+1}, E_i)} = 1$$

and

$$\frac{p(x_1, E_1)}{p(y, E_1)} \prod_{i=k}^{m} \frac{p(x_i, E_i)}{p(x_{i+1}, E_i)} = 1.$$

Multiplying these together gives $(*)$.

Finally, suppose $d_c \in E_1 \cap \cdots \cap E_m$ and $d_c \notin \{x_1, \ldots, x_m\}$. Since $x_1 = x_{m+1}$ we must show that

$$\frac{p(x_2, E_2)}{p(x_2, E_1)} \frac{p(x_3, E_3)}{p(x_3, E_2)} \cdots \frac{p(x_m, E_m)}{p(x_m, E_{m-1})} \frac{p(x_1, E_1)}{p(x_1, E_m)} = 1$$

which is equivalent to

$$\left[\frac{p(x_2, E_2)\, p(d_c, E_1)}{p(x_2, E_1)\, p(d_c, E_2)}\right] \left[\frac{p(x_3, E_3)\, p(d_c, E_2)}{p(x_3, E_2)\, p(d_c, E_3)}\right] \cdots \left[\frac{p(x_1, E_m)\, p(d_c, E_m)}{p(x_1, E_m)\, p(d_c, E_1)}\right] = 1.$$

This equation holds since each square-bracketed term is equal to 1 by IIA. Hence, the cycle c is consistent. □

Proof of Theorem 5. Since p satisfies positivity, all cycles are positive. The "only if" part is therefore immediate. We show the "if" part.

Suppose p satisfies IIA, and there exists a default option d such that $d \in E$ for all menus $E \in \mathcal{M}$. For any cycle $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$, we have $d \in E_1 \cap \cdots E_m$. Therefore, by Lemma 1, cycle c is consistent. □

Combining this result with Corollary 1 we have:

Corollary 2 *Suppose there is some fixed alternative $d \in X$ such that $d \in E$ for all non-singleton $E \in \mathcal{M}$. Let $p : X \times \mathcal{M} \to [0, 1]$ be an RCF that satisfies positivity. Then p has a LM if and only if it satisfies IIA.*

We can perform a similar exercise for the case with two potential default options. We say that $\mathcal{M}$ contains two potential default options if there exist $a, b \in X$ such that for all non-singleton $E \in \mathcal{M}$, either $a \in E$ or $b \in E$.

Lemma 2 *Let $p : X \times \mathcal{M} \to [0,1]$ be an RCF which satisfies positivity, and let $\mathcal{M}$ contain two potential default options, $a, b \in X$. Then p has a LM if and only if all cycles in the following categories are consistent: all 2-cycles and 3-cycles, and all 4-cycles of the form*

$$xE_1aE_2yE_3bE_4x \tag{♣}$$

Proof. The "only if" part is immediate from Corollary 1. We prove the "if" part.

We begin by defining some useful notation. If $q = \{(x_i, x_i', E_i)\}_{i=1}^m$ is a sequence with $\{x_i, x_i'\} \subseteq E_i \in \mathcal{M}$ for all $i \in \{1, 2, \ldots, m\}$ (not necessarily a connected sequence) we call the elements of the sequence *edges* and we define

$$v(q) = \prod_{i=1}^{m} \frac{p(x_i', E_i)}{p(x_i, E_i)}.$$

If q_1 and q_2 are two such sequences, let $q_1 \circ q_2$ denote their concatenation. Observe that $v(q_1 \circ q_2) = v(q_1)v(q_2)$.

Consider a cycle $c = x_1E_1x_2E_2 \cdots E_{m-1}x_mE_mx_1$. Since p satisfies positivity, c is a positive cycle so we must show that $v(q) = 1$ (i.e., c is consistent). We start by decomposing c as follows:

$$c = q_1 \circ q_2 \circ \cdots \circ q_k$$

where the connected subsequences $\{q_j\}_{j=1}^k$ are associated with an alternating sequence of defaults: there exist $d_j \in \{a, b\}$ for each $j \in \{1, 2, \ldots, k\}$ such that each E_i appearing in q_j contains d_j, and $d_j \neq d_{j+1}$ for each $j \in \{1, 2, \ldots, k-1\}$. Of course, many such decompositions may be possible; just fix one. Let $\mathcal{J}_a \subseteq \{1, 2, ..., k\}$ be defined by $j \in \mathcal{J}_a$ iff $d_j = a$, and let $\mathcal{J}_b$ be the complementary set of indices.

Define $E^j \in \{E_1, \ldots, E_m\}$ to be the first menu appearing in q_j and $F^j \in \{E_1, \ldots, E_m\}$ to be the last. Likewise, let $y^j \in \{x_1, x_2, \ldots, x_m\}$ be the first, and $z^j \in \{x_1, x_2, \ldots, x_m\}$ the last, alternative appearing in q_j. Hence $q_j = y^jE^j \cdots F^jz^j$. For each $j \in \{1, 2, \ldots, k\}$ define edges $e_j^1 = (z^j, d_j, F^j)$ and $e_j^2 = (d_j, y^j, E^j)$, and associated connected sequence $\pi_j = e_j^1 \circ e_j^2 = z^jF^jd_jE^jy^j$. Finally, let $\hat{q}_j = q_j \circ \pi_j$ and note that $\hat{q}_j$ is a positive cycle. This cycle follows q_j then loops back to the start of q_j by going through d_j (recall that each menu in q_j contains d_j). Since all 2-cycles are consistent, and so p satisfies IIA, Lemma 1 implies that $v(\hat{q}_j) = 1$ for each j. Hence

$$v(\hat{q}_1 \circ \hat{q}_2 \circ \cdots \circ \hat{q}_k) = 1 \tag{2}$$

We will use this fact to show that $v(c) = 1$.

The value of $v(\hat{q}_1 \circ \hat{q}_2 \circ \cdots \circ \hat{q}_k)$ is unaffected by re-arranging the edges in $\hat{q}_1 \circ \hat{q}_2 \circ \cdots \circ \hat{q}_k$. For each $j \in \mathcal{J}_a$ define a new connected sequence

$$\rho_j = \pi_j \circ e^1_{j \ominus 1} \circ e^2_{j \oplus 1}$$

where

$$j \ominus 1 = \begin{cases} j-1 & \text{if } j > 1 \\ k & \text{if } j = 1 \end{cases} \quad \text{and,} \quad j \oplus 1 = \begin{cases} j+1 & \text{if } j < k \\ 1 & \text{if } j = k \end{cases}$$

Since $y^j = z^{j \ominus 1}$ and $z^j = y^{j \oplus 1}$, and $j \in \mathcal{J}_a$ we see that

$$\rho_j = z^j F^j a E^j y^j F^{j \ominus 1} b E^{j \oplus 1} z^j$$

is a cycle of the form (♣). Hence $v(\rho_j) = 1$ for each $j \in \mathcal{J}_a$. Moreover, the sequence $\rho_1 \circ \rho_2 \circ \cdots \circ \rho_k$ is obtained from $\pi_1 \circ \pi_2 \circ \cdots \circ \pi_k$ by rearranging edges. We therefore have:

$$\begin{aligned} v(\hat{q}_1 \circ \hat{q}_2 \circ \cdots \circ \hat{q}_k) &= v(q_1 \circ \cdots \circ q_k) \prod_{j \in \mathcal{J}_a} v(\rho_j) \\ &= v(c) \end{aligned}$$

so $v(c) = 1$ from (2). □

Theorem 6 *Let $p : X \times \mathcal{M} \to [0,1]$ be an RCF that satisfies positivity and let $\mathcal{M}$ contain two potential default options, $a, b \in X$. Then:*

(i) p has a Luce model if and only if all 4-cycles are consistent.

(ii) If there exists a menu containing both default options, then p has a Luce model if and only if all 3-cycles are consistent.

Proof. (i) If p has a Luce model then p satisfies Corollary 1 so all cycles (hence all 4-cycles) are consistent. The converse holds by virtue of Lemma 2.

(ii) Let $F \in \mathcal{M}$ be such that $a, b \in F$.

The "only if" part follows from Corollary 1. Conversely, suppose all 3-cycles are consistent. To show that p has a Luce model we need to show that all cycles of the form (♣) are also consistent. To that end, let $c = xE_1aE_2yE_3bE_4x$. Now consider the following 3-cycles: xE_1aFbE_4x and aE_2yE_3bFa. Both are consistent so we have:

$$p(x,E_1)p(a,F)p(b,E_4) = p(a,E_1)p(b,F)p(x,E_4)$$
$$p(a,E_2)p(y,E_3)p(b,F) = p(y,E_2)p(b,E_3)p(a,F).$$

Multiplying, cancelling like terms, and re-arranging gives:

$$p(x,E_1)p(a,E_2)p(y,E_3)p(b,E_4) = p(a,E_1)p(y,E_2)p(b,E_3)p(x,E_4).$$

Hence c is consistent. □

5.2 Bounding the number of cycles

As Lemma 2 makes clear, there is some residual redundancy in Theorem 6(i). We do not need to check all 4-cycles in order to verify CI. There is similar redundancy in the product rule for the binary menus case, and in the Theorem 5 result. The latter redundancy is verified by Theorem 9 and Example 3 below; the following example shows that the set of 3-cycle equations is surplus to requirements in the binary menus case.

Example 1 *Let $X = \{a,b,c,d\}$ and let $\mathcal{M}$ consist of all non-empty subsets of X with cardinality no greater than two. Suppose p satisfies positivity and (PR) holds for*

$$(x,y,z) \in \{(a,b,c), (a,b,d), (b,c,d)\}.$$

Then (PR) also holds when $(x,y,z) = (a,c,d)$:

$$\begin{aligned}
\frac{p(a,c)}{p(c,a)}\frac{p(d,a)}{p(a,d)}\frac{p(c,d)}{p(d,c)} &= \frac{p(a,c)}{p(c,a)}\frac{p(d,a)}{p(a,d)}\left(\frac{p(c,b)}{p(b,c)}\frac{p(b,d)}{p(d,b)}\right) \\
&= \left(\frac{p(a,c)}{p(c,a)}\frac{p(c,b)}{p(b,c)}\right)\frac{p(d,a)}{p(a,d)}\frac{p(b,d)}{p(d,b)} \\
&= \frac{p(a,b)}{p(b,a)}\frac{p(d,a)}{p(a,d)}\frac{p(b,d)}{p(d,b)} \\
&= 1
\end{aligned}$$

where the first equality uses (PR) for $(x,y,z) = (b,c,d)$; the third uses (PR) for $(x,y,z) = (a,b,c)$ and the final equality uses (PR) for $(x,y,z) = (a,b,d)$.

We will shortly identify an "independent" set of cycle equations that suffice for CI in the binary menus case (Theorem 8).

In the context of the common prior problem, Rodrigues-Neto (2012, Corollary 2) uses graph-theoretic techniques to identify an upper bound on the number of cycle equations that need to be checked in order to verify the CI-analogue condition ($**$) in Theorem 4. This bound depends only on the structure of the knowledge partitions; it uses no information on posterior probabilities. Using the bridge constructed in Section 4, this result is readily translated into the stochastic choice context, where the bound now depends on the structure of the menu set, $\mathcal{M}$.

First, we define a suitable *multigraph* from $(X, \mathcal{M})$.[16] The vertex set is X and the edge set is $\{[xy; E] \,|\, \{x, y\} \subseteq E \in \mathcal{M}\}$. The edge $[xy; E]$ joins vertices x and y.[17] Hence, there are as many edges joining x to y as there are menus containing x and y. We say that edge $[xy; E]$ joins x to y within E. Let $G(X, \mathcal{M})$ denote this (undirected) multigraph. Each connected sequence determines a *walk* in the multigraph $G(X, \mathcal{M})$, and *vice versa*: the connected sequence $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ is associated with the walk:[18]

$$(x_1, [x_1x_2; E_1], x_2, [x_2x_3; E_2], x_3, ..., x_m, [x_mx_{m+1}; E_m], x_{m+1}).$$

A *simple walk* is a walk with no repeated edges; and a simple walk connecting two distinct vertices is a *path* if no vertex is encountered more than once along the walk. If each edge in a walk (or path) is within $E \in \mathcal{M}$ then we say that the walk (or path) itself is within E. A *connected component* of $G(X, \mathcal{M})$ is a set of vertices with the property that every distinct pair of vertices in the set is connected by a path, but there is no path from any element of the set to any vertex outside the set.

Rodrigues-Neto (2012) defines a *version* of $G(X, \mathcal{M})$ to be any subgraph obtained by deleting just enough edges to satisfy the following requirement: for every $E \in \mathcal{M}$ and every distinct $x, y \in E$ there is a unique path from x to y within E. Note that any version of $G(X, \mathcal{M})$ has the same connected components as $G(X, \mathcal{M})$ itself. The result of Rodrigues-Neto (2012), adapted to our setting, says the following: to ensure CI it suffices to check a number of cycle equations equal to the *cyclomatic number* (see

[16] See Appendix A for a review of the basic concepts from graph theory needed here.

[17] The expression "$[xy; E]$" is synonymous with "$[yx; E]$" – the multigraph is undirected. In terms of the notation in Appendix A.1, we may think of $[xy; E]$ as the *index* associated with the edge:

$$\mathcal{I} = \{[xy; E] \,|\, \{x, y\} \subseteq E \in \mathcal{M}\}.$$

[18] Provided, that is, the connected sequence contains no element of the form (x, x, E), since multigraphs, as we define them in Appendix A, exclude loops. For the purposes of verifying consistency, this restriction is innocuous.

Appendix B) of any version of $G(X, \mathcal{M})$. The cyclomatic number is the same for all versions and is equal to:

$$\left(\sum_{A\in\mathcal{M}} |A|\right) - |\mathcal{M}| - |X| + \kappa$$

where κ is the number of connected components in $G(X, \mathcal{M})$. Let us denote this quantity by $C(X, \mathcal{M})$. We therefore have:

Theorem 7 (Cf., Rodrigues-Neto, 2012) *Let $p : X \times \mathcal{M} \to [0, 1]$ be a random choice function satisfying positivity. There exists a set of cycle equations, with cardinality $C(X, \mathcal{M})$, such that p has a Luce model iff these cycle equations are satisfied.*

For completeness, we give a proof of Theorem 7 in Appendix B. It is a direct adaptation of the proof of Corollary 2 in Rodrigues-Neto (2012).

The following example computes $C(X, \mathcal{M})$ for the case of binary menus.

Example 2 *Suppose $|X| = n$ and $\mathcal{M}$ consists of all binary menus.*[19] *The cyclomatic number for this case is:*

$$\begin{aligned}\left(\sum_{A\in\mathcal{M}} |A|\right) - |\mathcal{M}| - |X| + \kappa &= 2|\mathcal{M}| - |\mathcal{M}| - |X| + \kappa \\ &= |\mathcal{M}| - n + 1 \\ &= \frac{n(n-1)}{2} - (n-1) \\ &= \frac{(n-1)(n-2)}{2}\end{aligned}$$

Example 2 reinforces our earlier observation about redundancy in the product rule conditions. The product rule rule specifies one equation for every three-element subset of X. The number of such subsets is

$$\frac{n(n-1)(n-2)}{6}$$

which exceeds $C(X, \mathcal{M})$ when $n > 3$. In particular:

$$\frac{(n-1)(n-2)}{2} = \frac{n(n-1)(n-2)}{6} - \frac{(n-1)(n-2)(n-3)}{6}$$

[19] The presence of singletons is inconvenient for present purposes so we exclude them from $\mathcal{M}$. The reader will easily verify that including them leads to the same answer by a slightly longer route.

so the product rule contains

$$\frac{(n-1)(n-2)(n-3)}{6}$$

redundant equations. This represents a fraction $(n-3)/n$ of the total. For example, 50% of the product rule equations are redundant when $n = 6$.

It is not difficult to identify a subset of the product rule conditions that achieves the bound in Example 2 and suffices for CI. The following may be proved using the algebra in Example 1.[20]

Theorem 8 *Let $X = \{x_1, x_2, \ldots, x_n\}$ and let $\mathcal{M}$ be the set of all binary subsets of X. A random choice function satisfying positivity has a Luce model iff it satisfies (PR) for all $\{x, y, z\} \subseteq X$ with $|\{x, y, z\}| = 3$ and $x_1 \in \{x, y, z\}$*

The CI condition is therefore guaranteed provided all instances of the product rule that include x_1 are satisfied. Of course, x_1 may be fixed arbitrarily. The number of three-element subsets of X that include x_1 is

$$\frac{(n-1)(n-2)}{2}$$

which is equal to the value of $C(X, \mathcal{M})$ computed in Example 2.

It is also interesting to revisit Theorem 5 and its corollary (Corollary 2) in the light of Theorem 7. When all menus contain a common default, Theorem 5 tells us that it suffices to check the cycle equations for 2-cycles. However, the proof of Theorem 5 makes it clear that we need only check a subset of these: just those 2-cycles which have the default option in the "middle" position and which are not contained within a single menu. The associated cycle equations embody a very intuitive condition: for any non-default option that is available in more than one menu, the likelihood of this option being chosen, relative to the likelihood of the default being chosen, should be the same across all menus in which the option appears. This condition is clearly necessary for the existence of a Luce model, and it takes but a little thought to convince oneself that it is also sufficient. This subset of the 2-cycles contains one cycle for every common element (other than d) of every pair of distinct menus in $\mathcal{M}$, so its cardinality is equal to that of the set

$$\{(x, \{E, F\}) \mid E, F \in \mathcal{M},\ x \in (E \cap F) \setminus \{d\} \text{ and } E \neq F\} \tag{$\spadesuit$}$$

Nevertheless, Theorem 7 implies that even this reduced set of 2-cycles is excessive to the purpose of verifying CI. Unless $|\mathcal{M}| < 3$, the cardinality of ($\spadesuit$) may exceed $C(X, \mathcal{M})$.

[20] Simply identify b with x_1 and $\{a, c, d\}$ with any three-element subset of X that excludes x_1.

Theorem 9 *Suppose there is some $d \in X$ such that $d \in E$ for all $E \in \mathcal{M}$.*[21] *Then the cardinality of (♠) weakly exceeds $C(X, \mathcal{M})$. They are equal if $|\mathcal{M}| < 3$.*

Proof. Suppose $\mathcal{M} = \{E_1, E_2, \ldots, E_m\}$. Let C_m denote the cardinality of (♠). Note that C_m is equal to:

$$\sum_{1 \leq i < j \leq m} |E_i \cap E_j| \; - \; \frac{m(m-1)}{2} \qquad (C_m)$$

In the present scenario:

$$\kappa \; = \; 1 \; + \; \left| X \setminus \bigcup_{i=1}^{m} E_i \right|$$

so $C(X, \mathcal{M})$ simplifies to

$$\sum_{i=1}^{m} |E_i| - \left| \bigcup_{i=1}^{m} E_i \right| - (m-1) \qquad (N_m)$$

If $m = 1$ then $N_1 \; = \; C_1 \; = \; 0$.

If $m = 2$ then $N_2 \; = \; C_2 \; = \; |E_1 \cap E_2| - 1$.

We next show that $N_m \; \leq C_m$ when $m > 2$. Note that $N_m \; \leq C_m$ iff

$$\left| \bigcup_{i=1}^{m} E_i \right| \; - \; \sum_{i=1}^{m} |E_i| \; + \; \sum_{1 \leq i < j \leq m} |E_i \cap E_j| \; \geq \; \frac{(m-1)(m-2)}{2} \qquad (\#)$$

To prove (#) we first define $X_k \subseteq X$ to be the set of alternatives that appear in exactly k menus in $\mathcal{M}$. Thus $d \in X_m$, $X_k \cap X_{k'} = \emptyset$ whenever $k \neq k'$, and $\bigcup_{k=1}^{m} X_k = X$ (though some X_k may be empty). We may therefore re-write the left-hand side of (#) as follows:

$$\sum_{k=1}^{m} \left[\left| \bigcup_{i=1}^{m} (E_i \cap X_k) \right| \; - \; \sum_{i=1}^{m} |E_i \cap X_k| \; + \; \sum_{1 \leq i < j \leq m} |E_i \cap E_j \cap X_k| \right]$$

The square-bracketed term is zero if $k \in \{1, 2\}$ and strictly positive if $k > 2$ (and $X_k \neq \emptyset$), since:

$$\left| \bigcup_{i=1}^{m} (E_i \cap X_k) \right| \; = \; |X_k|$$

$$\sum_{i=1}^{m} |E_i \cap X_k| \; = \; k \, |X_k|$$

and

$$\sum_{1 \leq i < j \leq m} |E_i \cap E_j \cap X_k| \; \begin{cases} = & 0 & \text{if } k = 1 \\ = & |X_k| & \text{if } k = 2 \\ > & (k-1)\,|X_k| & \text{if } k > 2 \end{cases}$$

[21] Here it is convenient, but inessential, to exclude singleton menus (other than $\{d\}$) from $\mathcal{M}$.

Now, for each $k < m$, create k (hypothetical) "replicas" of each alternative in X_k – one for each menu in which the alternative appears – and treat these replicas as distinct alternatives. Do the same for all elements of X_m except for the default, d, which remains unreplicated. Let X^* denote the expanded universe of alternatives.[22] For each $E_i \in \mathcal{M}$ define $E_i^* \subseteq X^*$ by relabelling each non-default alternative to the corresponding replica,[23] and define the corresponding menu set $\mathcal{M}^* = \{E_1^*, ..., E_m^*\}$. Therefore, d is common to every menu in $\mathcal{M}^*$ but every other alternative in X^* appears in a unique menu. Finally, define X_k^* for each $k \in \{1, ..., m\}$ to be the alternatives in X^* that appear in exactly k menus in $\mathcal{M}^*$. Hence, $|X_m^*| = 1$ and $|X_k^*| = 0$ if $1 < k < m$.

We now observe that the left-hand side of (#) weakly exceeds

$$\left|\bigcup_{i=1}^{m}(E_i \cap X_m)\right| - \sum_{i=1}^{m}|E_i \cap X_m| + \sum_{1\le i<j\le m}|E_i \cap E_j \cap X_m|$$

which in turn weakly exceeds

$$\begin{aligned}
&\left|\bigcup_{i=1}^{m}(E_i^* \cap X_m^*)\right| - \sum_{i=1}^{m}|E_i^* \cap X_m^*| + \sum_{1\le i<j\le m}\left|E_i^* \cap E_j^* \cap X_m^*\right| \\
&= \sum_{k=1}^{m}\left[\left|\bigcup_{i=1}^{m}(E_i^* \cap X_k^*)\right| - \sum_{i=1}^{m}|E_i^* \cap X_k^*| + \sum_{1\le i<j\le m}\left|E_i^* \cap E_j^* \cap X_k^*\right|\right] \\
&= \left|\bigcup_{i=1}^{m}E_i^*\right| - \sum_{i=1}^{m}|E_i^*| + \sum_{1\le i<j\le m}\left|E_i^* \cap E_j^*\right|
\end{aligned}$$

By the inclusion-exclusion formula

$$\left|\bigcup_{i=1}^{m}E_i^*\right| - \sum_{i=1}^{m}|E_i^*| + \sum_{1\le i<j\le m}\left|E_i^* \cap E_j^*\right| = \sum_{\substack{J\subseteq\{1,2,\dots,m\}\\|J|\ge3}}(-1)^{|J|+1}\left|\bigcap_{j\in J}E_j^*\right|.$$

Since $\bigcap_{j\in J}E_j^* = \{d\}$ for all $J \subseteq \{1, 2, \dots, m\}$ with $|J| \ge 3$ we have:

$$\begin{aligned}
\sum_{\substack{J\subseteq\{1,2,\dots,m\}\\|J|\ge3}}(-1)^{|J|+1}\left|\bigcap_{j\in J}E_j^*\right| &= (-1)\sum_{\substack{J\subseteq\{1,2,\dots,m\}\\|J|\ge3}}(-1)^{|J|} \\
&= (-1)\left[\sum_{J\subseteq\{1,2,\dots,m\}}(-1)^{|J|} - 1 + m - \frac{m(m-1)}{2}\right] \\
&= \frac{(m-1)(m-2)}{2} + (-1)\sum_{J\subseteq\{1,2,\dots,m\}}(-1)^{|J|} \\
&= \frac{(m-1)(m-2)}{2}
\end{aligned}$$

[22] Thus, $|X^*| = (\sum_{k=1}^{m} k\,|X_k|) - (m-1)$.
[23] Thus, $|E_i^*| = |E_i|$.

where the final equality uses Lemma 2.1 of Shafer (1976). We therefore deduce the required inequality (#). □

It is not difficult to see why there may be redundancy in (♠) when $|\mathcal{M}| \geq 3$. As the proof of Theorem 9 makes clear, there are redundant elements in (♠) whenever there exists some $x \in X \setminus \{d\}$ that is contained in at least three distinct menus. The following example illustrates.

Example 3 *Suppose* $X = \{x_1, x_2, x_3, y, d\}$ *and* $\mathcal{M} = \{E_1, E_2, E_3\}$ *with* $E_i = \{x_i, y, d\}$. *Let* p *be an RCF satisfying positivity and define*

$$p_i(a, b) = \frac{p(a, E_i)}{p(b, E_i)}$$

for each $i \in \{1, 2, 3\}$ *and each* $\{a, b\} \subseteq E_i$. *The set (♠) is*

$$\{(y, E_1, E_2),\ (y, E_2, E_3),\ (y, E_1, E_3)\}.$$

and element (y, E_i, E_j) *has associated cycle equation*

$$p_i(y, d) = p_j(y, d).$$

If $p_1(y, d) = p_2(y, d)$ *and* $p_2(y, d) = p_3(y, d)$ *then* $p_1(y, d) = p_3(y, d)$, *so one of the three cycle equations is redundant.*

5.3 Finding a "cycle basis"

To apply Theorem 7, we require a systematic way of identifying a set of $C(X, \mathcal{M})$ cycles such that CI holds iff the cycle equations of these $C(X, \mathcal{M})$ cycles are satisfied. As explained in Appendix B, this amounts to finding a basis for a particular vector subspace. Appendix C describes a simple *graph-theoretic* algorithm for finding such a basis. The algorithm fixes some version, $G^*(X, \mathcal{M})$, of $G(X, \mathcal{M})$, then proceeds as follows:

Initialisation: Let $G_1 = G^*(X, \mathcal{M})$ and $k = 1$

Step k Do while G_k has a cycle:

- **Step k-1** Choose a cycle of G_k and denote it by C_k
- **Step k-2** Choose an edge in C_k and denote it by e_k
- **Step k-3** Let G_{k+1} be the subgraph $G_k - e_k$
- **Step k-4** $k + 1 \leftarrow k$

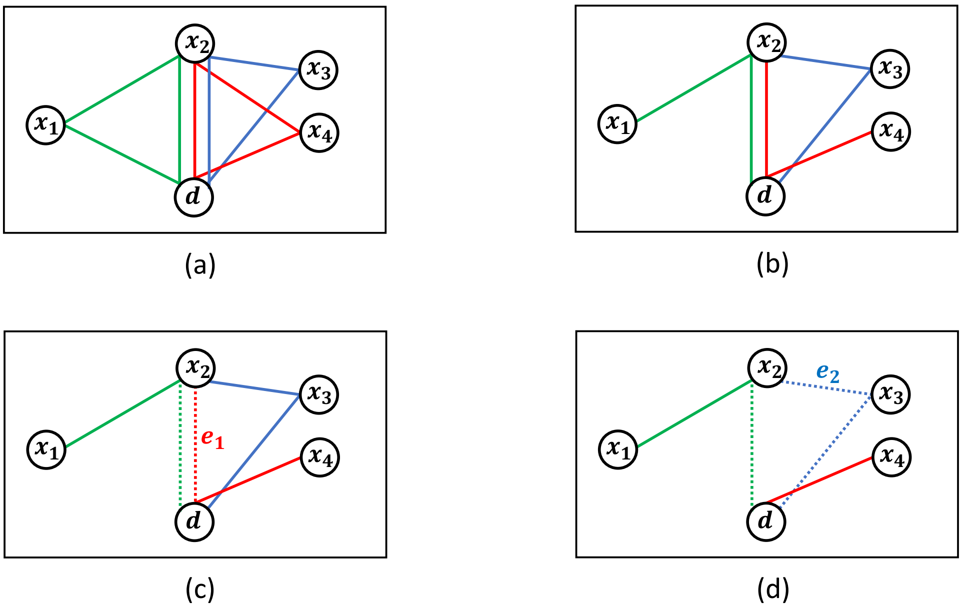

Figure 1: Graphs of Example 4.

This algorithm will terminate in finitely many steps, say K of them. It can be shown that $K = C(X, \mathcal{M})$. The set $\{C_1, C_2, ..., C_K\}$ of cycles is called a *cycle basis* for $G^*(X, \mathcal{M})$. Fixing any cycle basis for any version of $G(X, \mathcal{M})$, satisfaction of the cycle equations associated with the cycles in this cycle basis will suffice for CI. See see Appendix C for details.

We illustrate the algorithm with the following example.

Example 4 *Let* $X = \{d, x_1, x_2, x_3, x_4\}$ *and*

$$\mathcal{M} \;=\; \{\{d, x_1, x_2\}, \{d, x_2, x_3\}, \{d, x_2, x_4\}\}.$$

Note that every menu contains the "default" option, d*, and*

$$C(X, \mathcal{M}) \;=\; 9 - 3 - 5 + 1 \;=\; 2 \;>\; |(\spadesuit)| \;=\; 3.$$

Figure 1(a) depicts the multigraph, $G(X, \mathcal{M})$*, and Figure 1(b) depicts one version of* $G(X, \mathcal{M})$*. Edges of the same colour belong to the same menu.*[24] *Figure 1(c) indicates a possible choice of* C_1 *(dashed) and* e_1*. Figure 1(d) exhibits a subsequent possible choice of* C_2 *and* e_2*. The algorithm terminates at this point, so* $\{C_1, C_2\}$ *is a cycle basis.*

[24] The colours only matter for identifying versions. They are superfluous to the rest of the algorithm.

5.4 Independent cycles and uniqueness of Luce models

There is a close link between Theorem 7 and the uniqueness properties of Luce models. Fixing $(X, \mathcal{M})$, note that any (graph-theoretic) cycle in the muligraph, $G(X, \mathcal{M})$, must be contained within a connected component. These connected components determine, in the obvious fashion, partitions $\{X_k\}_{k=1}^{K}$ and $\{\mathcal{M}_k\}_{k=1}^{K}$ of X and $\mathcal{M}$ respectively. If $p : X \times \mathcal{M} \to [0,1]$ is an RCF, let p_k denote the restriction of p to $X_k \times \mathcal{M}_k$. It is clear that p is uniquely determined by $\{p_k\}_{k=1}^{K}$. It is equally clear that p has a Luce model iff there is a Luce model for each p_k. The following result establishes that a Luce model for p_k is unique up to multiplication by a strictly positive scalar.

Theorem 10 *Let $p_k : X_k \times \mathcal{M}_k \to [0,1]$ be an RCF satisfying positivity. Suppose there is a connected sequence joining any two distinct elements of X_k. Let v be a Luce model for p_k. Then u is a Luce model for p_k iff $u = \lambda v$ for some $\lambda > 0$.*

Proof. We already observed the "if" part. For the "only if" part, let v and u be Luce models for p_k, and let x and y be two distinct elements of X_k. Let $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ be a connected sequence with $x_1 = x$ and $x_{m+1} = y$. We therefore have:

$$\frac{p_k(x_1, E_1)}{p_k(x_2, E_1)} \frac{p_k(x_2, E_2)}{p_k(x_3, E_2)} \cdots \frac{p_k(x_m, E_m)}{p_k(x_{m+1}, E_m)} = \frac{v(x)}{v(y)} = \frac{u(x)}{u(y)}$$

Since this holds for every distinct x and y in X_k, the result follows. □

Thus, if p has a Luce model, then this model is unique up to one strictly positive multiplicative constant per connected component of $G(X, \mathcal{M})$. If p is decomposed as above, we may normalise any Luce model so that it sums to 1 on each X_k. This leaves $|X| - \kappa$ free parameters to specify. By Theorem 7:

$$|X| - \kappa = \left[\left(\sum_{A \in \mathcal{M}} |A|\right) - |\mathcal{M}|\right] - C(X, \mathcal{M})$$

The square-bracketed term is the number of free parameters in the specification of an RCF: for each $A \in \mathcal{M}$ we need to specify $|A| - 1$ values. The quantity $C(X, \mathcal{M})$ is precisely the difference between the degrees of freedom in the specification of an RCF and the degrees of freedom in the specification of a Luce model for p. Given positivity, we may specify p using log-probabilities, in which case the cycle equations are *linear* restrictions on the log-probabilities (see Appendix B). A Luce model imposes no more than $C(X, \mathcal{M})$ independent restrictions on this representation of an RCF.

6 Concluding remarks

Given positivity, CI is the empirical signature of Luce rationality (i.e., existence of a Luce model) when choice is observed from an arbitrarily fixed set of menus. We have shown that essentially the same insight underpins the so-called "cycles approach" to the common prior problem when X is finite. The literature on the latter problem provides graph-theoretic tools to reduce the number of cycle equations that need to be checked in order to verify CI. It also furnishes simple algorithms for identifying a suitable "cycle basis".

This theoretical apparatus reduces the empirical signature of Luce rationality to a small set of *linear* restrictions on the *logs* of choice probabilities (i.e., the log-transformed cycle equations). It remains to develop a suitable empirical methodology for testing these conditions. This is the subject of on-going research.

Appendix

A Graph theory toolkit

This appendix summarises the basic concepts and results from graph theory that are used in our analysis. The reader will find more thorough treatments in any introductory text (e.g., Diestel, 2010; Wallis, 2000). All results are well-known but we have included proofs to keep the paper self-contained. Our analysis deals with multigraphs, and proofs for this context are often omitted from texts or left as exercises for the reader.

A.1 Notation and definitions

For convenience, we use the term *graph* to include multigraphs. A graph has a finite vertex set, V, and a finite *family* (indexed set)

$$\mathcal{E} = \{e^i\}_{i \in \mathcal{I}}$$

of edges. Each edge is an unordered pair of distinct vertices. Edge $e^i = \{x, y\}$ joins the vertices x and y, which are said to be *incident* to the edge. Note that graphs are undirected by definition, and graphs have no loops.[25] Importantly, there may be multiple edges joining the same pair of vertices (i.e., our notion of a graph admits multigraphs). To avoid notational ambiguity, we use e^i to denote the edge with index $i \in \mathcal{I}$, while subscripts (e.g., e_i) are used to refer to arbitrary elements of $\mathcal{E}$. Also, if $e^i = \{x, y\} \subseteq V$ for some $i \in \mathcal{I}$, then we abuse notation and write $\{x, y\} \in \mathcal{E}$.

A *subgraph* of graph $G = (V, \mathcal{E})$ is a subset $V' \subseteq V$ of vertices, together with a subset $\mathcal{E}' \subseteq \mathcal{E}$ of vertices, such that any vertex incident to some edge in $\mathcal{E}'$ is included in V'. The subgraph $G' = (V', \mathcal{E}')$ is *spanning* if $V' = V$. If $e \in \mathcal{E}$, then $G - e$ denotes the spanning subgraph of G obtained by removing e from the edge set. If $G' = (V', \mathcal{E}')$ is a spanning subgraph of $G = (V, \mathcal{E})$ and $e \in \mathcal{E} \setminus \mathcal{E}'$ then $G' + e$ denotes the spanning subgraph of G obtained by adding the edge e to G'.

A *walk* of length ℓ in G is a sequence $(v_0, e_1, v_1, e_2, ..., e_\ell, v_\ell)$ with each $v_i \in V$ and each $e_i \in \mathcal{E}$ with $e_i = \{v_{i-1}, v_i\}$.[26] A walk is *simple* if its edges are all distinct elements of $\mathcal{E}$. A *path* is a simple walk whose vertices are all distinct. A walk of length ℓ is *closed*

[25] A *loop* is an edge joining some vertex to itself.

[26] If we excluded multigraphs, then a walk could just as well be described by its sequence of vertices.

if $v_0 = v_\ell$. A closed walk of length ℓ is a *cycle* (in the language of graph theory)[27] if

$$(v_0, e_1, v_1, e_2, ..., e_{\ell-1}, v_{\ell-1})$$

is a path. Any cycle has length at least two.[28] Note that we may think of paths and cycles as subgraphs.

A graph is *connected* if there is a walk (hence a path)[29] joining any two distinct vertices. The relation "is joined by a walk to" is a symmetric and transitive binary relation on V, and therefore its reflexive closure partitions the graph into maximal connected subgraphs known as its *components*. These can effectively be treated as separate graphs, so for most purposes it suffices to restrict attention to connected graphs.

A.2 Spanning trees

A *tree* is a connected graph with no cycles. A *bridge* is an edge whose removal from a graph increases the number of the graph's components. A bridge partitions V into two subsets, $\{V_1, V_2\}$, such that any walk from V_1 to V_2 (or from V_2 to V_1) must cross the bridge. It follows that no walk that includes a bridge can be a cycle: it would need to cross the bridge twice (once in each direction) so it would include a repeated edge. Hence, if G is connected and *every* edge is a bridge, then G is a tree. The converse is also true:

Lemma 3 *A connected graph is a tree iff every edge is a bridge.*

Proof. We already proved the "if" part. For the converse, let $G = (V, \mathcal{E})$ be connected and suppose $e = \{x, y\} \in \mathcal{E}$ is not a bridge. Then there is a walk – hence a path – from y to x (since G is connected) that does not include e (since e is not a bridge). Appending edge e to this path creates a cycle. □

As we noted, deletion of a bridge creates a two-element partition of the vertex set, and splits the graph into two disjoint subgraphs on the respective cells of the partition. A recursive application of this fact, together with Lemma 3, imply:

Corollary 3 *Let $G = (V, \mathcal{E})$ be connected. Then G is a tree iff $|\mathcal{E}| = |V| - 1$.*

[27] In the language of stochastic choice, "connected sequences" with $x_i \neq x_{i+1}$ for all i are analogous to walks, and the "cycles" of cyclical independence are analogous to closed walks (also known as *circuits*). See Appendix B for more on mapping cyclical independence into the vernacular of multigraphs.

[28] If we excluded multigraphs, cycles would have length at least three.

[29] If an edge is crossed twice in the same direction along the walk, then we can omit all of the walk in between plus one occurrence of this edge; if it is crossed twice in opposite directions, we can omit everything in between plus both occurrences.

A *spanning tree* for a graph is a subgraph that includes all the vertices of the graph and which is itself a tree. It is obvious that only connected graphs can possess a spanning tree. Moreover, all connected graphs do so:[30]

Lemma 4 *Any connected graph has a spanning tree.*

Proof. Let $G = (V, \mathcal{E})$ be a connected graph. If it has no cycle we are done. Otherwise, fix some cycle $C_1 = (V, \mathcal{E}^1)$ in G and some $i_1 \in \mathcal{I}$ such that $e^{i_1} \in \mathcal{E}^1$. Form the spanning subgraph $G_1 = G - e^{i_1}$ of G. Since e^{i_1} cannot be a bridge (recall that no cycle can include a bridge), it follows that G_1 is connected. If G_1 has no cycle, we are done. If not, fix some cycle $C_2 = (V, \mathcal{E}^2)$ in G_1 and some $i_2 \in \mathcal{I}$ such that $e^{i_2} \in \mathcal{E}^2$. Now form the subgraph $G_2 = G_1 - e^{i_2}$. Then G_2 is a connected spanning subgraph of G. And so on. This process must end in finitely many steps, at which point we have a spanning subgraph of G that is a tree. □

Suppose $\mathcal{N} = \{i_k\}_{k=1}^K \subseteq \mathcal{I}$ are the indices of the edges deleted from $\mathcal{E}$ in this process, with the subscript indicating the order of deletion (i.e., edge e^{i_1} is the edge deleted first). Let $\overline{\mathcal{N}} = \mathcal{I} \setminus \mathcal{N}$. Then

$$T = \left(V, \left\{e^i\right\}_{i \in \overline{\mathcal{N}}}\right)$$

is the spanning tree obtained at the end of the process. For each $k \in \{1, 2, \ldots, K\}$, let $G^k = T + e^{i_k}$.

Lemma 5 *The graph G^k contains a unique cycle and this cycle includes e^{i_k}.*

Proof. Any cycle in G^k must include e^{i_k} since T is a tree. There must exist a cycle in G^k since the vertices incident to e^{i_k} are joined by a path in T. Finally, there cannot be more than one such cycle: if there were two, then they must have edge e^{i_k} in common, so we could construct a cycle in T by removing this common edge.[31] □

The family $\left\{C^k\right\}_{k=1}^K$ will play an important role in our analysis. It is called the *fundamental basis* for the *cycle space* of G. As we show in Appendix C, all other cycles in G can be generated – in an appropriate linear sense – from this family.

[30] Hence, any graph possesses a *spanning forest*, which is a spanning subgraph comprising a disjoint union of trees.

[31] Let x and y denote the vertices incident to e^{i_k} and let C' and C'' be two cycles in G^k. Then C' contains a path in T from x to y and C'' contains a path in T from y to x. Concatenating these paths determines a circuit (i.e., closed walk) in T. It follows that T contains a cycle – recall footnote 29.

B Proof of Theorem 7

Cyclical independence requires that the following *cycle equation* be satisfied for any *positive* connected sequence $\{(x_i, x_{i+1}, E_i)\}_{i=1}^{m}$ with $x_1 = x_{m+1}$:

$$\prod_{i=1}^{m} \frac{p(x_i, E_i)}{p(x_{i+1}, E_i)} = 1 \qquad (\dagger)$$

It suffices to consider connected sequences that also satisfy $|\{x_1, x_2, ..., x_m\}| = m$. If the same alternative appears twice along the sequence then we can split it into two cyclical subsequences; provided the cycle equation for each subsequence is satisfied then the original cycle equation will be satisfied.

Recall that $G(X, \mathcal{M})$ is the (multi)graph with vertex set X and edge set indexed by

$$\mathcal{I} = \{[xy; E] \,|\, \{x, y\} \subseteq E \in \mathcal{M}\}.$$

We therefore need to check that (†) holds along any *cycle* – in the sense of Appendix A.1 – in $G(X, \mathcal{M})$. To be clear, graph $G(X, \mathcal{M})$ is undirected while equation (†) assumes a particular direction of travel around the cycle associated with the connected sequence. However, (†) says precisely that the product on the left is independent of the direction of travel (or the vertex at which the walk starts and ends), so we can unambiguously talk about equation (†) holding, or being satisfied, along a cycle in $G(X, \mathcal{M})$.

In fact, we can restrict attention to cycles within the spanning subgraph associated with some *version* of $G(X, \mathcal{M})$. If the subsequence from $x_j = x'$ to $x_k = x''$ (for $1 \leq j < k \leq m+1$) is *within* some menu (i.e., $E_j = E_{j+1} = \cdots = E_k$), then it may be replaced by any other subsequence from x' to x'' within the same menu without affecting the cycle equation, as is easily verified. Restricting attention to a version of $G(X, \mathcal{M})$ eliminates these redundancies in the cycle equations: for any menu E and any two disinct elements of E, a version contains exactly one path within E that joins these two elements. In particular, a version contains no cycle within any menu.

Let $G^* = (X, \mathcal{E}^*)$ be the spanning subgraph generated by a given version of $G(X, \mathcal{M})$. Let $\mathcal{I}^* \subseteq \mathcal{I}$ be the index set for $\mathcal{E}^*$. Assuming positivity of p, cyclical independence is equivalent to satisfaction of the cycle equation (†) associated with any cycle in G^*.

The proof of Theorem 7 now proceeds as follows. We first map each cycle of G^* to a vector in a suitable vector space, and consider the subspace spanned (in the linear sense) by these vectors. We then use a standard result from graph theory to establish the $C(X, \mathcal{M})$ is the dimension of this subspace. Finally, we show that satisfaction of the cycle equations associated with the cycles in a basis for this subspace suffices for cyclical independence.

To construct a suitable vector space, let

$$\iota : \mathcal{E}^* \to \{1, 2, ..., |\mathcal{E}^*|\}$$

be an enumeration of the edges. Define $\mathcal{V}$ to be the vector space $\mathbb{R}^{|\mathcal{E}^*|}$ with edge $[xy; E]$ associated to coordinate $\iota([xy; E])$. We next map each cycle in G^* to a vector in $\mathcal{V}$.[32] To do so, let $\overrightarrow{G}^*$ be a *directed* multigraph that adds an orientation to each edge of G^* (so "edges" become "arcs"). Following Berge (1962), each cycle in G^* is mapped to a vector in $\mathcal{V}$ by setting the value of the $\iota([xy; E])$ coordinate equal to: 1 if the cycle traverses edge $[xy; E]$ in the direction of the corresponding arc in $\overrightarrow{G}^*$; -1 if the cycle traverses edge $[xy; E]$ in the opposite direction; and 0 if the cycle does not traverse $[xy; E]$. We call this the *cycle vector* for the given cycle in G^*. Let $\mathcal{C}$ denote the subspace of $\mathcal{V}$ spanned by the zero vector together with the cycle vectors for all cycles in G^*. This subspace is called the *cycle space* of $\overrightarrow{G}^*$.

By definition, the dimension of $\mathcal{C}$ is the *cyclomatic number* of G^*. This dimension is independent of the particular $\overrightarrow{G}^*$ used to fix the orientations of edges.[33] A classical result in graph theory (see, for example, Berge, 1962, Chapter 4, Theorem 2)[34] shows that the cyclomatic number of G^* is equal to $|\mathcal{E}^*| - |X| + \kappa$. In our case,

$$|\mathcal{E}^*| \;=\; \sum_{A \in \mathcal{M}} (|A| - 1) \;=\; \left(\sum_{A \in \mathcal{M}} |A| \right) - |\mathcal{M}|$$

so the cyclomatic number of G^* is $C(X, \mathcal{M})$. Since κ is the same for any version of $G(X, \mathcal{M})$, *all versions have the same cyclomatic number*, so we call this common value the cyclomatic number of $G(X, \mathcal{M})$.

[32]Modern treatments of graph theory (e.g., Diestel, 2010; Wallis, 2000) encode the cycles in a vector space defined over the two-element field $\mathbb{Z}_2$ comprising the integers $\{0, 1\}$ together with addition and multiplication modulo-2, rather than a real vector space. However, the cycle equation (†) is linear in the Euclidean sense, so the older set-up of Berge (1962) is more convenient for our purposes.

[33]Let $\{x^1, x^2, ..., x^K\}$ be a linearly dependent collection of cycle vectors for a particular orientation. Then there exists $\beta \in \mathbb{R}^K$ with $\beta \neq 0$ and $\sum_{k=1}^K \beta_k x^k = 0$. If we change orientations, we simply reverse the signs of the x^k vectors in specified components; the corresponding rows of $\sum_{k=1}^K \beta_k x^k$ get multiplied by -1. Therefore, letting the resulting set of vectors be denoted $\{\hat{x}^1, \hat{x}^2, ..., \hat{x}^K\}$ we still have $\sum_{k=1}^K \beta_k \hat{x}^k = 0$. Thus, a given set of cycles has linearly dependent cycle vectors for one orientation iff it has a linearly dependent set of cycle vectors for all orientations.

[34]Berge considers the space of vectors associated with all *circuits*, not just cycles. For circuits that are not cycles, the $\iota([xy; E])$ coordinate of the associated vector is equal to (zero plus) the number of times that edge $[xy; E]$ is traversed in the positive orientation less the number of times it is traversed in the opposite direction. It is evident that this vector is the sum of the vectors associated with sub-cycles into which the original circuit may be decomposed. Therefore, the dimension of the cycle subspace is the same whether or not we include the vectors associated with circuits.

Next, we re-write the cycle equation (†) in the following form:

$$\sum_{i=1}^{m} \left(\ln\left[p\left(x_i, E_i\right)\right] - \ln\left[p\left(x_{i+1}, E_i\right)\right]\right) \;=\; 0 \tag{††}$$

Let $v \in \mathcal{V}$ be the vector representation for the associated cycle in $\overrightarrow{G}^*$. Then equation (††) may be expressed $\alpha \cdot v = 0$, where $\alpha \in \mathcal{V}$ is defined as follows:

$$\alpha_{\iota([xy;E])} \;=\; \ln\left(p\left(\underline{z}, E\right)\right) - \ln\left(p\left(\overline{z}, E\right)\right)$$

with $\{\underline{z}, \overline{z}\} = \{x, y\}$ and the arc in $\overrightarrow{G}^*$ corresponding to edge $[xy; E]$ in G^* is oriented from $\underline{z}$ to $\overline{z}$. The cycle equations associated with **all** cycles of G^* are therefore satisfied iff α is in the **orthogonal complement** of $\mathcal{C}$. For this, it suffices that α is orthogonal to each vector in a basis for $\mathcal{C}$. The dimension of $\mathcal{C}$ gives the number of elements in a basis.

In summary, any basis for $\mathcal{C}$ has $C(X, \mathcal{M})$ elements, and CI is satisfied for $G\left(X, \mathcal{M}\right)$ iff each of these elements is orthogonal to α. This is equivalent to cycle equation (††) being satisfied for the $C(X, \mathcal{M})$ cycles associated with the basis vectors.

C Identifying a “cycle basis”

Let $\mathcal{C} \subseteq \mathcal{V}$ be the cycle space for a given version, $G^* = (X, \mathcal{E}^*)$, of $G\left(X, \mathcal{M}\right)$. A *cycle basis* is a basis for the subspace $\mathcal{C}$. Any collection of $C(X, \mathcal{M})$ linearly independent vectors in $\mathcal{V}$ will constitute such a basis.

Suppose that G^* is connected (i.e., $\kappa = 1$). Let $T = (X, \hat{\mathcal{E}})$ be the spanning tree for G^*, and let $\{e^{i_k}\}_{k=1}^{K} \subseteq \mathcal{E}^*$ be the edges removed to construct T. Note that the associated *fundamental basis*, $\left\{C^k\right\}_{k=1}^{K}$, is linearly independent: if $z^k \in \mathcal{C}$ is the vector associated with C^k then $z^k_{i_\ell} \neq 0$ iff $k = \ell$. By Corollary 3, $|\hat{\mathcal{E}}| = |X| - 1$ so

$$K \;=\; |\mathcal{E}^*| - |X| + 1 \;=\; C(X, \mathcal{M}).$$

It follows that the $\left\{C^k\right\}_{k=1}^{K}$ forms a basis for $\mathcal{C}$.

It is useful to note that the family $\{C_k\}_{k=1}^{K}$ of cycles used in the construction of T generates another cycle basis. If $\hat{z}^k \in \mathcal{C}$ is the vector associated with C_k then the family $\left\{\hat{z}^k\right\}_{k=1}^{K}$ is also linearly independent: $\left|\hat{z}^k_{i_k}\right| = 1$ while $\hat{z}^j_{i_k} = 0$ for all $j > k$.

Now suppose $\kappa > 1$. The components of G^* determine subgraphs which are individually connected and together partition G^*. Any cycle in G^* will be contained in one of those subgraphs. Fixing a cycle basis for each subgraph, their union will therefore be a cycle basis for G^*.

REFERENCES

Ahumada, A. and Ülkü, L. (2018). "Luce rule with limited consideration," *Mathematical Social Sciences* 93, 52-56.

Berge, C. (1962). *The Theory of Graphs and its Applications.* Methuen London. [Reprinted by Dover, 2001]

Cerreia-Vioglio, S., Lindberg, P.O., Maccheroni, F., Marinacci, M., and Rustichini, A. (2021). "A canon of probabilistic rationality," *Journal of Economic Theory* 196, 105289.

Diestel, R. (2010). *Graph Theory (4th edition).* Springer.

Echenique, F. and Saito, K. (2019). "General Luce model," *Economic Theory* 68, 811-826.

Estes, W. K. (1960) "A random-walk model for choice behavior," in Arrow, K.J., Karlin, S. and Suppes, P. (eds), *Mathematical Methods in the Social Sciences, 1959: Proceedings of the First Stanford Symposium.* Stanford University Press.

Fiorini, L.C. and Rodrigues-Neto, J.A. (2017). "Self-consistency, consistency and cycles in non-partitional knowledge models," *Mathematical Social Sciences* 87, 11-21.

Harsanyi, J.C., (1967-1968). "Games with incomplete information played by Bayesian players, I-III" *Management Science* 14, 159-182; 320-334; 486-502.

Hellman, Z. Samet, D. (2012). "How common are common priors?", *Games and Economic Behavior* 74, 517-525.

Hellwig, M.F. (2013). "From posteriors to priors via cycles: An addendum," *Economics Letters* 118, 455-458.

Kovach, M., and Tserenjigmid, G. (2022). "The focal Luce model." *American Economic Journal: Microeconomics* 14(3), 378-413.

Luce, R.D. (1959). *Individual Choice Behaviour: A Theoretical Analysis.* John Wiley & Sons, New York.

Luce, R. D., and Suppes, P. (1965). "Preferences, utility and subjective probability," in Luce, R.D., Bush, R.R. and Galanter, E. (eds) *Handbook of Mathematical Psychology, vol. III.* John Wiley & Sons: New York.

McCausland, W. J., Davis-Stober, C., Marley, A. A., Park, S., and Brown, N. (2020). "Testing the random utility hypothesis directly." *The Economic Journal* 130(625), 183-207.

McFadden, D. (1973). "Conditional Logit analysis of qualitative choice behavior," in Zarembka, P. (ed.) *Frontiers in Econometrics.* Academic Press.

Rodrigues-Neto, J.A. (2009). "From posteriors to priors via cycles," *Journal of Eco-*

nomic Theory 144, 876-883.

Rodrigues-Neto, J.A. (2012). "The cycles approach," *Journal of Mathematical Economics* 48, 207-211.

Rodrigues-Neto, J.A. (2014). "Monotonic models and cycles," *International Journal of Game Theory* 43, 403-413.

Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press.

Wallis, W. D. (2000). *A Beginner's Guide to Graph Theory*. Birkhäuser.