

ONLINE APPENDIX: RANDOM UTILITY MODELS WITH ORDERED TYPES AND DOMAINS

JOSE APESTEGUIA[†] AND MIGUEL A. BALLESTER[‡]

1. PRACTICAL IMPLEMENTATION OF \mathcal{T} -RUM

We now show how to implement the \mathcal{T} -RUM in practice, discussing each of the steps required for the analysis of a dataset. We use the experimental dataset analyzed in Apestegua and Ballester (2021) that involves choices over lotteries and, since the purposes are purely expository of the framework, we adopt the most convenient empirical implementation strategy, namely, we assume that all decisions are made following the same distribution over CRRA expected utility types. To emphasize, the selection of expected utilities with CRRA monetary utilities, and the consideration of aggregate data with the assumption that all individuals are equally stochastic has only an illustrative purpose; we simply aim for a better understanding of how to implement our results in a given economic setting.

We start by describing the dataset. Eighty-seven UCL undergraduates were asked to choose lotteries from menus of sizes 2, 3, and 5, from the nine equiprobable monetary lotteries described in Table 1. Each participant faced a total of 108 different menus of lotteries, including all 36 binary menus; 36 menus with 3 alternatives, out of the possible 84; and another 36 menus with 5 alternatives, out of the possible 126. Random individual processes, without replacement, were used for the selection and order of presentation of the menus of 3 and 5 alternatives, and for the location of the lotteries on the screen and the monetary prizes within a lottery. There were two treatments, but for present purposes we merge all the data from the experiment.¹

Date: March, 2023.

[†]ICREA, Universitat Pompeu Fabra and Barcelona GSE. E-mail: jose.apestegua@upf.edu.

[‡]University of Oxford. E-mail: miguel.ballester@economics.ox.ac.uk.

¹Experimental payoffs were determined by randomly selecting one menu, and awarding subjects in accordance with their choice from that menu.

TABLE 1. Lotteries

$l_1 = (17)$	$l_4 = (30, 10)$	$l_7 = (40, 12, 5)$
$l_2 = (50, 0)$	$l_5 = (20, 15)$	$l_8 = (30, 12, 10)$
$l_3 = (40, 5)$	$l_6 = (50, 12, 0)$	$l_9 = (20, 12, 15)$

Step 1 (Set of Types). The first step involves the selection of an ordered family of ordinal preferences $\{U_t\}_{t \in \mathcal{T}}$. For example, one may start with a well-known family of utilities that is sufficiently rich, and look for the parameter values that yield indifference between a pair of alternatives in the dataset. Then, \mathcal{T} could be specified by considering preferences right below and above the parameters that create indifference. Note that, in \mathcal{T} -RUMs, all that matters is the ordinal preference of alternatives, and two utility functions that are not separated by any of these thresholds are associated with the same ordinal preference.

Given that our dataset involves lottery choices, here we adopt the most standard practice and use expected utility with CRRA monetary utilities, where $u_\omega(x) = \frac{x^{1-\omega}}{1-\omega}$, whenever $\omega \neq 1$, and $u_1(x) = \log x$, where ω represents the risk-aversion coefficient.² Consider every risk aversion coefficient ω for which two lotteries in the experiment are indifferent. This corresponds to the 29 finite values reported in Columns 2 and 6 of Table 2. Now, to specify the family \mathcal{T} consider one utility in each of the 30 intervals determined by these values. Notice, moreover, that this implies that these 30 ordinal preferences account for all the possible CRRA ordinal preferences, given the set of lotteries. Columns 3 and 7 in Table 2 describe the ordinal preferences per type, beginning with the preference of the first type and then specifying the pair(s) of alternatives that flip from the preference order of the previous type.

Step 2 (Ordered domain). Next, it is necessary to check whether the domain of menus is ordered, given the selected family of utilities. Thus, one needs to identify in each menu $j \in \mathcal{J}$ and for every type $t \in \mathcal{T}$ the alternative $x \in A_j$ that is maximal. Now, for every pair of maximal alternatives in a menu, one could check whether there is a unique $t^* \in \mathcal{T} \setminus \{T\}$ such that one alternative is preferred to the other if, and only if, $t \leq t^*$. Alternatively, one could follow Claim 1 in the characterization theorem and

²Since lotteries l_2 and l_6 involve 0 payoffs, we assume a small fixed positive background consumption.

TABLE 2. Preferences and Estimation Results

t	ω	U_t	ψ	t	ω	U_t	ψ
1	-4.148	2-6-3-7-4-8-5-9-1	0.1166	16	0.374	(7,2)	0
2	-0.518	(1,9)	0	17	0.408	(8,2)	0
3	-0.313	(3,6)	0.1009	18	0.443	(9,2)	0.0915
4	-0.083	(4,7)	0	19	0.516	(4,3),(8,7)	0.0079
5	0.065	(5,8)	0	20	0.607	(6,2)	0
6	0.154	(4,6)	0	21	0.652	(5,3),(9,7)	0
7	0.209	(1,8)	0	22	0.808	(1,3)	0
8	0.229	(3,2),(7,6)	0.0561	23	0.844	(8,3)	0
9	0.258	(5,6)	0	24	1.000	(9,3)	0.1388
10	0.262	(1,6)	0.0274	25	1.124	(5,4),(9,8)	0
11	0.273	(5,7)	0	26	1.309	(1,4)	0
12	0.339	(4,2),(8,6)	0	27	2.000	(7,3)	0.0497
13	0.342	(1,7)	0	28	2.826	(9,4)	0.0489
14	0.358	(5,2),(9,6)	0	29	4.710	(1,5)	0
15	0.363	(1,2)	0	30	∞	(8,4)	0.3622

λ : 0.2516 log-likelihood: -10.7168

the discussion of interval domains in Section 4.3, and check whether, for every menu $j \in \mathcal{J}$ and alternative $x \in A_j$, $\mathcal{T}(x, j)$ is an interval.

It is routine to check that our domain of menus is ordered, given our set of types.

Step 3 (Identifiability). Having specified $\{U_t\}_{t \in \mathcal{T}}$ and shown that the domain is ordered, we now obtain the set of fully identifiable types \mathcal{T}^I . To do this, it is enough to determine those types that are the largest maximal type for an alternative in a menu. Given the ordered domain, it is sufficient to select every type whose maximal alternative in a menu is different from that of the next type.

In our case, we have selected all the ordinal CRRA preferences for the lotteries involved in the experiment, and, since our dataset includes all the binary menus, all 30 types are identifiable.

Step 4 (Characterizing property). In the main result of the paper we show that \mathcal{T} -Monotonicity is a necessary and sufficient property for \mathcal{T} -RUMs. The analyst may be

interested in evaluating whether the property is satisfied within the sample. In Section 4.2 we provide a procedure for verifying this property in practice. If the property is fully satisfied, then the analyst can directly obtain a distribution over the type space that is consistent with the observed frequencies, as explained in the proof of Theorem 1, immediately after Claim 2.

In general, there are two possible relevant violations of \mathcal{T} -Monotonicity. The first, which involves the choice of dominated alternatives. This type of violation cannot be explained by sampling issues, and hence, if observed, the appropriate approach involves the use of the trembling version of the model, studied in Section 5. In this case, one can work with the conditional choice probabilities \bar{p} defined in Section 5, and check whether \mathcal{T} -Monotonicity is satisfied for \bar{p} . Again, if \mathcal{T} -Monotonicity holds over \bar{p} , the analyst can obtain a distribution over types that is consistent with the conditional choice data.

\mathcal{T} -Monotonicity can also be violated in other ways, perhaps due to finite sampling, as studied in Section 6. This can be evaluated by focusing on the conditional choice probabilities \bar{p} introduced in Section 5, and analyzing \mathcal{T} -Monotonicity using the procedure described above.

In our dataset, it is rather direct that \mathcal{T} -Monotonicity holds neither over the observed choice probabilities nor over the conditional ones. In the case of the former, there are several menus for which some alternatives are not maximal for any type but are nevertheless chosen in the data. In the binary menu $\{l_5, l_9\}$, for example, l_5 first-order stochastically dominates l_9 and l_9 is chosen with probability .17. In the case of the latter, consider the binary menus $\{l_4, l_7\}$ and $\{l_2, l_8\}$. The riskier lottery, l_7 (respectively, l_2), is maximal for types up to type 3 (respectively, type 16). \mathcal{T} -Monotonicity requires that the choice probability of l_7 in $\{l_4, l_7\}$ must be lower than that of l_2 in $\{l_2, l_8\}$. However, the observed choice probabilities are .38 and .22, respectively, which violates the property.³

Step 5 (Estimation). When \mathcal{T} -Monotonicity is not satisfied, an estimation exercise will find the closest distribution over types, as outlined in Section 6. This requires the selection of a particular estimator from the broad class of estimators shown in Section 6 to be strongly consistent. In addition, depending on whether dominated alternatives

³Note that these choice probabilities also apply to the conditional ones, since there are no dominated options in these menus.

are selected, the estimation may involve the adoption of the trembling version of the model of Section 5. When this is the case, the analyst needs to decide which trembling structure to be adopted from the general class of specifications used in Section 5.

For our analysis, we adopt the standard maximum likelihood estimator, and, since dominated alternatives are observed to be chosen, we implement the trembling version of the model. We choose to go with the simplest version and use a constant tremble parameter across all choice problems, and a uniform selection of non-maximal alternatives.⁴ Table 2 reports the estimated densities of the \mathcal{T} -RUMT (Columns 4 and 8). The results reveal a high degree of heterogeneity, including a very significant proportion of highly risk-averse types. The percentage of types with close to or higher than logarithmic curvature (type 24 and above) is 60%, with more than a third of all decisions belonging to the highest type, type 30, which exhibits risk-aversion levels as high as 4.7 and above. The results also show that a relevant proportion of the decisions reflect highly risk-seeking attitudes (22%), and even extreme risk-seeking, (12% of all decisions correspond to the lowest type, type 1, that is, risk-aversion coefficients below -4.8). The estimated probability of tremble is .25, thus confirming the behavioral relevance of non-maximal alternatives.

Step 6 (Statistical testing). The final step is the statistical testing of the model, as elaborated in Section 6. This involves comparing the observed choice frequencies with the predictions given by the estimated model, using the Pearson statistic. The analysis can be performed at the menu level with $|A_j| - 1$ degrees of freedom, or at the aggregate level with $\sum_{j \in \mathcal{J}} |A_j| - J$ degrees of freedom.

In the dataset, based on the Pearson statistic per each of the 246 menus, we are unable to reject the null hypothesis of equality between observed and predicted choices in 72% (85%) of the menus at the 5% (1%) significance level.⁵ Note that the degrees of freedom in the individual menus are 1, 2 and 4 for the menus with 2, 3 and 5 alternatives. The aggregated Pearson statistic across menus enables us to reject the null hypothesis of equality at conventional significance levels. In this case, the degrees of freedom are 708.

⁴The data and estimation programs are available for use on our websites.

⁵Recall that the purpose of this exercise is to illustrate the applicability of the model, and for this purpose we chose the simplest possible implementation, involving a representative agent approach with CRRA expected utilities. Naturally, accounting for inter- and intra-personal variability and using other families of utility functions may improve these results.

REFERENCES

- [1] Apesteguia, J. and M.A. Ballester (2021). “Separating Predicted Randomness from Residual Behavior,” *Journal of the European Economic Association*, 19(2):1041–1076.