On Maximization, Learning, and Predictions

Ido Erev, Technion and Univ. of Warwick

Based on research with Alvin E. Roth, Kinneret Teodorescu, Eyal Ert, and Ori Plonsky PNAS: http://iew.technion.ac.il/lad/files/Erev-Roth-Feb14-2014.pdf

Review of the most successful inventions and discoveries in the social sciences reveals an apparent inconsistency. The best inventions (e.g., trading, markets, contracts, law enforcement, banks, auctions, typical recommendation systems) are methods to reduce conflicts under the assumption that people maximize expected return, and the most elegant discoveries are demonstrations of deviations from maximization (e.g., Kahneman & Tversky, 1979).

This gap led mainstream research to focus on gentle generalizations of the maximization assumption. For example, Bernoulli (1734) added a risk aversion parameter, and Kahneman and Tversky (1979) added 5 different parameters.

We try to clarify the value of an alternative approach.

We suggest that gap is a reflection of the fact that maximization based inventions are successful when they create incentive that facilitate quick learning to select the desirable actions.

Instead of adding parameters to a weak maximization model, **we try to clarify the conditions under which learning leads to maximization.**

The clicking paradigm (Barron & Erev, 2003)

The current experiment includes many trials. Your task, in each trial, is to click on one of the two keys presented on the screen. Each click will be followed by the presentation of the keys' payoffs. Your payoff for the trial is the payoff of the selected key.

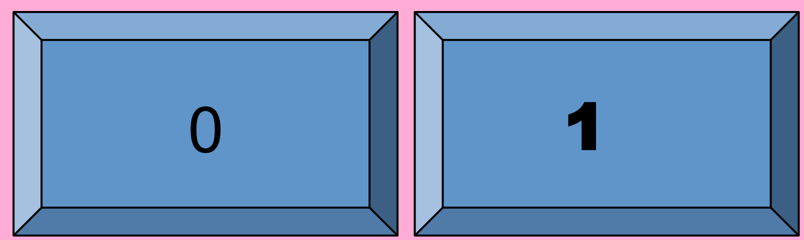

You selected Right. Your payoff in this trial is **1**
Had you selected Left, your payoff would be 0

The results reveals fast convergence to maximization in some settings, and consistent deviations from maximization in other settings. **A particular large and robust deviation from maximization is documented when the best option on average is worst most of the time (Barron & Erev, 2003)**

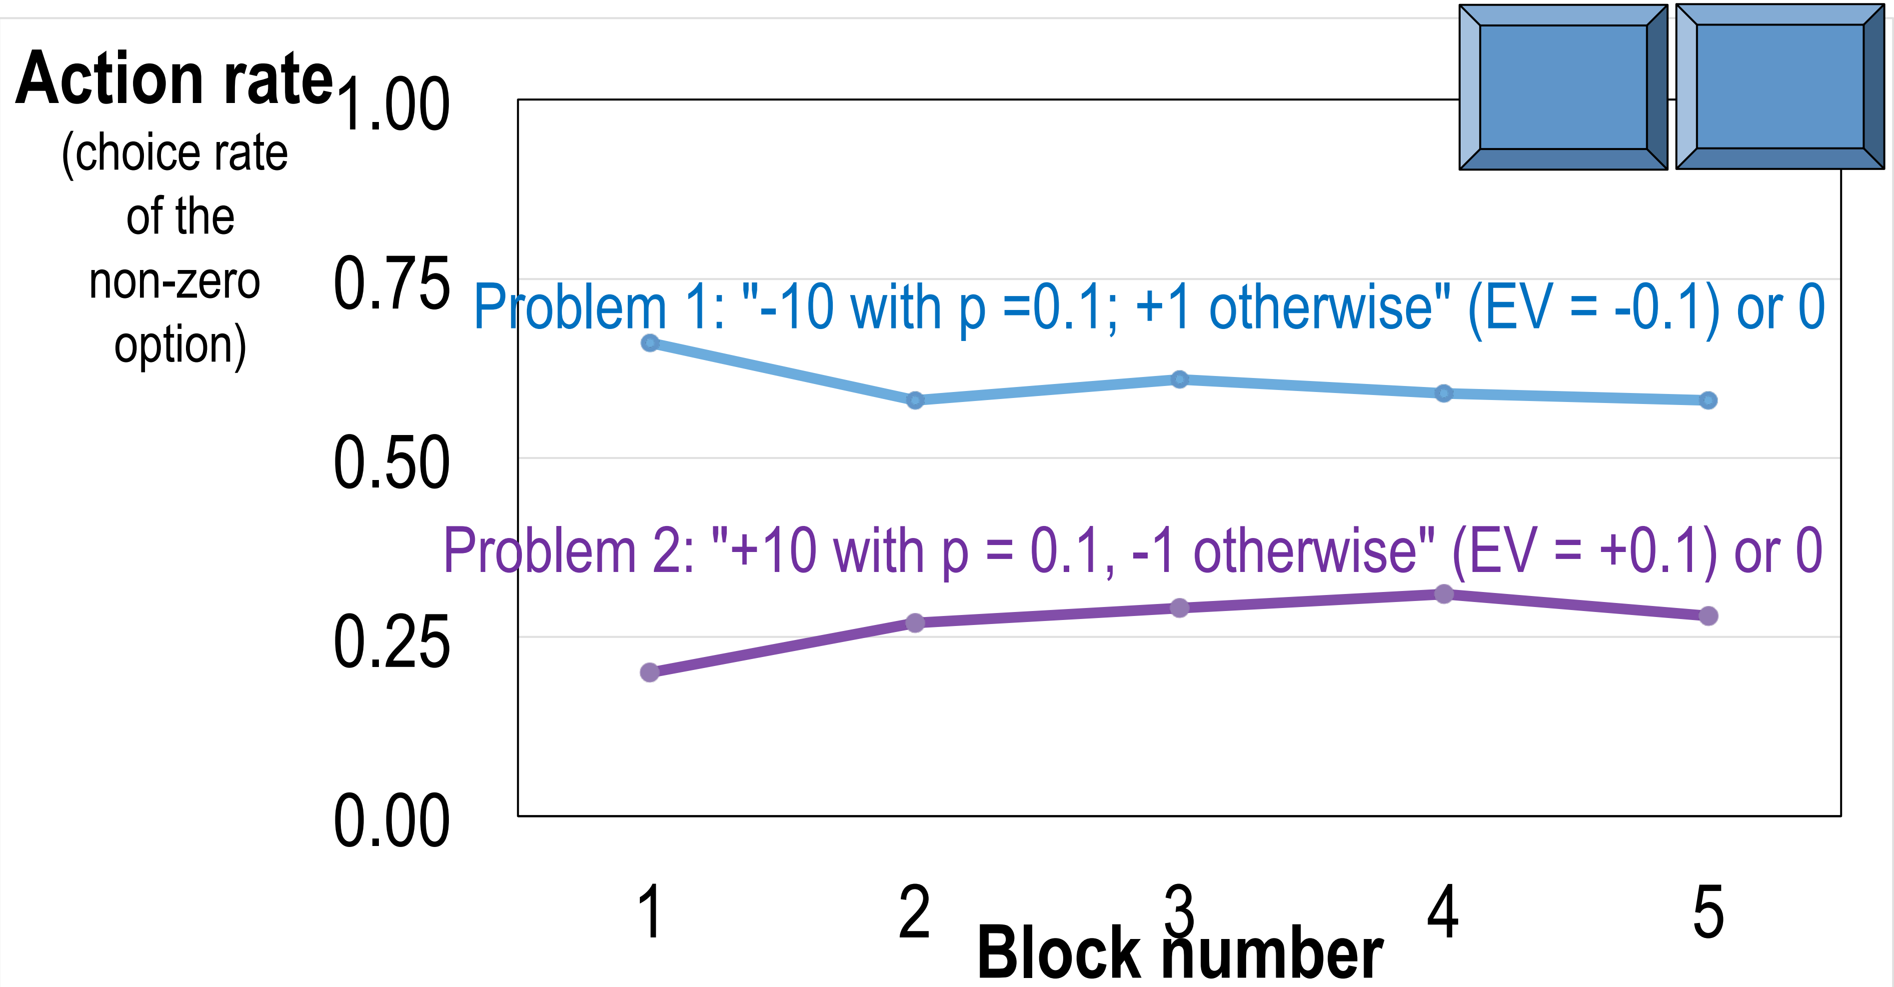

Underweighting of rare events

The tendency underweight rare events is rather robust. It was documented in:

- Perceptual decisions; Signal detection tasks (Barkan, Erev & Zohar, 1998)
- Repeated decisions with limited feedback (Barron & Erev, 2003)
- One shot decisions based on free sampling (Hertwig et al., 2004; Rakow & Newell, 2010).
- Multiple alterenatives (Ert & Erev, 2007)
- Different cultures (Di Guida, Marchiori & Erev, 2014)
- Bees decisions (Shafir et al., 2008)
- Decisions among defaults (Di Guida, Erev & Marchiori, 2012)
- Decisions based on description and experience (Yechiam et al., 2005; Jessup, Bishara, & Busemeyer, 2008; Lejarraga & Gonzelez, 2011).
- Investment decisions (Selten, Pittnauer, Hohnisch, 2014; Taleb, 2007)

Another robust class of deviations from maximization is documented in studies that examine the effect of payoff variability

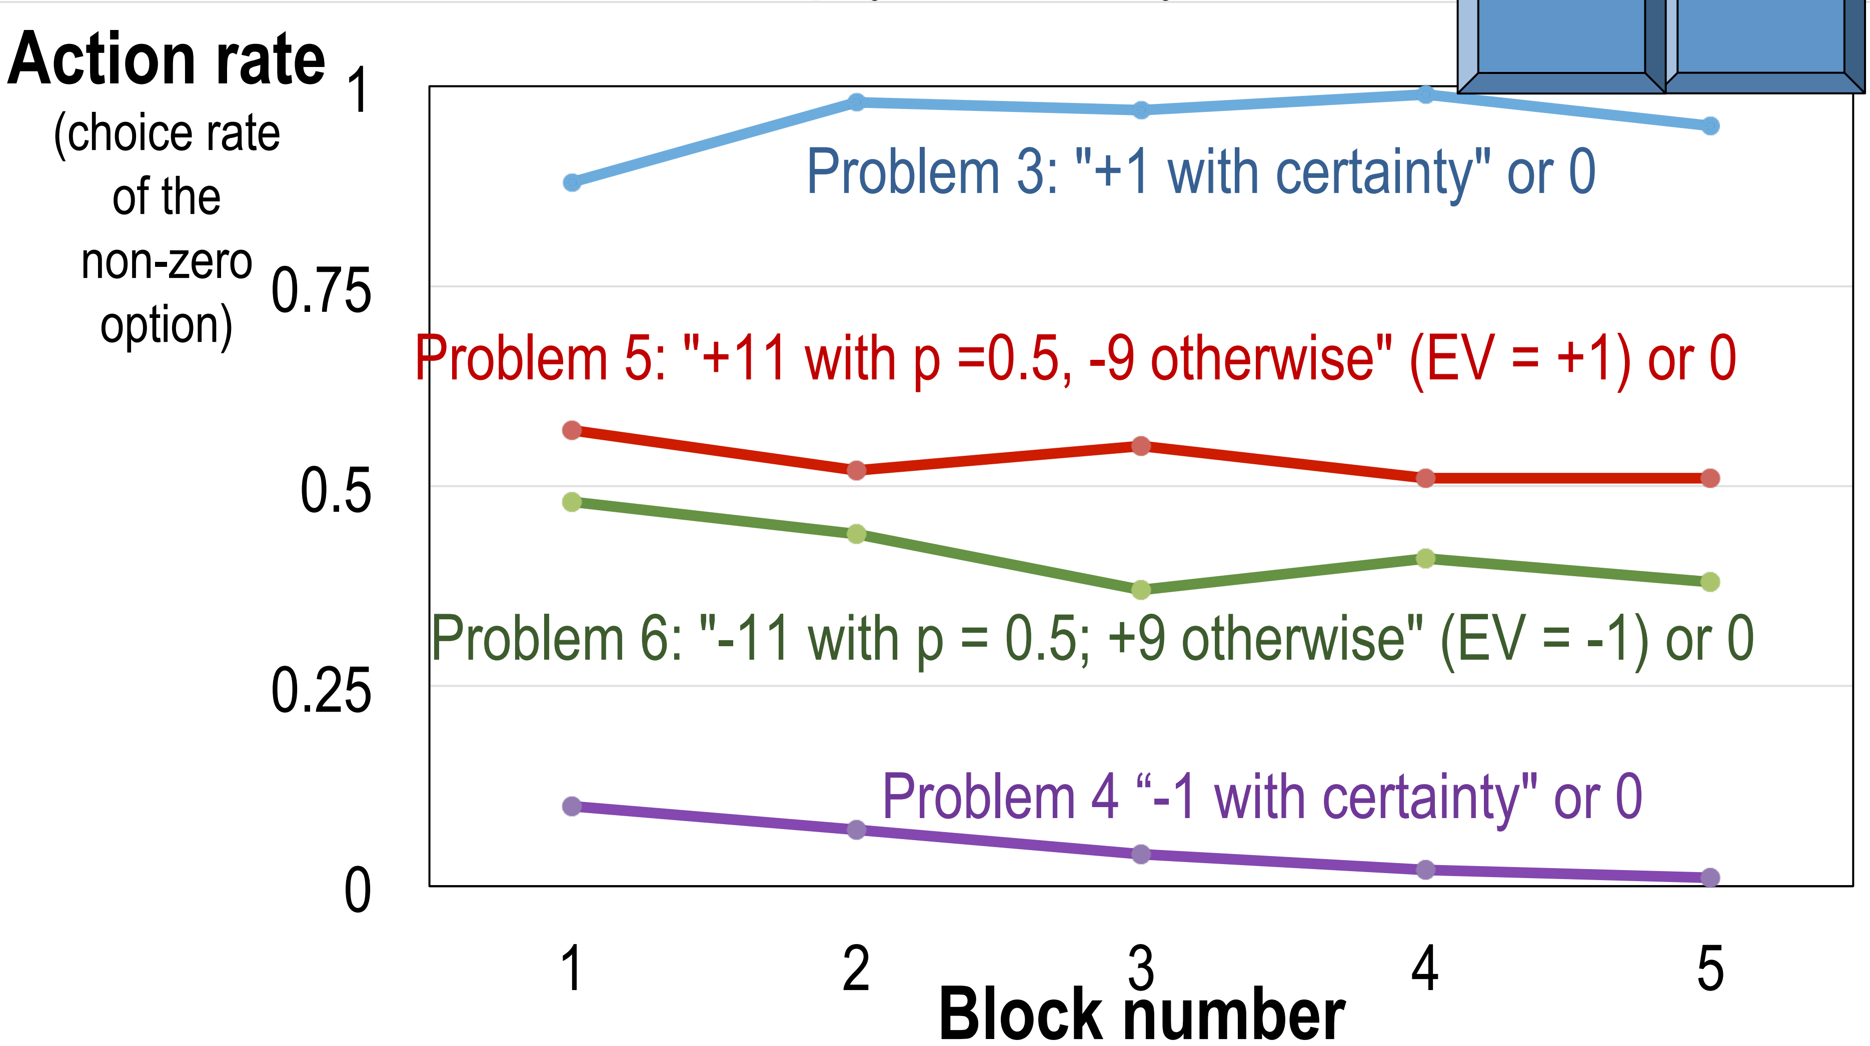

The payoff variability effect (Busemeyer and Townsend, 1993; Diederich and Busemeyer, 1999).

The reliance on small sample explanation of the basic properties of decisions from experience (Hertwig et al., 2004, and see related observations in Kareev 2000; Fiedler, 2000)

Action (the alternative to the status quo)	EV	Action-rate	Prediction of a 'sample of 5' model
(-10, 0.1; +1, 0.9)	-0.1	60%	62%

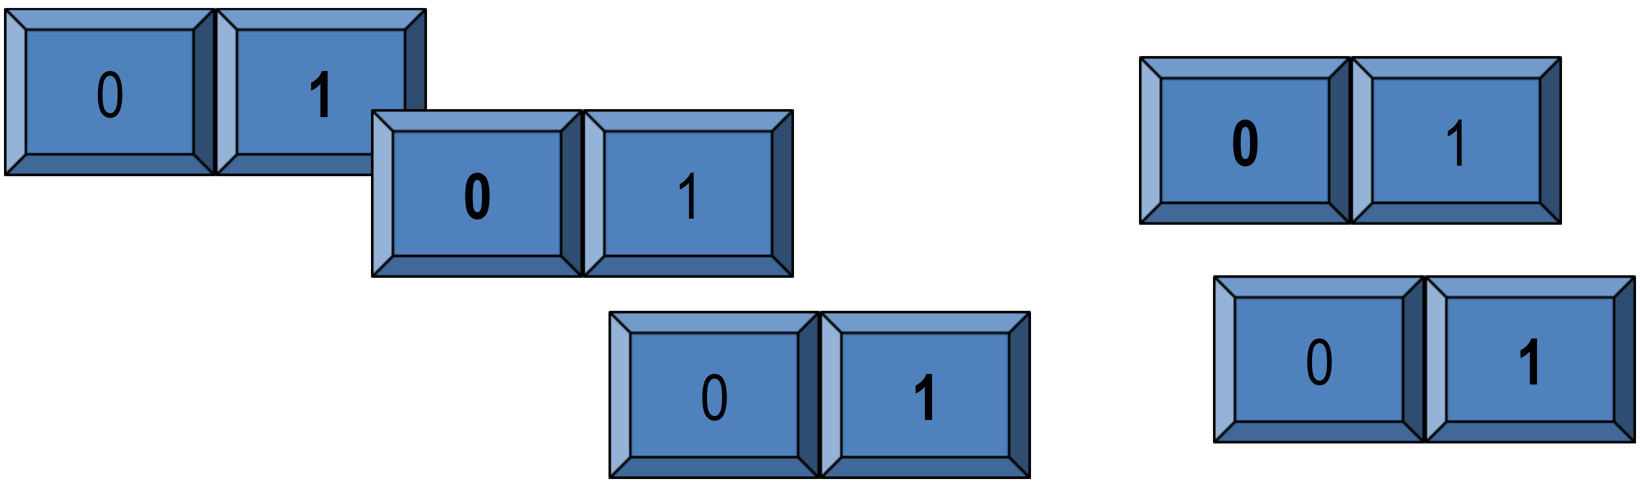

A sample of size k will include the rare event with probability below .5 when P(no rare) = $(1-p)^k < 0.5$. This inequality implies that $k < \text{Log}(.5)/\text{Log}(1-p)$. For example, when p = .1, when k < 6.57.

The reliance on small sample explanation of the basic properties of decisions from experience (Hertwig et al., 2004, and see related observations in Kareev 2000; Fiedler, 2000)

Action (the alternative to the status quo)	EV	Action-rate	Prediction of a ‘sample of 5’ model
(-10, 0.1; +1, 0.9)	-0.1	60%	62%
(+10, 0.1; -1,0.9)	+0.1	27%	38%
+1 with certainty	+1	95%	99%
-1 with certainty	-1	5%	1%
(+11, 0.5; -9, 0.5)	+1	53%	50%
(+9, 0.5; -11, 0.5)	-1	42%	50%

Why small samples (Plonsky, Teoderescu & Erev, 2014)? Consider a variant of the clicking experiment in which each trial starts with a presentation of a color signal. What would you select in trial 16?

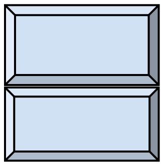

Pre-choice signal:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
	-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

Why small samples (Plonsky, Teoderescu & Erev, 2014)? Consider a variant of the clicking experiment in which each trial starts with a presentation of a color signal. What would you select in trial 16?

Pre-choice signal:

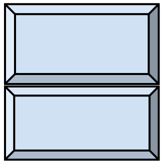

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

Most subjects select Up, even without the color signals:

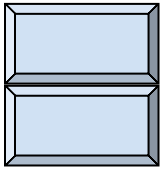

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	+2	-1	-1	-1	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

Thus, they behave as if they use on 3 past experiences. And in trial 8 of the following example, many behave as if they use only 1 past experience:

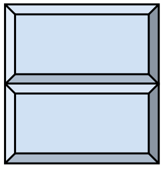

1	2	3	4	5	6	7	8
-1	-1	-1	+2	-1	-1	-1	
0	0	0	0	0	0	0	

Relationship to machine learning and recommendation systems

We propose that people, like leading machine learning algorithms, tend to select the options that worked best in the most similar situations in the past. This tendency approximates the optimal behavior when there is strong structure in the environment (e.g., payoffs are a function of the state of nature, and the state is determined by a Markov process), but implies underweighting of rare events (exceptions).

In other words, we suggest that designers of external recommendation systems should consider the possibility that people has internal recommendation systems that decide which external systems to use, and how to response to the external recommendations.

Skinner, Chomsky and Norvig

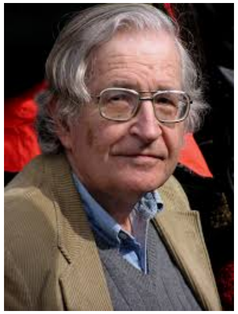

Multiple alternatives:
Chasing past returns (and fleeing from past losses)
The Big Eye effect (Ben Zion et al., 2010, Grosskopf, Yechiam & Erev., 2006)

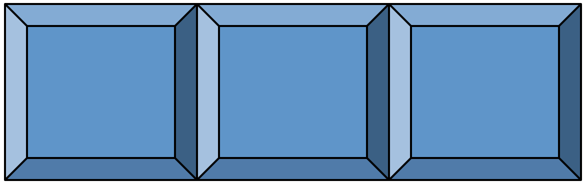

x ~ N(0,300), y ~ N(0, 300)

R1: x (EV = 0)
R2: y (EV = 0)
M: Mean(R1,R2) + 5 (EV = 5)

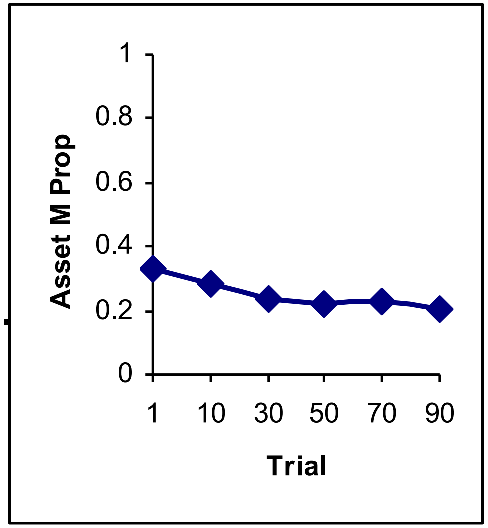

Deviation from: maximization, risk aversion, loss aversion.
Implies under-diversification
Robust to prior information

The hot stove effect (Mark Twain; Denrell & March, 2001; Hogarth & Einhorn, 1992).

Giving up too early and over-exploration (Teodorescu & Erev, 2013): Many behavioral problems appear to reflects insufficient exploration, but in other cases people appear to exhibit over-exploration.

	EV
Problem Rare disasters: 10% disasters (-10), 90% rewards (+1)	**-0.1**
Problem Rare treasure: 10% treasures (+10), 90% disappointments (-1)	**+0.1**

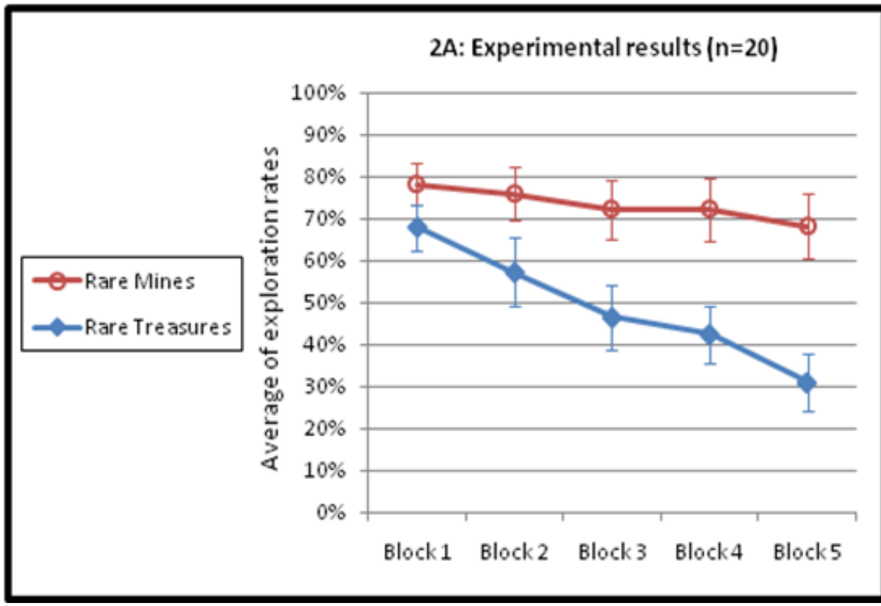

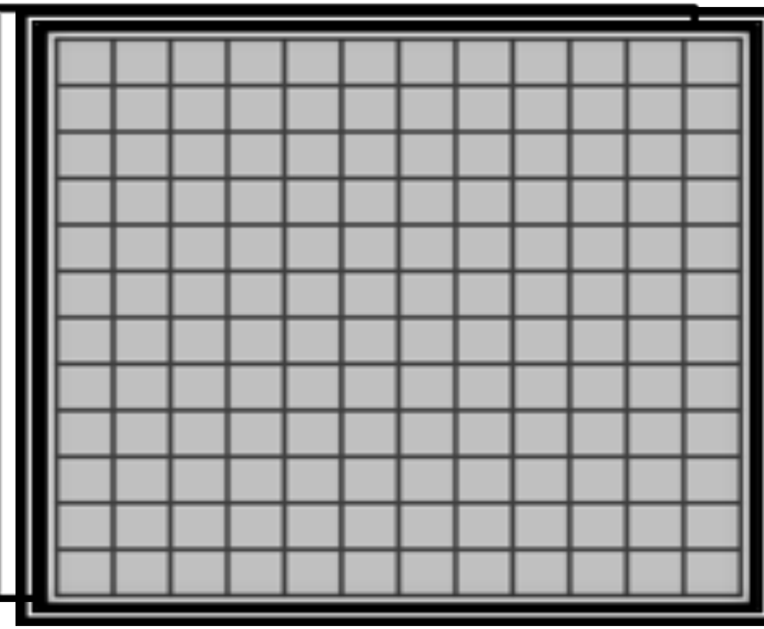

The two stage explanation: An initial choice whether to explore (partially based on a small sample), and then a choice between the alternatives.

Learned helplessness (Seligman, 1975)

	Experimental group	
Stage	With Control	Yoked
Training Task 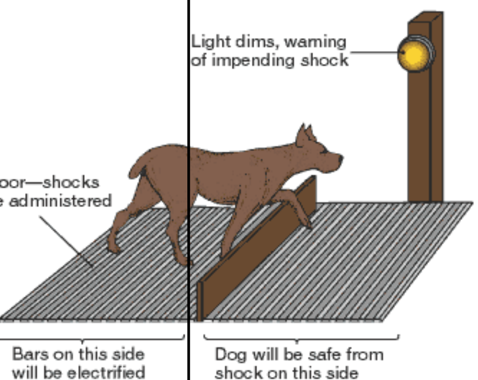	Dogs were exposed to shocks applied to their foot pads while restrained in a hammock. Flat panels located on either side of the dog would immediately terminate shock if pressed by side-to-side movement of its head.	Each dog in this group was wired in parallel with a dog from the "With Control" group to control for equal durations and timing of shocks. These dogs experienced the shocks as inescapable.
Identical Test Task	Both groups were placed in a shuttle box, exposed to escape-avoidance training. The required response in order to avoid shock was jumping over a hurdle into an adjacent compartment.	
Results	90% of dogs readily learned the escape avoidance response.	2/3 of dogs failed to escape shock. Instead of jumping over the hurdle they laid down huddling passively in the corner.

Two explanations:

1. Learning that the agent has no control
2. Underweight rare events

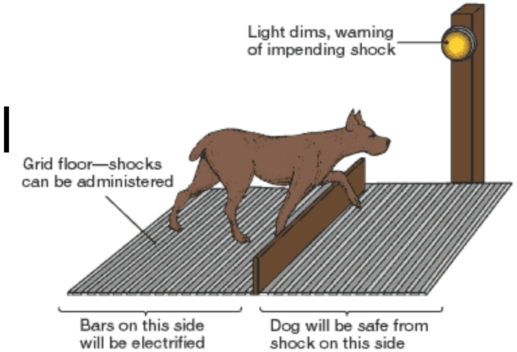

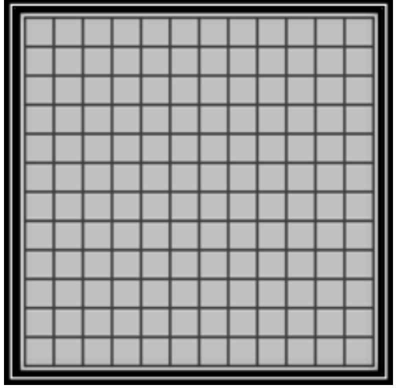

Group	Practice		Test	
	Cost of exploration	Potential benefit	Cost of exploration	Potential benefit
No control	Loss of 1	+10 if **another** agent finds a treasure (x% of cells)	Loss of 1	+10 if the agent finds a treasure (x% of cells)
With Control	Loss of 1	x+10 if the agent finds a treasure (x% of cells)	Loss of 1	+10 if the agent finds a treasure (x% of cells)

X	Practice w/o control	Practice with control
10	40,40	45, 40
20	40,45	70,70
100	40 90	90, 90

The experience-description gap (Barron & Erev, 2003; Hertwig & Erev, 2009).

Kahneman and Tversky's (1979) classical study of decisions under risk reveals overweighting of rare events

	Choice rate consistent with overweighting
S: 5 with certainty **R: 5000 with probability 1/1000, 0 otherwise** (EV = 5)	72%
S: -5 with certainty R: -5000 with probability 1/1000, 0 otherwise (EV = -5)	80%

The study of decisions from experience reveals the opposite pattern

S: 0 with certainty R: 10 with probability 1/10, -1 otherwise (EV = +0.1)	27%
S: 0 with certainty **R: -10 with probability 1/10, 1 otherwise** (EV = -0.1)	42%

Evaluation of the experience-description gap shows that it can emerge within the same task when the subjects can use both description and experience:

Lejarraga & Gonzelez (2011, Yechiam, Barron & Erev, 2005): 100 trials with feedback and complete description S: 3 with certainty R: 64 with probability 1/20, 0 otherwise (EV = 3.2)	R-rate (choice rate consistent with overweighting of rare events)
First trial (decision from description)	55%
Trials 81-100 (based on description and experience)	30%

We believe that the experience-description gap reflects the existence of two classes of deviations from maximization.

First, initial decisions reflect overgeneralizations from past experiences in similar, but not identical situations (biases). Initial overweighting of rare events, and many of the other demonstrations of deviations from maximization in decisions from description belong to this class.

Second, when people gain repeated experience in a given setting they rely on the most similar past experiences. This tendency tends to increase maximization, but imply reliance on small samples and three robust deviations from maximizations:

- Underweighting of rare events
- The big eyes effect
- Hot stove effect

Four classes of general implications : Many avoidable inefficiencies

Deviations from maximizations

Car inspection (description, defaults or enforcement)

Using safety devices and Gentle rule enforcement (Schurr et al., 2014)

Not checking (Y. Roth Waenke & Erev, 2014)

A (10 or -10), B (-10 or 10), RecSys (1 for sure)

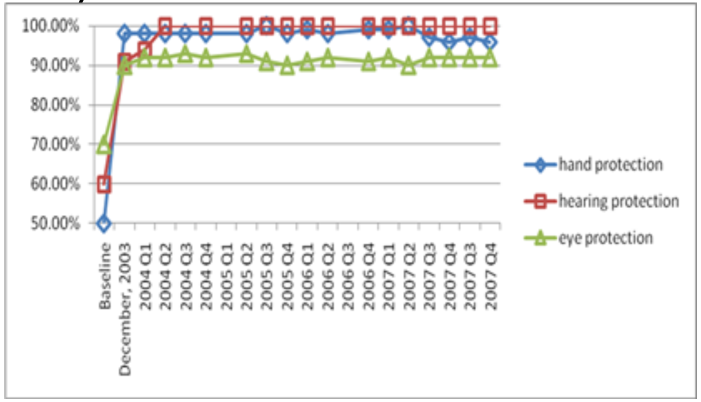

Slow learning

First vs. Second price auctions

New middle east

Convergence to inefficient self confirming equilibria

Corporal punishment (Skinner, 1953; Sobolev & Erev)

Dinners in conferences

Convergence to inefficient Nash equilibria

Cheating in exams

Vandalism (Hreib & Erev)

More specific implications to recommendation systems and users experience

Steve Job's "it is not the consumer Job's to know what they want"

Microsoft strategy and Windows 8

Stable preferences, constructed preferences, and learning.

e-contracts

Converging to few sites/applications

Learned helplessness and the value of free gifts and friends

Part II: Nine classical choice phenomena, and a choice prediction competition

(with Ori Plonsky, Technion
And Eyal Ert, Hebrew Univ.)

Summary of the part I:

Many successful social inventions can be described as methods to reduce conflicts under the assumption that people maximize expected return.

We hypothesize that their success is a result of the fact that they create conditions that facilitate learning to maximize.

For example, they insure that the maximizing strategy minimizes probability of regret.

This condition insures quick learning even by agents that rely on small samples of past experiences in similar situations (and exhibit underweighting of rare events, like our subjects).

The current paper tries to clarify the relationship between the basic properties of decisions from experience and nine classical choice phenomena. We consider the following phenomena

1. The St. Petersburg paradox (Risk aversion)
2. The Allais paradox (Certainty effect)
3. The Ellsberg paradox (Ambiguity aversion)
4. Overweighting of rare events (Buying lotteries and insurance).
5. The reflection effect (risk aversion in the gain, risk seeking in the loss domain)
6. Contingent loss aversion (rejecting moderate stakes mixed gambles)
7. Break even effect (more risk seeking if this action can prevent loss)
8. Get something effect (more risk aversion when this action guarantees gain)
9. Regret and correlation effect

Method

Study 1: Replicating the classical phenomena under a single experimental setting that involves binary decisions from description (no feedback) with small real payoffs, and then exploring the effect of feedback in this setting.

- 30 problems that replicate the classical demonstrations.
- 25 choices per problem, feedback was provided after the 6th trial.
- Payoff in Shekels for one (of the 750) randomly selected trial.

Study 2: Same design as study 1, with 60 randomly selected problems. The results, and baseline models, will be published on the web. We will challenge other researchers to participate in the competition to predict the results of Study 3. The participants will be asked to submit a computer program that reads the parameters of the problems, and derives the predicted choice rate as output. To qualify, it has to capture the classical phenomena, and to be clear. The winner is the submission with the lowest prediction error (Mean Squared Deviation score).

Study 3: Same as Study 2, with a different sample of 90 problems.

Example of the basic experimental task, 1

Please select one of the following options:

A:
3 with certainty

B:
4 with $p = 0.8$
0 with $p = 0.2$

A	B

Example of the basic experimental task, 2 (trial 2, no feedback)

A:
3 with certainty

B:
4 with p = 0.8
0 with p = 0.2

A

B

You selected B

Example of the basic experimental task, 3

Please select one of the following options:

A:
3 with certainty

B:
4 with $p = 0.8$
0 with $p = 0.2$

A	B

Example of the basic experimental task, 3 (with feedback)

A:
3 with certainty

B:
4 with p = 0.8
0 with p = 0.2

A	**B**
3	**4**

You selected B, your payoff is 4
Had you selected A your payoff would be 3

The St. Petersburg “paradox” (after Bernoulli, 1734)

Coin,	Block	1, No FB	5, with FB
S	9 with certainty		
R	Play the following game: A coin will be flipped until the first fall on Head, but not more than 8 time. Your payoff will be 2^k where k is the number of flips	37%	

List			
S	9 with certainty		
R	2 with p=.5 4 with p=.25 8 with p=.125 16 with p=.0625 32 with p=.03125 64 with p=.0015625 128 with p=.00078125 256 with p=.00078125	38%	

Risk aversion.

The St. Petersburg “paradox” (after Bernoulli, 1734)

Coin,	Block	1, No FB	5, with FB
S	9 with certainty		
R	Play the following game: A coin will be flipped until the first fall on Head, but not more than 8 time. Your payoff will be 2^k where k is the number of flips	37%	31%

List			
S	9 with certainty		
R	2 with p=.5 4 with p=.25 8 with p=.125 16 with p=.0625 32 with p=.03125 64 with p=.0015625 128 with p=.00078125 256 with p=.00078125	38%	36%

Risk aversion. The added feedback appears to increase risk aversion

The Allais paradox, common ratio, certainty effect (Allais, 1953, K&T, 1979)

Common ratio ¼	Block	1, No FB	5, with FB
S	3 with certainty		
R	4 with p= .8, 0 otherwise	31%	

Common ratio 1/4			
S’	3 with p =.25, 0 otherwise		
R’	4 with p= .2, 0 otherwise	52%	

Certainty effect from description.

The Allais paradox, common ratio, certainty effect (Allais, 1953, K&T, 1979)

Common ratio ¼	Block	1, No FB	5, with FB
S	3 with certainty		
R	4 with p= .8, 0 otherwise	31%	61%

Common ratio 1/4			
S'	3 with p =.25, 0 otherwise		
R'	4 with p= .2, 0 otherwise	52%	59%

Certainty effect from description.
The addition of feedback increases maximization and eliminates the paradox

The Ellsberg paradox (Ellsberg, 1961)

Ellsberg 50	Block	1, No FB	5, with FB
S	10 with p= 0.5; 0 otherwise		
R	10 or 0	32%	

Ellsberg 90			
S	10 with p= 0.9; 0 otherwise		
R	10 or 0	11%	

Ellsberg 10			
S	10 with p= 0.1; 0 otherwise		
R	10 or 0	86%	

Ambiguity aversion plus a bias toward uniform priors.

The Ellsberg paradox (Ellsberg, 1961)

Ellsberg 50	Block	1, No FB	5, with FB
S	10 with p= 0.5; 0 otherwise		
R	10 or 0	32%	51%

Ellsberg 90			
S	10 with p= 0.9; 0 otherwise		
R	10 or 0	11%	32%

Ellsberg 10			
S	10 with p= 0.1; 0 otherwise		
R	10 or 0	86%	66%

Ambiguity aversion plus a bias toward uniform priors.
The addition of feedback eliminates ambiguity aversion. The effect of the initial description is larger when the uniform resolution is wrong.

Insurance, lotteries and underweighting of rare events (Kahneman &Tversky, 1979

	Block	1 (No FB)	5 (with FB)
S	1 with certainty		
R	20 with p= .05, 0 otherwise	40%	

S	1 with certainty		
R	100 with p= .01, 0 otherwise	48%	

S	2 with certainty		
R	101 with p= .01, 1 otherwise	55%	

Some overweighting of rare events before feedback,

Insurance, lotteries and underweighting of rare events (Kahneman &Tversky, 1979

	Block	1 (No FB)	5 (with FB)
S	1 with certainty		
R	20 with p= .05, 0 otherwise	40%	26%

S	1 with certainty		
R	100 with p= .01, 0 otherwise	48%	39%

S	2 with certainty		
R	101 with p= .01, 1 otherwise	55%	43%

Some overweighting of rare events before feedback, feedback leads to underweighting of rare events (get something masks overweighting)

Reflection effect

	Block	1 (No FB)	5 (with FB)
S	3 with certainty		
R	4 with p= .8, 0 otherwise	31%	

S	-3 with certainty		
R	-4 with p= .8, 0 otherwise	52%	

Risk aversion in the gain and weak risk seeking in the loss domain,

Reflection effect

		Block	1 (No FB)	5 (with FB)
S	3 with certainty			
R	4 with p= .8, 0 otherwise		31%	59%

S	-3 with certainty		
R	-4 with p= .8, 0 otherwise	52%	34%

Risk aversion in the gain and weak risk seeking in the loss domain, feedback eliminates this pattern and increases maximization

Get something (Payne, 2005), and break even (Johnson & Thaler, 1991)

	Block	1 (No FB)	5 (with FB)
S	1 with certainty		
R	2 with p= .5, 0 otherwise	35%	

S	2 with certainty		
R	3 with p= .5, 1 otherwise	45%	

S	-1 with certainty		
R	-2 with p= .5, 0 otherwise	55%	

S	-2 with certainty		
R	-3 with p= .5, -1 otherwise	52%	

Initial get something and break even,

Get something (Payne, 2005), and break even (Johnson & Thaler, 1991)

	Block	1 (No FB)	5 (with FB)
S	1 with certainty		
R	2 with p= .5, 0 otherwise	35%	50%
S	2 with certainty		
R	3 with p= .5, 1 otherwise	45%	53%
S	-1 with certainty		
R	-2 with p= .5, 0 otherwise	55%	51%
S	-2 with certainty		
R	-3 with p= .5, -1 otherwise	52%	51%

Initial get something and break even, but feedback eliminates these effects

Contingent loss aversion (Ert & Erev, 2014)

	Block	1 (No FB)	5 (with FB)
S	0 with certainty		
R	+1 with p= .5, -1 otherwise	49%	

S	0 with certainty		
R	+50 with p= .5, -50 otherwise	34%	

Contingent loss aversion without feedback,

Contingent loss aversion (Ert & Erev, 2014)

	Block	1 (No FB)	5 (with FB)
S	0 with certainty		
R	+1 with p= .5, -1 otherwise	49%	38%

S	0 with certainty		
R	+50 with p= .5, -50 otherwise	34%	38%

Contingent loss aversion without feedback,
feedback increases loss or risk aversion even with low stakes (this effect is inconsistent with previous study of pure decisions from experience that shows risk neutrality).

Regret (Loomes & Sugden, 1982) and correlation (Diederich & Busemeyer, 1999) effects.

P(E) = 0.5

	Block	1 (No FB)	5 (with FB)
S	6 if E, 0 otherwise		
R	**9 if E-not**, 0 otherwise	96%	85%
S	6 if E, 0 otherwise		
R	**8 if E**, 0 otherwise	96%	98%

No sensitivity for regret/correlation without feedback, feedback increases the regret correlation effect.

A baseline model: Best Estimation And Sampling Theory (**BEAST**)

Choose R at trial t iff, (EV[R] – EV[L]) + error term + sample difference(t) > 0

The error term is drawn from N(0,**σ**), and sample size is **κ**

The sample is taken using 4 procedures (one unbiased, and three biased):

- Unbiased
- Step function (all outcomes are replaced by their sign)
- Equally likely (all probabilities are set to 1/m)
- Conditional Pessimistic (if the difference between the ration of the worst possible outcome is large enough, and a gain is possible, the draw is the worst possible outcome, in other cases all the outcomes are drawn with equal probability)

Ambiguity leads to slight overweighting (η) of the worst outcome.

Feedback increases the probability biased sampling.

P(Bias) = $\lambda / (1+r^{\theta})$ where r = number of trials with feedback.

Best fit is obtained with σ = 4, κ =4, λ =.6, θ = .5, η = 0.04

Competition details

Study 2 (60 problem randomly sampled, estimation study)

WE will post the results on the web (by November, 2014), and challenge decision scientist to participate is a competition to predict the results of study 3.

Study 3 will be run by December, 2014.

The submission deadline will be April 2015.

Can be used for class projects (Interested instructors will get their students MSD scores).

Summary:

It is important to distinguish between two classes of behavioral phenomena that affect choice behavior.

A wide set of initial overgeneralizations

A tendency to rely on small samples that implies near maximization following but can lead to underweighting of rar event, the hot stove effect, and the big eyes effect.

Successful recommendation systems, like other successful invension, appear to create a situation in which people learn to select the option that maximize their return (and lead to good social outcomes).

We challenge you to help us improve our understanding of this phenomena by participating in our competition