Random Utility with Experimentation

Jack Adeney

December 17, 2025

PROOF OF CONCEPT

The Random Utility Model (McFadden, 1974) is one of the canonical models used to rationalize stochastic choice data. Intuitively, it can be thought of as fitting a distribution over decision makers such that the expected choices over this distribution are equal to the actual choice proportions observed in the data. However, not all stochastic choice datasets have a random utility representation. For example, take a set of three alternatives, $X = \{x, y, z\}$, and let the probability of choosing alternative x from the menu $A \subseteq X$ be denoted $\rho_A(x)$. Choice data involving "stochastic cycles", such as $\rho_{\{x,y\}}(x) = 0.6$, $\rho_{\{y,z\}}(y) = 0.6$, and $\rho_{\{x,z\}}(z) = 0.6$ have no distribution over "rational" decision makers such that the expected choice proportions match this data. Further, the literature is yet to unanimously support a measure of distance to rationality when no representation exists. In this note, I generate an algorithm to rationalize stochastic choice data with a mixture of both rational and irrational choice functions. This representation is designed to place more weight on the rational actors within the system, and then allocate residual probabilities to choice functions with increasing distance from the rational choice functions. This method has an intuitive interpretation when thinking about these data as aggregations of individual decision maker choices.

Let X represent the finite set of alternatives over which individuals will be choosing. Assume $|X| = n$, and the set of all menus, denoted $\mathcal{A}$ is the set of all non-empty subsets $A \subseteq X \backslash \emptyset$ over X. A choice function $c : \mathcal{A} \to X$ is a mapping from a menu to an alternative within that menu. Choice functions can be either rational or irrational in the traditional sense. We denoted these sets of choice functions as $\mathcal{C}_R$ and $\mathcal{C}_I$ respectively.

For all $c \in \mathcal{C}_R$, there exists some (strict) linear order π on X such that, for all $A \in \mathcal{A}$

$$c(A) = \{a \in A : a \succ_\pi b \ \ \forall \ \ b \in A \backslash \{a\}\}$$

In words, the choice function represents an agent with a complete, transitive, and asymmetric ranking over X, and from each menu A, picks the best element in A according to that ranking. As there are $n!$ strict linear orders on a set of cardinality n, there are also $n!$ distinct choice functions in $\mathcal{C}_R$. $\mathcal{C}_I$ is the set of all choice functions on X not in $\mathcal{C}_R$.

We will focus on stochastic choice datasets that are complete, meaning that we can observe choice proportions for all alternatives over all menus. Let this dataset be represented as $\{\rho_A(a)\}_{\{a\in A, A\in\mathcal{A}\}}$. The dataset is random utility rationalizable if there exists some probability measure $\mu \in \Delta(C_R)$ such that, for all $a \in A$ and $A \in \mathcal{A}$

$$\rho_A(a) = \mu(\{c \in \mathcal{C}_R : c(A) = a\})$$

But what happens when a dataset is not random utility rationalizable? In that case there are a number of methods that we could use in order to interpret the 'consistency' of such choice data. A natural approach would be to expand the support of the rationalizing distribution from the set of all rational choice functions into the set of irrational choice functions. Given that this set now contains all possible choice functions, there will be a rationalizing distribution for every dataset, and so we require conditioning in order to make the rationalizing meaningful.

To do this, consider a single linear order $\pi \in \Pi$, and let the corresponding choice function be labelled $c_\pi(\cdot)$. Define the set C^1_π as the set of all choice functions that agree with c_π over all menus besides one, in which they pick a different alternative. For example, if $X = \{x, y, z\}$ and $\pi = x \succ y \succ z$ then the choice function c_π could be represented by the vector in *Table 1*. The matrix of choice vectors that are one mistake away from c_π are shown in *Table 2*.

Alternative	Menu	$c_\pi(\cdot)$
x	$\{x, y\}$	1
y	$\{x, y\}$	0
x	$\{x, z\}$	1
z	$\{x, z\}$	0
y	$\{y, z\}$	1
z	$\{y, z\}$	0
x	$\{x, y, z\}$	1
y	$\{x, y, z\}$	0
z	$\{x, y, z\}$	0

Table 1: Rational Choice Vector: $x \succ y \succ z$

Alternative	Menu	$c^1_\pi(\cdot)$				
x	$\{x, y\}$	0	1	1	1	1
y	$\{x, y\}$	1	0	0	0	0
x	$\{x, z\}$	1	0	1	1	1
z	$\{x, z\}$	0	1	0	0	0
y	$\{y, z\}$	1	1	0	1	1
z	$\{y, z\}$	0	0	1	0	0
x	$\{x, y, z\}$	1	1	1	0	0
y	$\{x, y, z\}$	0	0	0	1	0
z	$\{x, y, z\}$	0	0	0	0	1

Table 2: Irrational Choice Matrix C^1_π

Let $C_1 = \cup_{\pi \in \Pi} C_\pi^1$. In the same way that the original random utility rationalization searched for a probability measure over the matrix C_R, we can now look for a probability measure over the matrix $[C_R, C_1]$.

To put this concretely, suppose that C_R is an $i \times j$ matrix and C_1 is an $i \times k$ matrix. This means that we have a maximal support cardinality of $j + k$. The weights can therefore also be written as $\mu_1, ..., \mu_j, \mu_{j+1}, ..., \mu_{j+k}$. The solution to this new problem takes the form of some $\mu' \in \Delta([C_R, C_1])$ such that

$$\rho_A(a) = \mu'(\{c \in C_R \cup C_1 : c(A) = a\})$$

Due to the rate at which the number of columns in C_1 increases with n, the likelihood of there existing a probability measure on $[C_R, C_1]$ is large. Further, a rationalizing probability measure with most weight on these 'mistaken' choice functions does not have a particularly generalizable interpretation. Therefore, it makes sense to push weight onto the rational choice functions where possible. Fortunately, one of the major features of random utility models is that they often provided a multiplicity of rationalizing measures. This can be used in our model to select for the rationalzing measure that places most weight on the rational choice functions.

We can place penalties on weights λ_0 and λ_1 such that $\lambda_0 > \lambda_1$ and stack them into a vector $\vec{\lambda}$ where the first j elements are λ_0 and the last k elements are λ_1. We can therefore look for a following solution:

$$\begin{aligned} \max_{\mu} \quad & \vec{\lambda}\mu \\ s.t. \quad & \rho = [C_R, C_1]\mu \end{aligned}$$

This will search among the rationalizing measures μ' to find the one that places most weight on the rational choice functions.

If there exists no such μ' then we can continue iterating the process to $[C_R, C_1, C_2, ...]$ with penalty terms $\lambda_0 > \lambda_1 > ...$ until eventually we find a rationalizing measure.

In order to test this model, I have simulated a set of stochastic choice data in which a pre-specified proportion of decisions are made according to rational choice functions, and the remaining proportion are chosen at random. For example, if the proportion of rational choice functions, $f_{rat} = 0.5$ and set set of available alterantives $X = \{x, y, z\}$ then we randomly select 3 of the 6 rational choice functions and build a dataset that may look something like this:

Alternative	Menu	$x \succ y \succ z$	$y \succ x \succ z$	$z \succ x \succ y$	$c_{I,1}$	$c_{I,2}$	$c_{I,3}$
x	$\{x, y\}$	1	0	1	$p_{4,1}$	$p_{5,1}$	$p_{6,1}$
y	$\{x, y\}$	0	1	0	$1 - p_{4,1}$	$1 - p_{5,1}$	$1 - p_{6,1}$
x	$\{x, z\}$	1	1	0	$p_{4,2}$	$p_{5,2}$	$p_{6,2}$
z	$\{x, z\}$	0	0	1	$1 - p_{4,2}$	$1 - p_{5,2}$	$1 - p_{6,2}$
y	$\{y, z\}$	1	1	0	$p_{4,3}$	$p_{5,3}$	$p_{6,3}$
z	$\{y, z\}$	0	0	1	$1 - p_{4,3}$	$1 - p_{5,3}$	$1 - p_{6,3}$
x	$\{x, y, z\}$	1	0	0	$p_{4,4}$	$p_{5,4}$	$p_{6,4}$
y	$\{x, y, z\}$	0	1	0	$p_{4,5}$	$p_{5,5}$	$p_{6,5}$
z	$\{x, y, z\}$	0	0	1	$1 - p_{4,4} - p_{4,5}$	$1 - p_{5,4} - p_{5,5}$	$1 - p_{6,4} - p_{6,5}$

Table 3: Simulation of irrational data, $f_{rat} = 0.5$

The first three columns of *Table 3* correspond to rational choice functions, whereas the final three consist of probabilities generated at random. We then take a random convex combination of these columns to generate our simulated stochastic choice data $\hat{\rho}$.

We first search over $[C_R]$ to see if there is a random utility representation in the traditional sense. If not, then we move to searching for a probability measure over $[C_R, C_1]$. Typically, one will exist when $n = 3$. This gives us a measure μ, and, because each choice function in C_1 can be mapped to a choice function in C_R, we can break up the weights on linear order π as $\mu^\pi = \mu_R^\pi + \mu_I^\pi$. The first term corresponds to the weight placed on the rational choice function, c_π, and the second term corresponds to the weight placed on the irrational choice function one mistake away from c_π.

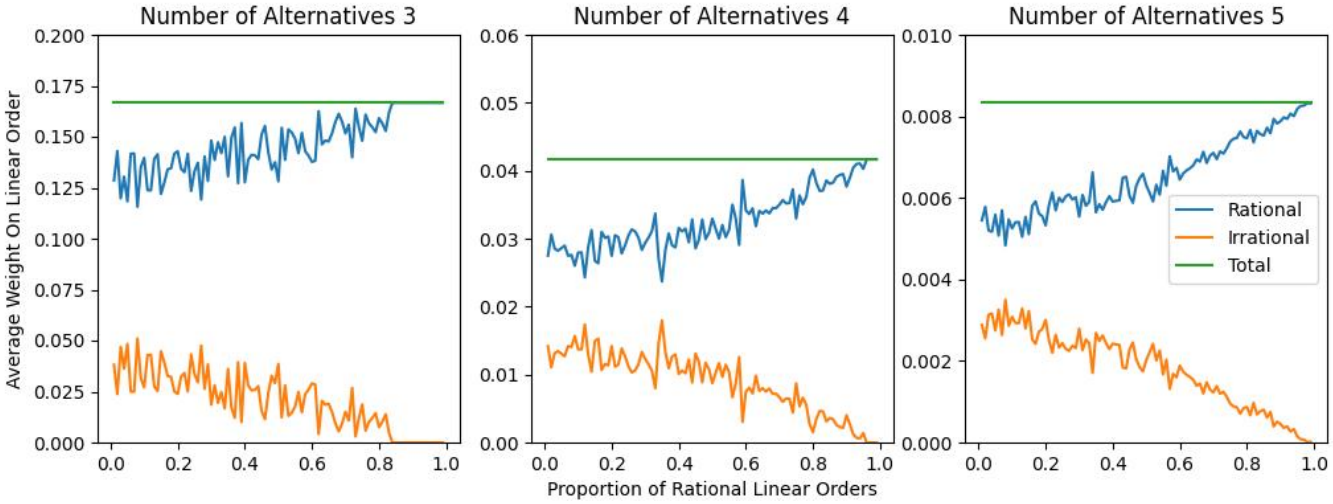

Figure 1: Rational and Irrational Weight Distributions

Figure 1 shows the average split of μ_R and μ_I as we increase the proportion of rational choice functions/linear orders. Clearly, as the proportion of rational choice functions increases, the average weight placed on rational choice functions increases and the average weight placed on irrational choice functions decreases.

The figure above shows the aggregated split of rational and irrational choice weights. However, our method of choice data simulation means that we actually know the distribution of rational and irrational choices. This means that we can observe how much irrational versus rational weight is placed on each function conditional on whether they were rational or irrational in the data simulation.

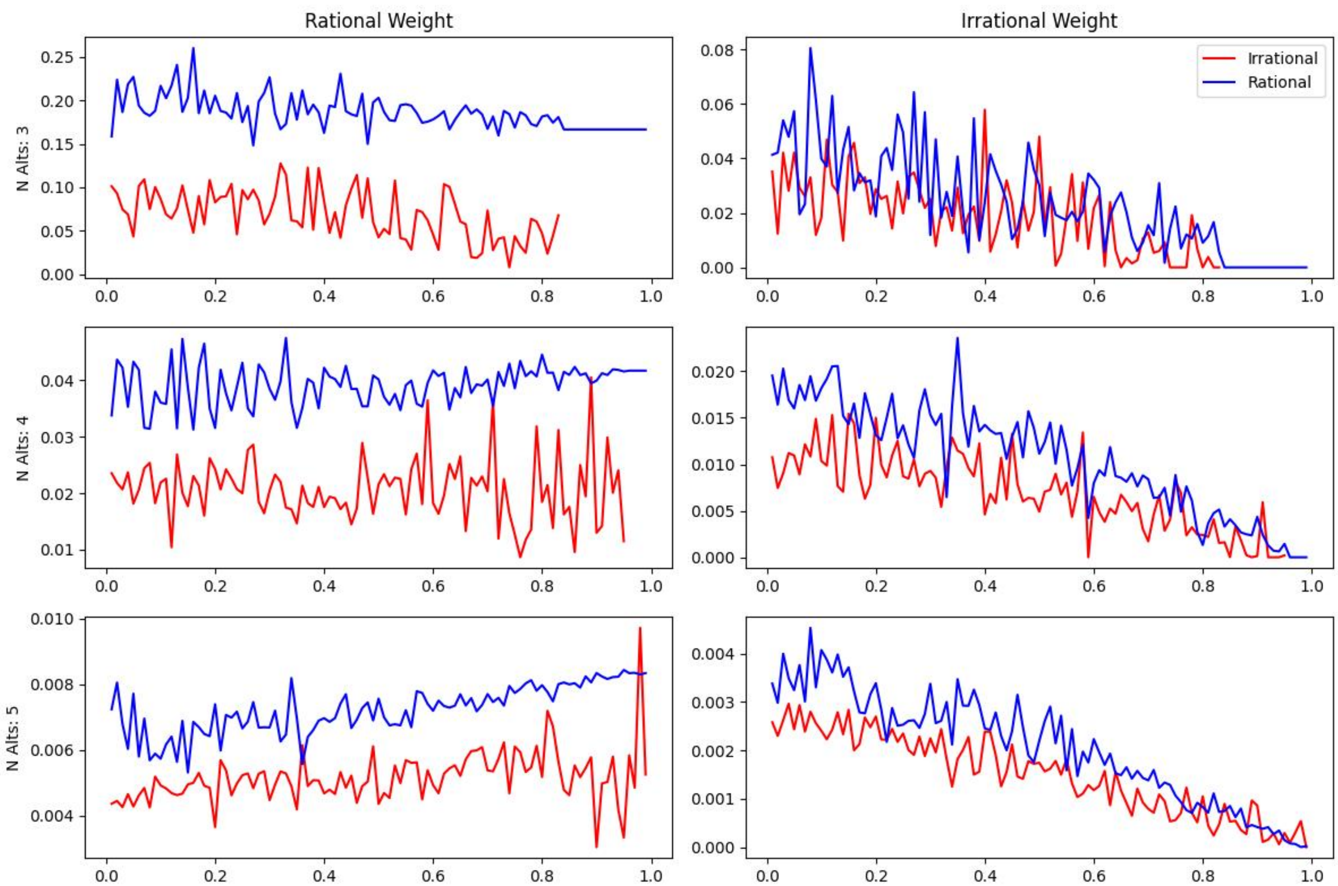

Figure 2: Rational and Irrational Weights by Choice Function Type

Figure 2 shows exactly this. The left-hand column of figures shows the rational and irrational weights placed on choice functions that were in fact rational in our dataset (for example, the first three columns of *Table 2*). The x-axis once again represents the proportion of functions that were rational. The right-hand column shows the rational and irrational weights placed on the irrational choice functions. The first thing to note is that the majority of weight is placed on the rational part of the function where possible, as demonstrated by the blue line being above the red line in all figures. Interestingly, the difference between the weight placed on rational versus irrational parts is greater for rational choice functions than for irrational choice functions. Intuitively, one might expect the irrational proportion of weight placed on rational choice functions to be equal to 0. However, spillover from the irrational choice functions results in this not being the case. Similarly, the rational weight placed on irrational choice functions is non-zero due to the system extracting 'rational behavior' from the irrational choice probabilities.

References

D. McFadden. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, 1974.