Max Wertheimer
Minerva Center

Technion
Israel Institute of
Technology

On the tendency to rely on small samples

Ido Erev and Kinneret Teoderescu

Based on a research proposal with Ralph Hertwig, MPI Berlin

Previous research reveals a gap between decisions that are made based on a description of the payoff distributions, and decisions from experience (Barron & Erev, 2003; Hertwig et al., 2004; Yechiam et al., 2005) .

The clearest deviation from maximization in decisions from description reflects overweighting of rare events (Kahneman & Tversky, 1979). Experience leads to the opposite bias

S:0 with certainty
R: +1 in 90% of the trials; -10 in the other 10%

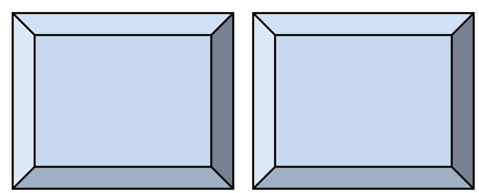

Two choice prediction competitions suggest that the counterproductive effect of experience is best captured as the product of the tendency to rely on small samples (Erev et al., 2010a, 2010b) .

For example the tendency to prefer the bad gamble here
(+1 in 90% of the trials; -10 in the other 10%)
can be the product of the fact that most samples of 6 or less experiences with this gamble are positive: they do not include the -10 outcome.

We try to clarify the relationship between this observation and mainstream research in psychology.

In Hertwig et al. (2004) we show that in some cases the tendency to rely on small sample can be a result of cognitive limitations and/or laziness.

The current research focuses on another contributor to this tendency. We hypothesize that it can be a byproduct of **adaptive perceptual organization and recognition rules** (like the rules discussed yesterday in Baruch's and Halamish's talks).

The basic idea behind our analysis is the assumption that human learning was evolved to address dynamic environments. Thus, people use rules that approximate the optimal strategy in a wide set of dynamic settings, but leads to deviations from maximization in the subspace of static settings.

Example:

S: 0 with certainty

R: V_{gain} if the state of nature is G1 or G2; V_{loss} otherwise

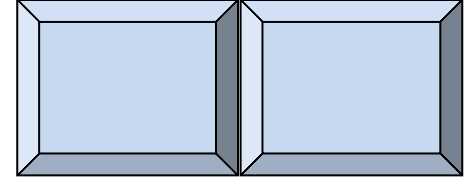

And the state of nature is determined by the following Markov chain:

		State at trial t+1			
		G1	L1	G2	L2
State at trial t	G1	$P_{1,1}$	$P_{1,2}$	$P_{1,3}$	$P_{1,4}$
	L1	$P_{2,1}$	$P_{2,2}$	$P_{2,3}$	$P_{2,4}$
	G2	$P_{3,1}$	$P_{3,2}$	$P_{3,3}$	$P_{3,4}$
	L2	$P_{4,1}$	$P_{4,2}$	$P_{4,3}$	$P_{4,4}$

When the payoff rule is unknown, the computation of the optimal strategy is very difficulty (the solution is unknown). Nevertheless, it is easy to approximate this strategy with simple classification rules.

To clarify the logic behind these rules consider the choice at trial 11 after observing the following sequence:

Trial	1	2	3	4	5	6	7	8	9	10	11	
Payoff	L	G	G	G	L	L	L	L	L	G		
Class1		G	G	G	L	L	L	L	L	G		
Class2			G	G	L	L	L	L	L	G		

Rule Class1, classifies all trials based on the prior payoff. The target class is 3, 4 & 5 (all trials after G), and the rest. The decision is made based on the average payoff in the target class

Rule Class2, classify all trials based on the 2 prior payoffs. The target class is 3(all trials after LG), and the rest.

Performance of five classification rules. Rule Class-k classifies all trials that follow the same sequence of k trials to the same category. FIBA is a fully informed Bayesian agent. The agent's performance represents an upper bound for best performance. Performance was computed for a 1,000 trial experiments.

Strategy	Mean payoff	Prop. Of choices Consist with FIBA
Fully informed rational agent (FIBA)	20.69	1
Class-1	18.66	0.91
Class-2	20.42	0.98
Class-3	20.45	0.97
Class-4	20.25	0.97
Class-5	20.36	0.96

Notice that these rules imply reliance on small samples. For example, Class-5 rely on 1/32 of the observations. The reliance on small samples is not costly the dynamic features of the environment are important, but it leads to deviations from maximization when the environment is static.

For example:
S:0 with certainty
R: +1 if G1 or G2; -10 otherwise

Class-5 prefer R
(the low EV Option) in 60%

		State at trial t+1			
		G1	L1	G2	L2
State at trial t	G1	.9	.1	0	0
	L1	.9	.1	0	0
	G2	.9	.1	0	0
	L2	.9	.1	0	0

But how do people select among the different classification rules?

One interesting meta-rule is Take-the-Best (TTB)

This meta-rule selects the classification rule with the highest discrimination score. The discrimination score is the absolute difference between the G-rate in the two classes.

Trial	1	2	3	4	5	6	7	8	9	10	11	Discrimination Score
Payoff	L	G	G	G	L	L	L	L	L	G		
Class1		G	G	G	L	L	L	L	L	G		ABS(2/3-2/6)=1/3
Class2			G	G	L	L	L	L	L	G		ABS(1/1-2/7)=5/7

Pilot results suggest that an algorithm that uses TTB to select among the classification rules outperforms the best stable classification rule, but enhances underweighting of rare events.

Mainstream research of simple cognitive rules (fast and frugal heuristics, perceptual organization and recognition rules) tries to discover the rules that people bring to the lab in relatively complex setting.

Our research tries to shed light on the emergence of similar rules in simple settings.

We believe that similar processes have underlie the evolution of more complex rule in natural settings.