Order-Dependent Choice and Non-Binary Preferences

T.C.A Madhav Raghavan

January 5, 2011

Abstract

In this note we consider a model of successive choice based on complete and transitive preferences. When we weaken the conditions of completeness, we find that choice becomes order-dependent. We introduce a *tie-breaker rule* to solve the problem of incompleteness, but retain order-dependence. The combined effect of order-dependence and this tie-breaker rule will lead us to a generalisation of *non-binary preferences* (i.e., composite preferences over sets of more than two elements). This generalisation is probabilistic in nature, and in many cases matches common intuition.

1 Successive Choice

The standard model of successive choice is presented as example 5 in Rubinstein and Salant (2006). Algorithmically:

> Successive choice: The primitive of this choice procedure is a binary relation R over X, where xRy is interpreted as "x rejects y". Given a list L = $(a_1, ..., a_K)$, the decision maker stores a_1 in a "register"; at stage t of the computation, $1 \leq t < K$, the decision maker replaces the register value y with a_{t+1} if $a_{t+1}Ry$. When the list ends the decision maker chooses the alternative in the register.

Claim **1** *This choice function (let's called it D) satisfies the axioms of Partition Independence (PI), List Independence of Irrelevant Alternatives (LIIA), and Order Invariance (OI).*

Proof: To show that D satisfies PI, we have to show that there exists a $\succsim$ (weak ordering) and a unique δ (priority indicator function) such that D = $D_{\succsim,\delta}$. Firstly, we note that $\succsim$ is determined by the binary relation R. This follows from the facts that R is complete, and since 'Rejects' is a strict relation, that R is also transitive. We can now define the strict component of $\succsim$ as equivalent to R, and set indifference everywhere else.
Now, let δ be the priority indicator function that assigns label 1 to every indifference class.

It follows from these two that D = $D_{\succsim,\delta}$, and so D satisfies PI. It also, therefore, satisfies LIIA.
It can also be easily shown that D chooses the maximal element regardless of the order of the list, and so it satisfies Order Invariance (OI). ■

2 Extensions

In what follows, we will weaken some of the conditions on this successive choice function. This will make it order-dependent.
In particular, we will examine the following two notions:

- The completeness of the underlying binary relation.
- The priority indicator function, δ.

Incomplete Preferences

A complete preference ordering is one that is defined for every pair in the choice set X. But it is possible that some alternatives simply cannot be compared - they are like apples and oranges.
Note that this is very different from the notion of indifference, which requires an active comparison and evaluation. The idea we are after is that of pairs that cannot be compared at all, not that they are compared and found to be equivalent. Importantly, decisions made over these incomparable pairs are made either randomly, or according to a heuristic, with consequences for the choice function and the chosen alternatives.
For example, suppose a_1Pa_2 and a_2Pa_3. It is not possible for a_1 to be indifferent to a_3 without violating transitivity. But if a_1 and a_3 cannot be compared at all, then a choice of a_1 can still be made by a heuristic that favours a_1. We shall explore such heuristics, and this violation of transitivity, in the model below.
Formally, there is a finite set X of alternatives. $\mathcal{M}$ is a (possibly infinite) set of criteria or motives. To each x_i from X, there is associated a finite subset M_i of $\mathcal{M}$, which is the set of relevant criteria. Comparable alternatives are defined as any pair (x_i, x_j) from X such that $M_i \cap M_j \neq \phi$. That is, there is a non-empty set of criteria that is common to alternatives x_i and x_j. We assume that this is a basis for a preference expressed over them. Incomparable choices, on the other hand, are all pairs (x_k, x_l) from X such that $M_k \cap M_l = \phi$, and so in this case there is no basis for comparison. [1]

[1] This terminology is borrowed from Dietrich and List (2010), who look at reason-based choice. Using two simple axioms, they derive a relationship between preference over objects and suitably defined underlying sets of reasons relevant to those objects. This is done via a 'weighing' relation between sets of reasons. However, they do not approach the problem of non-comparability. We shall explore these concepts in further

Consider an example. Let X $= \{a_1, a_2, a_3, a_4\}$. Suppose $a_3 P a_4$, $a_4 P a_1$, $a_1 P a_2$, and suppose that we cannot say anything about any of the other pairs. We denote this by $a_1 T a_3$, $a_2 T a_3$, $a_2 T a_4$, respectively, the operator T signifying incomparability (See Table 1). Note that a_3 is indirectly the best alternative, but because preferences are not complete, it is not the direct best.

Table 1: An Incomplete Preference Ordering

Strict Preference			Incomparable
a_3			
$\downarrow$			$a_1 \leftrightarrow a_3$
a_4	a_4		
	$\downarrow$		$a_2 \leftrightarrow a_3$
	a_1	a_1	
		$\downarrow$	$a_2 \leftrightarrow a_4$
		a_2	

Suppose that the choices are encountered in the following order: $a_2 a_1 a_3 a_4$
At the first stage, a_2 is compared with a_1, and since $a_1 P a_2$, a_1 is carried forward. At this point, since a_1 and a_3 are incomparable, we make an arbitrary choice between them in order to proceed. Say we pick at random, and choose a_1. The alternative a_3 is then eliminated, and in the last stage, because $a_4 P a_1$, a_4 is the eventual winner. Note that this happens even though a_3 is indirectly the best element in X. Of course, this is the consequence of the tie-breaker, which in this case was arbitrary. But it illustrates how, in the absence of complete preferences, a tie-breaker could result in picking alternatives that are not always maximal.

The Tie-Breaker Rule

The tie-breaker need not always be arbitrary. Indeed, the δ used by Rubinstein and Salant is one such tie-breaker rule, and it is well defined (the assigning of the label 'first' or 'last' uniformly within an indifference class). Alternatively, it is possible that the incompleteness in preferences is gradually eliminated as choices are made. Suppose $a_1 P a_2$ and $a_2 P a_3$, but a_1 and a_3 are incomparable, and the order is $a_1 a_2 a_3$. A learning agent might be able to deduce, in the second round, that a_1 should be preferred to a_3. This is because by then all information has been presented to him or her - in particular, that $a_1 P a_2$ and $a_2 P a_3$. We will consider this possibility in later work. But for now, as we saw above, the presence of such a TBR does allow the possibility of picking an alternative that is not preference maximal.
The TBR is used every time an incompatible pair is encountered and, depending on the preferences, sometimes more than once in a list. An example of a commonly used TBR is

work.

the heuristic that says 'pick the status quo'. This applies to the cases where alternatives are sequentially received over time, such that comparisons are made between the current alternative and a 'new' one.

3 An Extended Choice Model

So we define a successive choice function as follows: Given a list, start with the first two elements. If a preference over the pair exists, then pick the more preferred element. When an incomparable pair is encountered, use a tie-breaker rule to pick one of them, and proceed. When the last element has been compared, the alternative that remains is the choice. This is very similar to the Rubinstein and Salant example, except for the additional conditions that the underlying preferences are not complete, and that a TBR is employed to resolve incomparabilities.

Claim **2** *This choice function D^* does not satisfy PI or LIIA, or even OI.*

Proof: To show that it does not satisfy PI, we show that $D^*(< L_1, L_2 >) \neq D^*(D^*(L_1), D^*(L_2))$. Consider the same example used above. Let the TBR be 'pick the status quo,' i.e., the first element. Now, $C\{a_1, a_2, a_3, a_4\} = C\{a_1, a_3, a_4\} = C\{a_1, a_4\} = a_4$
But, suppose we partition the list down the middle, then $C\{a_1, a_2\} = a_1$ and $C\{a_3, a_4\} = a_3$, and thus $C\{a_1, a_3\} = a_1$ is the final choice.
Since a_1 is different from a_4, PI is not satisfied. The choice from the best elements of the partitioned lists did not match the choice from the full list. It follows that the choice function does not satisfy LIIA either.
Also, since when we take a different order of the elements, say $C\{a_4, a_1, a_2, a_3\} = a_3 \neq a_4$, the choice function also does not satisfy OI. ■

4 Virtual Orderings

The choice function described above is order dependent. This is not a problem as long as the list is known, because we can work through the preference rules to evaluate the final choice. But in the case that the actual order of the list is not observable (what Rubinstein and Salant call *virtual orderings*), then it is not possible to definitely establish which alternative is the winner. Note that this is a consequence of the incomplete nature of these preferences. If they were complete, and transitive, then this choice function would become the standard successive choice model, which satisfies OI. And so the best element, i.e., a_3, would be chosen every time.

Sum-Over-Paths / Probability Approach

So in order to get a measure of which alternatives tend to be selected, we consider all permutations of the set $\{a_1, a_2, a_3, a_4\}$ as lists that could be encountered. It is assumed that there is no way to know beforehand which list will appear, and so that each list is equally likely. There are 24 such permutations. We proceed pair-wise through each list, using the underlying preferences and the TBR where appropriate. The final selected alternatives in the various scenarios are shown in Table 1.

Representing the relative frequencies of wins (out of 24) for each alternative, we get:

Table 2: Incomplete Preferences

Alternative	Frequency	Percentage
a_1	3	12.5%
a_2	0	0.0%
a_3	16	66.7%
a_4	5	20.8%

It is clear that the 'best' alternative a_3 does not win every time. This is because it gets eliminated by the TBR whenever it appears after a_1 or a_2, to which is it incomparable. Also, the indirect worst element never gets picked. To fully highlight these observations, let us consider the two extreme cases. First, suppose preferences are complete. Then, as argued above, a_3 should win regardless of the order. Indeed, this is what happens (see Table 3).

Table 3: Complete Preferences

Alternative	Frequency	Percentage
a_1	0	0.0%
a_2	0	0.0%
a_3	24	100.0%
a_4	0	0.0%

On the other hand, suppose that all choices are incomparable, then we would expect that any of the alternatives is as likely to be chosen as any other. That is what we find (see Table 4).

These present the two limiting cases of no- and full-information. In general, when preferences are incomplete, we get an intermediate result, which is an order-dependent likelihood of picking the various options, and interior probabilities such as in Table 2.

5 The Tie-breaker Rule and its Properties:

Table 4: Empty Preferences

Alternative	Frequency	Percentage
a_1	6	25.0%
a_2	6	25.0%
a_3	6	25.0%
a_4	6	25.0%

Preliminary Observations

We assume that the strict preference relation (PR) is transitive, wherever defined.

• When presented with two incomparable alternatives x and y, if the TBR were to always pick x over y, then this would be equivalent to having a preference of x over y. So to qualify as a TBR, the rule is essentially a randomisation between x and y based on whether the choices are presented as xy or as yx. There must be some chance of picking either.

• Rubinstein and Salant's δ function is a 50-50 randomisation. Always picking the status quo is another. So is always picking the new alternative. Further work may look at what a general 'mixed strategy' would mean in this context.

• A TBR that says 'pick the status quo' minimises risk. This is because a chosen alternative is replaced only when an explicitly superior alternative emerges. It produces a zero probability of the worst alternative being chosen.

• On the other hand, a TBR that says 'pick the new alternative' (pick 'last') is more experimental. It raises the probability of picking the worst outcomes, and thus reduces the probability of picking the best one. See Table 5.

Table 5: TBR 'pick the latter alternative', incomplete preferences

Alternative	Frequency	Percentage
a_1	7	29.2%
a_2	4	16.7%
a_3	8	33.3%
a_4	5	20.8%

• Since a_2 and a_3, and a_2 and a_4 are incomparable, the TBR picks a_2 whenever it follows a_3 or a_4. The non-zero probability corresponding to a_2 in the table arises out of the fact that a_2 appears last on some permutations, and thus becomes the final choice.

• Any 'mixed' TBR consisting of these two rules assigned to different incomparable pairs also produces non-zero probabilities for the 'worst' alternative.

• The combined effect of the PR and the TBR is to provide a probability distribution for the set of alternatives. This probability distribution can be viewed as representing non-binary preferences over the set. That is, these preferences are of a probabilistic nature. They represent the relative likelihoods of picking each element out of a set.

• The probability distributions are sensitive to the arrangements of the PR and the TBR. Complete preferences, for example, result in a 100 per cent probability of the maximal element being chosen. Empty preferences result in equiprobable distributions.

• Being probabilities, they are inter-related for all elements of the set, and thus might not satisfy many of the classical binary independence axioms. Classical IIA, for example, fails to be satisfied.

• Since preferences are not complete, our primitive PR here is defined on a subset of $\mathcal{X}$, the set of all pairs (x,y) from X. Each pairwise comparison is in a sense delinked, since we cannot use normal transitivity. The question of multiple agents may be addressed by the set-theoretic approach of Dietrich and List.

6 Related Papers

Salant (2003) uses a model of successive choice and incorporate memory. Depending on how many previous rounds are remembered, the information set present in any round changes. Including this possibility might allow us to consider learning in this model, removing incomparabilities and thereby limiting the use of the TBR.

Fishburn and Rubinstein (1986) use attributes to partition sets into equivalence classes, and then define aggregators over these partitions. Here we are aggregating probabilities over permutations to produce distributions over the set, and are thus using a class of aggregator. The notions of consistent and conjunctive aggregators defined there might be of use.

Dietrich and List (2010) is discussed in Footnote 2.

References

Dietrich, F. and C. List (2010): "A Reason-Based Theory of Rational Choice," Working paper. http://www.franzdietrich.net/Papers/DietrichList-ReasonBasedRationalChoice.pdf.

Fishburn, P. and A. Rubinstein (1986): "Aggregation of Equivalence Relations," *Journal of Classification*, 3.

Rubinstein, A. and Y. Salant (2006): "A Model of Choice from Lists," *Theoretical Economics*, 1.

Salant, Y. (2003): "Limited computational resources favour rationalizability," Discussion paper 320, Hebrew University of Jerusalem. http://ratio.huji.ac.il/dp/dp320.pdf.