On the Likelihood of Cyclic Comparisons[*]

Ariel Rubinstein[†]   Uzi Segal[‡]

November 22, 2011

Abstract

One problem caused by cycles of choice functions is indecisiveness
— decision makers will be paralyzed when they face choice sets with
more than two options. We investigate the procedure of "random sam-
pling" where the alternatives are random variables. When comparing
any two alternatives, the decision maker samples each of the alterna-
tives once and ranks them according to the comparison between the
two realizations. We show that while this procedure may lead to vio-
lations of transitivity, the probability of such cycles is bounded from
above by $\frac{8}{27}$. Even lower bounds are obtained for some other related
procedures.

KEYWORDS: Transitivity, preference formation, the paradox of nontransitive
dice

[†]University of Tel Aviv Cafes, University of Tel Aviv, and New York University
[‡]Department of Economics, Boston College (`segalu@bc.edu`)

1 Introduction

A standard argument that is used to justify the transitivity assumption in decision theory is indecisiveness — a decision maker with cyclical preferences on a triple of alternatives will not be able to choose from this set. Suppose that $A \succ B$, $B \succ C$, and $C \succ A$. If the decision maker has to choose from the set $\{A, B, C\}$ he will be freezed: for each alternative he may choose he will find a better one. The freeze may bring the decision maker to re-evaluate his preferences and probably to change them so that the cycles will be eliminated.

This discussion presupposes the existence of a deterministic binary relation, that is, a relation that presumably depends on the decision maker's psychological and physiological evaluation of the different alternatives. But what if the decision maker is using a random procedure to determine his choice? Such a procedure may yield indecisiveness and the higher is its probability the more likely it is that the decision maker will change the procedure so that cycles will not emerge.

In this paper we discuss the indecisiveness argument in the context of random preferences. We focus on a nondeterministic procedure of preference formation which we call *Random Sampling procedure*. When comparing two lotteries, the decision maker samples once from each lottery and ranks them according to the two realizations. (This concept is related to the S-1 procedure proposed in Osborne and Rubinstein (1998)).

The main message of our paper is that when applied to random preferences, the scope of the indecisiveness argument is limited. Whereas the argument is always applicable for deterministic procedures which yield cycles, the random procedures which we study would be less vulnerable to the indecisiveness argument. Our formal analysis provides a characterization of the upper bound on the probability that the random procedure we study yields indecisiveness and shows that this bound is quite low.

Our first result refers to the case where when choosing from three random variables the decision maker independently compares each pair of them. He starts by comparing some alternatives A and B, continues into comparing B and C, and then finally compares A and C. In each of the three stages he draws new samples from the relevant pair of random variables and does not use the values he observed before. We find that the bound on the probability that this random sampling procedure yields a cycle is 8/27 (Claim 1). This is somewhat lower than 1/3, which we show in Claim 3 to be the bound for

2

the Block and Marschak's (1960) *Random Ordering* procedure. According to this procedure, which is different from those discussed by us, the decision maker has in mind a set of orderings. When comparing two alternatives, he randomly samples one of the orderings and ranks the two alternatives according to that ordering.[1]

We then turn to the case where the decision maker activates the three comparisons in a pre-determined order, starting by comparing the alternatives A and B and continues with comparing B and C and then C and A, but unlike the random sampling procedure, he partially recalls past observations. In the second comparison the decision maker remembers the value of the observed sampling from B which he got in the first comparison, and in the last comparison he recalls the value of C which was used in the second comparison. However, in the third comparison he samples afresh from A. In other words, this decision maker remembers the outcomes of the last comparisons, but not what he has seen two stages ago. We show that whatever are the three alternatives and whatever is the order of comparisons, the bound on the probability that this procedure yields a cycle is $1/4$ (Claim 5).

In the last part of the paper we show that the probability of indecisiveness can be reduced significantly if someone else (for example, an agent who wants the consumer to make a quick choice) can control the order at which the decision maker compares the alternatives. We show that the upper bound on the probability of a cycle is reduced to $1/16$ for the case of binary lotteries (Claim 6) and to 0.091 for lotteries with at most three outcomes. Moreover, if the external agent's choice of order could depend on the realizations in the first comparison, then he can eliminate cycles altogether for binary lotteries (Claim 7) and he can reduce the bound on the probability of a cycle for lotteries with three outcomes to $1/32$.

Thus, we show that nondeterministic procedures of choice, applied to three alternatives, yield transitivity with a fairly high probability. Therefore, the mere fact that choice is "almost" well behaved and only a small number of cycles is observed does not necessarily prove that decision makers are using deterministic transitive preferences (while making occasional mistakes). Such behavior can also emerge when choice is based on some variants of random sampling where decision makers do not employ preference relations and certainly do not change them to avoid indecisiveness.

---

[1]For a recent discussion of how the random ordering procedure can explain data which exhibits intransitivity, see Regenwetter, Dana, and Davis-Stober (2011).

3

We suggest two conclusions: First, one has to be careful in interpreting the emerged transitive relation as exhibiting the welfare of the decision maker. (This point is related to the recent literature on eliciting notions of welfare from data which is conflicting with rationality. See Bernheim and Rangelk (2009) and Rubinstein and Salant (2011)). Secondly, indecisiveness emerging while applying random procedures may push decision makers to re-evaluate the procedures in order to avoid indecisiveness. But the pressure to change under some of these random procedures is dramatically less severe than when deterministic procedures are used.

2 Random Sampling

The main procedure we discuss in this paper is *random sampling*: To compare two random variables the decision maker draws a fresh sample from each and ranks them according to the sampled values.

Throughout the paper, all triples of random variables have finite and disjoint supports. Denote by $s(A)$ the support of the lottery A and by $\Pr(A > B)$ the probability that the realization of A is higher than the realization of B. By the disjoint supports assumption, $\Pr(A > B) + \Pr(B > A) = 1$. Let $\Pi(A, B, C)$ be the probability of a cycle being created by the decision maker's procedure. Applied to the random sampling procedure we have:

$$\begin{aligned} \Pi(A, B, C) & = \Pr(A > B) \Pr(B > C) \Pr(C > A) + \\ & \quad \Pr(A > C) \Pr(C > B) \Pr(B > A) \end{aligned}$$

Claim 1 *The maximal probability that the procedure of random sampling yields a cycle is $\frac{8}{27}$.*

Proof: Consider the three random variables presented in the following table:

Value	A	B	C
4	1/3		
3		2/3	
2			1
1	2/3		
0		1/3	

4

In this case, $\Pr(A > B) = \frac{5}{9}$, $\Pr(B > C) = \frac{2}{3}$, $\Pr(C > A) = \frac{2}{3}$ and the probability of a cycle $\Pi(A, B, C)$ is $\frac{20}{81} + \frac{4}{81} = \frac{8}{27}$.

In order to prove that this is the upper bound, let $x_1 > x_2 > \ldots > x_n$ be the values in the supports of the three random variables A, B and C. Denote by $X_i \in \{A, B, C\}$ the random variable that contains x_i in its support. Let $\pi_i = \Pr(X_i = x_i) > 0$.

First, we assume without loss of generality that for all i, $X_i \neq X_{i+1}$; otherwise, if $X_i = X_{i+1} = A$, let A' be the random variable which differs from A by $\Pr(A' = x_i) = \pi_i + \pi_{i+1}$ and $\Pr(A' = x_{i+1}) = 0$. Then, $\Pi(A', B, C) = \Pi(A, B, C)$.

Next, assume that for some i, $X_i = X_{i+2} \neq X_{i+1}$ (without loss of generality $X_i = A$ and $X_{i+1} = B$). Then we can (weakly) increase the probability of a cycle by replacing A with A_ε, a random variable which differs from A by either moving a probability mass $\varepsilon > 0$ from x_i to x_{i+2} or from x_{i+2} to x_i. Clearly, $\Pr(C > A_\varepsilon) = \Pr(C > A)$ and $\Pr(A_\varepsilon > B)$ is linear in ε. Since

$$\begin{aligned}\Pi(A_\varepsilon, B, C) &= \Pr(A_\varepsilon > B)[\Pr(B > C)\Pr(C > A)] + \\ &\quad (1 - \Pr(A_\varepsilon > B))\Pr(A > C)\Pr(C > B)\end{aligned}$$

shifting probability mass from x_{i+2} to x_i or the other way around (according to the sign of $\Pr(B > C)\Pr(C > A) - \Pr(A > C)\Pr(C > B)$) will (weakly) increase the probability of a cycle.

Thus, without loss of generality we can assume that the sequence $\{X_i\}$ is of the form $\ldots A, B, C, A, B, C, \ldots$ ending with $X_{n-2} = A$, $X_{n-1} = B$ and $X_n = C$.

Next we show that if the three random variables (A, B, C) maximize Π and if $n > 6$, then there is a triple of random variables that maximizes Π with less than n values in their joint supports. First note that:

$$\begin{aligned}\Pi(A, B, C) &= \Pr(C > A)[\Pr(B > C)\Pr(A > B) - \\ &\quad \Pr(B > A)\Pr(C > B)] + \Pr(B > A)\Pr(C > B)\end{aligned}$$

Changing C does not affect $\Pr(B > A)$. Consider the set of all C' with a support that is a subset of C such that $\Pr(B > C') = \Pr(B > C)$. For all such C', denote by γ_i the probability that C' yields the outcome x_i. This is the set of all vectors $(\gamma_i)_{x_i \in s(C)}$ such that $\gamma_i \geqslant 0$ for all i and the following two linear equations hold:

$$\sum_{x_i \in s(C)} \gamma_i \;\; = \;\; 1$$

$$\sum_{x_i \in s(C)} \gamma_i \;\; \times \sum_{j<i \text{ and } x_j \in s(B)} \pi_j \;\; = \;\; \Pr(B > C)$$

Since $n > 6$ and $X_n = C$, there are at least $m \geqslant 3$ points in the support of C. The set C' is therefore non empty and is given by the intersection of R^m_{++} and the above two $m-1$ dimensional hyperplanes. The two hyperplanes intersect at C, thus the set is the intersection of R^m_{++} and a linear space of dimension $m - 2 > 0$.

Replacing C with C' will increase the probability of a cycle if $\Pr(B > C) \Pr(A > B) - \Pr(B > A) \Pr(C > B)$ and $\Pr(C' > A) - \Pr(C > A)$ have the same sign. The expression

$$\Pr(C' > A) = \sum_{x_i \in s(C)\gamma_i \times \sum_{j>i \text{ and } x_j \in s(A)} \pi_j}$$

is a linear function in $(\gamma_i)_{x_i \in s(C)}$. Therefore, we can (weakly) increase $\Pr(C' > A)$ by moving in some direction until we reach the boundary where $\gamma_i = 0$ for some x_i in the support of C.

We can therefore narrow down our attention to the sequence of variables $(X_i)_{i=1\ldots6}$ which is of the form A, B, C, A, B, C. Denote by α, β, γ the probabilities that the variables A, B and C obtain the highest prize in their supports. Then,

$$\begin{aligned} \Pi(A, B, C) &= (1 - \beta + \alpha\beta)(1 - \gamma + \beta\gamma)(\gamma - \alpha\gamma) + \\ &\quad (\beta - \alpha\beta)(\gamma - \beta\gamma)(1 - \gamma + \alpha\gamma) = \\ &= \gamma^2(1 - \alpha)(\beta - 1) + \gamma(1 - \alpha)(\alpha\beta - \beta^2 + 1) \end{aligned}$$

Assuming that both $1 > \alpha$ and $\beta > 0$, the last expression is strictly increasing in γ within the interval $[0, 1]$. Thus, it attains its maximum at $\gamma = 1$. We conclude that in the optimum, one of the three variables must be degenerate and without loss of generality the sequence $(X_i)_{i=1\ldots5} = (B, C, A, B, C)$. Then,

$$\Pi = \gamma^2(\beta - 1) + \gamma(-\beta^2 + 1) = \gamma^2\beta - \gamma^2 - \gamma\beta^2 + \gamma$$

This expression has a unique maximum point at $\beta = \frac{1}{3}$ and $\gamma = \frac{2}{3}$ and a maximization value of $\Pi = \frac{8}{27}$. ∎

In Claim 1 we obtained the upper bound on the probability that the procedure of random realizations yields one of the two possible cycles $A \succ C \succ B \succ A$ or $A \succ B \succ C \succ A$. Claim 2 identifies the highest probability that the procedure yields a particular cycle.

Claim 2 *The maximal probability that the procedure of random sampling yields a particular cycle is $\frac{1}{4}$.*

Proof: Let A to be the random variable that receives the values 3 or 0 with equal probabilities. Let $B \equiv 2$ and let $C \equiv 1$. Then, $\Pr(A > B)\Pr(B > C)\Pr(C > A) = \frac{1}{4}$.

Now, let A, B and C be random variables with $n > 6$ values in their joint support. We show that there is another triple of random variables that yields the cycle $A \succ B \succ C \succ A$ with at least as high a probability and with less than n values in their joint support.

As in the proof of Claim 1 we can easily reduce n if $X_i = X_{i+1}$. If $X_i = A$, then $X_{i+1} = B$ (and similarly if $X_i = B$, then $X_{i+1} = C$ and if $X_i = C$, then $X_{i+1} = A$) since if $X_{i+1} = C$ we can increase $\Pr(C > A)$ without affecting $\Pr(A > B)$ and $\Pr(B > C)$ by shifting the probability mass in A from x_i to x_{i+1} and shifting the probability mass in C from x_{i+1} to x_i.

Without loss of generality, let $X_n = C$. As in the proof of Claim 1, we can modify C to C' such that $\Pr(B > C') = \Pr(B > C)$ and increase (at least weakly) $\Pr(C > A)$ until we reach the boundary where probability 0 is assigned to one of the outcomes in the support of C. Thus, we can assume that there is a triple which maximizes the probability of the cycle with $n \leqslant 6$ points in their joint support. Let α, β, and γ be the the probabilities that the variables $X_1 = A$, $X_2 = B$, and $X_3 = C$ obtain the highest prize in their supports. Using the inequality of geometrical and arithmetic averages

we obtain that

$$
\begin{aligned}
&\Pr(A > B)\Pr(B > C)\Pr(C > A) = \\
&(1 - \beta + \alpha\beta)(1 - \gamma + \beta\gamma)(\gamma(1 - \alpha)) = \\
&[(1 - \beta + \alpha\beta)\gamma][(1 - \gamma + \beta\gamma)(1 - \alpha)] \leqslant \\
&\left(\frac{(1 - \beta + \alpha\beta)\gamma + (1 - \gamma + \beta\gamma)(1 - \alpha)}{2}\right)^2 = \\
&\frac{(1 - \alpha + \alpha\gamma)^2}{4} \leqslant \frac{1}{4} \qquad\qquad\qquad \blacksquare
\end{aligned}
$$

Comments:

(a) The problem we dealt with in this section is related to the so-called "paradox of nontransitive dice" (see Gardner (1970) who credits it to the statistician Bradley Efrom). This "paradox" involves three independent random variables: A, B, and C, where $Pr(A > B)$, $Pr(B > C)$, and $Pr(C > A)$ all exceed 0.5.[2] Savage (1994) further proved that $\max_{A,B,C} \min\{\Pr(A > B), \Pr(B > C), \Pr(C > A)\} = (\sqrt{5} - 1)/2$.

(b) It follows from Claims 1 and 2 that for every three distributions F,G, and H with a bounded domain and which do not have an atom in the same point:

$$
\int FdG \int GdH \int HdF + \int FdH \int HdG \int GdF \leqslant \frac{8}{27}
$$

and

$$
\int FdG \int GdH \int HdF \leqslant \frac{1}{4}
$$

(c) When a decision maker applies the ordering sample procedure to a set of size n, the maximum probability that his ranking is acyclic goes to zero as the number of alternatives increases to infinity. To see it consider n random variables which are uniform on the interval $[0, 1]$ (and obviously could be approximated by random variables with finite and disjoint supports). For any two of these random variables, the probability that the realization of one is higher than of the other is $\frac{1}{2}$. By Moon and Moser (1962), the probability that the realized tournament is irreducible (i.e., there are no two non-empty disjoint sets such that every node in one set "beats" every node in the other)

---

[2]See http://singingbanana.com/dice/article.htm for an entertaining demonstration of this setup.

goes to 1 as $n \to \infty$. By Moon (1966), a tournament with n nodes has a cycle of length n (and therefore is not acyclic) if and only if it is irreducible. Thus, the probability that the decision maker's comparisons of n uniform random variables yields a cycle of size n goes to 1 as $n \to \infty$.

3 The Random Ordering Procedure

In the random ordering procedure (Block and Marschak (1960)) the decision maker is characterized by π, a probability measure over the six orderings of the three alternatives A, B, and C. When comparing any pair of alternatives, the decision maker draws an ordering that will determine his ranking of these alternatives. Thus, he might apply different orderings in ranking two different pairs of alternatives. In this section we show that the bounds we obtained in the previous section are lower than the bounds on the probability of a cycle in the random ordering procedure.

Claim 3 *The maximal probability that the random ordering procedure yields a cycle is $\frac{1}{3}$.*

Proof: Consider π to be a probability measure on the orderings that assigns equal probabilities to the three orderings $A \succ_1 B \succ_1 C$, $B \succ_2 C \succ_2 A$ and $C \succ_3 A \succ_3 B$. Then, $\Pr(A \succ B) = \Pr(B \succ C) = \Pr(C \succ A) = \frac{2}{3}$ and the probability of a cycle is $\frac{8}{27} + \frac{1}{27} = \frac{1}{3}$.

To see that $\frac{1}{3}$ is indeed the bound, note that by the inequality of arithmetic and geometric means:

$$
\begin{aligned}
\Pi(A, B, C) \;=\; & \Pr(A \succ B)\Pr(B \succ C)\Pr(C \succ A) + \\
& \Pr(A \succ C)\Pr(C \succ B)\Pr(B \succ A) \leqslant \\
& \Pr(A \succ B) + \Pr(B \succ C) + \Pr(C \succ A)]^3/27 + \\
& [\Pr(A \succ C) + \Pr(C \succ B) + \Pr(B \succ A)]^3/27
\end{aligned}
$$

Since every ordering must satisfy at least one and at most two of $A \succ B$, $B \succ C$ and $C \succ A$, we obtain: $1 \leqslant [\Pr(A \succ B) + \Pr(B \succ C) + \Pr(C \succ A)] \leqslant 2$. The function $x^3 + (3-x)^3$ is convex in the interval $[1, 2]$ and obtains its maximum at $x = 1$ and $x = 2$. Thus $\Pi(A, B, C) \leqslant \frac{1}{27} + \frac{8}{27} = \frac{1}{3}$. ∎

Claim 4 *The maximal probability that the procedure of random ordering yields a particular cycle is $\frac{8}{27}$.*

9

Proof: The above example attains the bound. To prove that the bound is $\frac{8}{27}$, note that $\Pr(A \succ B)\Pr(B \succ C)\Pr(C \succ A) \leqslant \Pr(A \succ B) + \Pr(B \succ C) + \Pr(C \succ A)]^3/27$. The function x^3 in the interval $[1,2]$ attains the maximum at 2 and thus the inequality follows. ∎

Note that the above example is the only one in which the probability of a cycle is $\frac{8}{27}$. To see this, count the six orderings: $A \succ_1 B \succ_1 C$, $B \succ_2 C \succ_2 A$, $C \succ_3 A \succ_3 B$, $A \succ_4 C \succ_4 B$, and $B \succ_5 A \succ_5 C$, $C \succ_6 B \succ_6 A$. Denote by π_i the probability of $\succ_i$. Then, $\Pr(A \succ B)\Pr(B \succ C)\Pr(C \succ A) = (\pi_1 + \pi_3 + \pi_4)(\pi_1 + \pi_2 + \pi_5)(\pi_2 + \pi_3 + \pi_6)$. The maximum is attained only when $\pi_4 = \pi_5 = \pi_6 = 0$ and $\pi_1 = \pi_2 = \pi_3 = \frac{1}{3}$.

4 The Random Sampling Procedure with Partial Recall

In the procedure discussed in Section 2 each comparison is done independently of the other two comparisons. A decision maker who compares first A and B and moves to compare B and C does not recall the previous value of B. Thus the existence of a cycle did not depend on the order by which the comparisons were done. In contrast, in this section we assume that the decision maker carries out the comparisons sequentially in three stages and at each stage he remembers the realizations of the previous stage, but not those of two stages earlier. In other words, he applies the procedure of *Random sampling with partial recall*. It is applied to the sequence of three lotteries (A, B, C) in the following way:

(i) Compare A and B by sampling each once.

(ii) Compare B and C by sampling C once and compare the outcome with that of the previous-stage sampling of B.

(iii) Compare C and A by sampling A again and compare the outcome with that of the previous-stage sampling of C.

The probability that the procedure yields a cycle is

$$\Pi(A, B, C) = \Pr(A_1 > B > C > A_2) + \Pr(A_2 > C > B > A_1)$$

10

where A_1 and A_2 are copies of A, i.e., they are i.i.d and distributed like A. Note that $\Pi(A, B, C)$ might differ from $\Pi(B, A, C)$ but that $\Pi(A, B, C) = \Pi(A, C, B)$.

Claim 5 *The maximal probability that the random sampling procedure with partial recall yields a cycle is $\frac{1}{4}$.*

Proof: Even though we use here a different procedure, the probability of a cycle is $\frac{1}{4}$ for the same triple of variables used at the beginning of the proof of Claim 2. To see that $\frac{1}{4}$ is the bound, denote by Π_b the probability of a cycle given that the value of B is b:

$$\begin{aligned}
\Pi_b \;=\; & \Pr(A_1 > b > C > A_2) + \Pr(A_2 > C > b > A_1) \leqslant \\
& \Pr(A_1 > b > C)\Pr(b > A_2) + \Pr(A_2 > b)\Pr(C > b > A_1) = \\
& \Pr(A > b)\Pr(b > C)\Pr(b > A) + \Pr(A > b)\Pr(C > b)\Pr(b > A) = \\
& \Pr(b > A)\Pr(A > b)[\Pr(b > C) + \Pr(C > b)] \leqslant \frac{1}{4}
\end{aligned}$$

Since $\Pi_b \leqslant \frac{1}{4}$ for every possible realization of B, $\Pi(A, B, C) \leqslant \frac{1}{4}$ as well. $\blacksquare$

Imagine now that the order in which the alternatives are presented to the decision maker is determined by a "master of ceremonies" (MC) who wants the decision maker having a clear ordering of the alternatives. Let $V(A, B, C) = \min\{\Pi(A, B, C), \Pi(B, C, A), \Pi(C, A, B)\}$ be the probability of a cycle given that the MC chooses the order of the comparisons of the three variables A, B and C in order to minimize the probability of the cycle. In the example used in the above proof $\Pi(A, B, C) = \frac{1}{4}$ but $\Pi(B, C, A) = 0$ and thus $V(A, B, C) = 0$. On the other hand, if A, B, C are uniformly distributed over $[0, 1]$ then $V(A, B, C) = \Pi(A, B, C) = \frac{1}{12}$ (each ordering of four identical random variables has the same probability of $\frac{1}{24}$ and therefore $\Pr(A_1 > B > C > A_2) + \Pr(A_2 > C > B > A_1) = \frac{1}{12}$). We succeeded to find the bound on V for only a limited family of random variables.

Claim 6 *The maximal $V(A, B, C)$ for three binary random variables is $\frac{1}{16}$.*

Proof: First note that for the following three variables $V(A, B, C) = \frac{1}{16}$.

Value	A	B	C
5	1/2		
4		1/2	
3			1/2
2	1/2		
1		1/2	
0			1/2

If the three variables are such that between the two values of one of the lotteries, say A, there are no values of another lottery, say C, then $\Pi(A, B, C) = 0$. Thus, we need to consider only the case in which the values of the three lotteries can be ordered as A,B,C,A,B,C. Denote by α,β,γ the probabilities of the highest value of each of the three lotteries A,B,C respectively. Then, $\Pi(A, B, C) = \alpha\beta\gamma(1 - \alpha)$, $\Pi(B, A, C) = \beta\gamma(1 - \alpha)(1 - \beta)$ and $\Pi(C, A, B) = \gamma(1 - \alpha)(1 - \beta)(1 - \gamma)$.

Note that by the continuity of Π, at a maximum point of $V(A, B, C)$ it must be that two of the terms $\Pi(A, B, C)$, $\Pi(B, C, A)$, and $\Pi(C, A, B)$ are equal and are weakly less than the third. If $\Pi(B, A, C)$ is minimal then $\Pi(B, C, A) = \beta\gamma(1 - \alpha)(1 - \beta) = \min\{\alpha\beta\gamma(1 - \alpha),\ \gamma(1 - \alpha)(1 - \beta)(1 - \gamma)\}$. It follows that $1 - \beta \leqslant \alpha$ and $\beta \leqslant 1 - \gamma$ and thus, $\Pi(B, C, A) \leqslant \beta(1 - \beta)(1 - \alpha)\alpha \leqslant \frac{1}{16}$. If $\Pi(B, C, A)$ is not minimal then at the maximum point of V, $\beta\gamma(1 - \alpha)(1 - \beta) > \alpha\beta\gamma(1 - \alpha) = \gamma(1 - \alpha)(1 - \beta)(1 - \gamma)$, hence $1 - \alpha > \beta$ and $\beta > 1 - \gamma$. The maximum with respect to β of the function $\alpha\beta\gamma(1 - \alpha)$ (which is linear in β) given the linear constraints $\alpha\beta = (1 - \beta)(1 - \gamma)$ and $(1 - \alpha) \geqslant \beta \geqslant (1 - \gamma)$ must be obtained where either $\beta = 1 - \alpha$ or $\beta = 1 - \gamma$. In the former case $\alpha\beta\gamma(1 - \alpha) = (1 - \beta)(1 - \gamma)\gamma\beta \leqslant \frac{1}{16}$ while in the latter $\alpha\beta\gamma(1 - \alpha) = \alpha(1 - \gamma)\gamma(1 - \alpha) \leqslant \frac{1}{16}$. ∎

When the support of each of the random variables has at most three points, numerical methods prove that the maximum of $V(A, B, C)$ is roughly

0.0910 and is attained near the triple of random variables:

Value	A	B	C
6	0.19		
5		0.37	
4			0.63
3		0.63	
2	0.62		
1			0.37
0	0.19		

We do not know what are the upper bounds for the cases of at most n outcomes in each lottery for $n > 3$. It is worthwhile mentioning, though, that almost all experiments in the literature utilize lotteries with no more than three different outcomes each as lotteries with more prizes are often difficult to absorb. (For a survey of this literature see Starmer (2000)).

The probability of a cycle can be reduced even further if the MC can choose the first couple of alternatives and only after he observes their realizations he determines which of the two alternatives will be compared with the third one at the second stage. Using numerical methods we conclude that for any triple of lotteries with no more than three outcomes the MC can present the comparisons such that the probability of a cycle is not greater than $\frac{1}{32}$. Moreover, if each lottery has at most two outcomes cycles can be eliminated:

Claim 7 *Let A, B, and C be three binary random variables. If the decision maker follows the Random Sampling procedure with partial recall then the MC who observes the realizations can arrange the order of comparisons so that no cycles emerge.*

Proof: Suppose that between the outcomes of one lottery, say A, there are no outcomes of another lottery, say B. Then the MC will ask the decision maker to compare A and B and then B and C. Assume $B \succ A$. If $B \succ C$ then there is no cycle. If $C \succ B$ then the fresh realization of C is higher than both values of A and at the third stage $C \succ A$. The case that $A \succ B$ is similar.

Suppose that the outcomes are ordered $a_1 > b_1 > c_1 > a_2 > b_2 > c_2$. The MC's instructions could be the following: Start by comparing A and B. Then,

1. If the realization of A is a_1 continue with comparing A and C. Whatever is the realization of C, $A \succ C$ and hence no cycle.

2. If the realizations are a_2 and b_1 ($B \succ A$) then continue by comparing B and C. Whatever is the realization of C, $B \succ C$, hence no cycle.

3. If the realizations are a_2 and b_2 then $A \succ B$. Proceed to compare B and C. If the realization is c_1 then $C \succ B$, hence no cycle. If the realization is c_2 then $B \succ C$ and when A and C are compared (using c_2) then $A \succ C$ and there is no cycle. ∎

5   Comment

This paper discussed the relative stability of random sampling procedures, but our results may be relevant for experimental choice theory as well. Researchers in the field should be aware of the fact that procedures of choice which are not based on transitive preferences may still yield transitive behavior over three alternatives with fairly high probability. Low frequency of cycles over three alternatives could be explained by a model of rational decision makers who, with some probability, make independent errors in each comparison. But it may also be consistent with a situation where the subjects are using procedures of random choice and are manipulated by researchers who wish to show that people are rational. For example, consider an experimenter who would like to prove that people hold transitive preferences. He asks a psychologist, who thinks otherwise, to suggest ten triples of lotteries that in his view are likely to lead to cycles. He requires that no two lotteries in the same triple will have a common outcome and for simplicity he also requires that each lottery will have three outcomes at most. The psychologist provides the experimenter with ten triple of lotteries $\{A_i, B_i, C_i\}_{i=1,\dots,10}$. Each of the subjects is asked to make the thirty binary choices, three for each triple. A person is said to reveal a cycle in triple i if his choices from $\{A_i, B_i\}$, $\{B_i, C_i\}$, and $\{A_i, C_i\}$ are either A_i, B_i, and C_i, or B_i, C_i and A_i. For each subject, the experimenter counts the number of cycles (out of a possible ten), and reports the following results:

# of cycles	0	1	2	3	4	5	6	7	8	9	10
% of subjects	73	23	3	1	0	0	0	0	0	0	0

The experimenter claims that the data can be nicely explained by a theory according to which the decision maker activates a transitive preference relation and there is a 3% chance that he makes a mistake when making a choice. The meaning of our results here is that the data can also be explained by the decision maker following a random sampling procedure with partial recall and the experimenter manipulating the order by which he requires the subjects to make the comparisons (see the analysis preceding Claim 7).

In this paper we were interested in the burden put on a nontransitive decision maker by his indecisiveness. This is different from the money pump argument (discussed in Yaari (1998)) which is based on the existence of a manipulator who is using the cycle to pump out the decision maker's resources. This argument is based on the assumption that the decision maker is myopic and unaware of the consequences of his trades. Our analysis is directed at decision makers who are fully aware of their situation. It is this awareness that prevents them from making decisions. As we showed, comparisons that are based on random selections lead to reasonable upper bounds on the probability of cycles, and clever external agents can push these probabilities down to close to zero.

References

[1] Bernheim, B. Douglas and Antonio Rangel (2009). "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics*, 124:51–104.

[2] Block, H. D. and Jacob Marschak (1960). "Random Orderings and Stochastic Theories of Responses," in *Contributions to Probability and Statistics*, (I. Olkin et al. eds.) Stanford University Press.

[3] Gardner, Martin. (1970). "The Paradox of the Nontransitive Dice and the Elusive Principle of Indifference," *Scientific American*, 223:110–114.

[4] Moon, J. W. (1966). "On Subtournaments of a Tournament," *Canadian Mathematical Bulletin*, 9:297–301.

[5] Moon, J. W. and L. Moser. (1962). "Almost all Tournaments are Irreducible," *Canadian Mathematical Bulletin*, 5:61–65.

[6] Osborne, Martin and Ariel Rubinstein (1998). "Games with Procedurally Rational Players," *American Economic Review*, 88:834–847.

[7] Regenwetter Michel, Jason Dana, and Clintin P. Davis-Stober (2011). "Transitivity of Preferences," *Psychological Review*, 118:42–56.

[8] Rubisntein, Ariel and Yuval Salant (2011). "Eliciting Welfare Preferences from Behavioral Datasets," *Review of Economic Studies*, forthcoming.

[9] Savage, Richard P. (1994). "The Paradox of Nontransitive Dice," *The American Mathematical Monthly* 101:429–436.

[10] Starmer, Chris (2000). "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature* 38:332–382.

[11] Yaari, Menahem E., (1998). "On the Role of 'Dutch Books' in the Theory of Choice under Risk," in *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures 1983–1997*, (D. Jacobs et al. eds.) Cambridge University press.