האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

NEWCOMB`S PROBLEM: PARADOX LOST

By

MENAHEM YAARI

Discussion Paper # 635 April 2013

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

NEWCOMB'S PROBLEM: PARADOX LOST

Menahem Yaari[1]

April 2013

ABSTRACT

An agent needs to decide which of two available actions, **A** or **B**, to take. The agent's payoffs are such that **A** *dominates* **B**, i.e., taking **A** yields a better payoff than taking **B**, in *every* contingency. On the other hand, the agent's *expected* payoffs, given the action taken, are in the reverse order, i.e., E(payoff | **B**) > E(payoff | **A**) , which can happen if the probabilities of the various contingencies are not independent of the action being taken. What should the agent do? This dilemma has come to be known as *Newcomb's Paradox* (Nozick, 1969). The present essay shows that the rule "keep away, as much as possible, from any dominated action" is perfectly consistent with actually taking the dominated action, when appropriate. No paradox.

[1] I wish to thank Ariel Rubinstein for pointing out an error I had made in an earlier draft. Sincere thanks are due also to the Institute for Advanced Study in Berlin (Wissenschaftskolleg zu Berlin) for its hospitality.

1. Introduction

The Theory of Games came into being when thinking about interactive situations – even the simplest ones, like Matching Pennies – gave rise to noxious puzzles and paradoxes: The fleeing villain has but two escape routes, say **V** (go to Victoria Station) or **W** (go to Waterloo Station). The sleuth, in hot pursuit, must decide whether to go **V** or **W** to apprehend the fleeing villain. If **V** is the sleuth's correct action then the villain, knowing this, would go **W**, i.e. **W** is then the sleuth's correct action. And conversely. It thus appears, intolerably, that **V** is the action to take if, and only if, **W** is the action to take. In 1921, Emile Borel pointed out, in effect, that the paradox would disappear if the interacting agents' aspirations were to be seen differently. Instead of the villain seeking an escape route, he should be regarded as seeking to enhance the *chance* of escape. Similarly, instead of the sleuth seeking simply to apprehend the villain, she should be regarded as seeking to enhance the *chance* of apprehending him. Once this change of outlook is adopted, it becomes apparent that the villain has a course of action that secures a 50% chance of escape and the sleuth has a course of action that secures a 50% chance of apprehending him. And, concomitantly, no actions exist that guarantee the respective agents a greater chance of achieving their goals. Seven years later, John von Neumann [1928] proved the Minimax Theorem, effectively generalizing this result to all 2-person zero-sum games. The key to this breakthrough lay in allowing the interacting agents to *randomize*, i.e., to use mixed strategies.

About forty years after the successful resolution of the impasse besetting zero-sum games, a new conundrum appeared in the intellectual landscape, when Robert Nozick [1969] brought up Newcomb's problem. Nozick asked us to consider the following choice situation: You see before you two boxes, one transparent and one opaque, and you need to decide whether to take the opaque box alone or to take both boxes. You can see clearly that the transparent box contains $1,000. You can't see what's in the opaque box, but you know that it contains either nothing at all or $1million. Which of the two possibilities it is – 0 or $1M – is determined by a being who had either placed $1M in the (otherwise empty) opaque box or had refrained from doing so. The being's rule regarding its action is to place the money in the box if, and only if, it predicts that you will pick the opaque box alone. The being is known to possess excellent predictive ability regarding the choices that people make and, in particular, regarding your own upcoming choice. The money you find in the box(es) you pick is yours to keep. What should you do? Dominance (or Savage's [1954, p. 21] "Sure Thing Principle") dictates that you must pick *both* boxes, regardless of anything else. Yet if you believe that the being's predictive powers are even moderately better than chance[2], you should pick the opaque box *alone*, to obtain the higher expected payoff. What are you (we) to do?

Rivers of ink flowed in the streets of Academia in attempts to resolve this difficulty (see, e.g., Campbell and Sowden, eds. [1985]) or to explain it away. Here, I shall argue that just as in the case of two-person zero-sum games, the key to resolving the paradox lies in allowing the two interacting agents to randomize. Now, in his original paper, Nozick [1969, footnote 1] explicitly *forbids* the interacting agents' use of randomization. But this decree turns out to be a mistake. It's like telling the villain and the sleuth, above, that they are not allowed to randomize (it is acknowledged, though, that the roles of randomization in the two settings are quite different).

[2] with $1million denoted M and $1,000 denoted 1, we find that (M+1)/2M is the probability of correct prediction above which picking the opaque box alone yields the higher expected payoff.

2. The Setting

Let us consider two interacting agents, one to be known as "nozick" and the other as "the being". Each agent must choose one of two available actions. The following table specifies nozick's payoffs, given the two agents' actions:

		the being	
		IN	**NO**
nozick	**OP**	M	0
	BO	M+1	1

where M stands for a million dollars and 1 stands for a thousand dollars; **OP** stands for "take the *opaque* box alone" and **BO** for "take *both* boxes"; **IN** stands for "million *in* box" and **NO** for "*no* money in box".

The action **BO** strongly dominates the action **OP** so, under conventional decision theory, nozick should pick **BO**, regardless of what the being chooses to do. But this is not conventional decision theory, because the two agents' actions are taken to be **correlated**. Specifically, the more likely nozick is to pick **OP**, the more likely the being is to pick **IN**. This "pushes" the interaction towards the main diagonal of the array above, where choosing **OP** yields M while choosing **BO** yields just 1. What should nozick do?

This is the problem which Robert Nozick (the philosopher) posed in 1969, citing William Newcomb. A few years before that, Richard Jeffrey [1965, pp.8-9] had posed what was essentially the same problem, where a country needs to decide between arming and disarming when its decision is known to be correlated with whether war will break out or not.

Let us re-state Newcomb's problem as follows:

- nozick sends in a coin that falls Heads with probability p and Tails with probability $1-p$.

- the being sends in a coin that falls Heads with probability q and Tails with probability $1-q$.
- the two coins are positively **correlated**, with the correlation coefficient given by r $(0 \leq r \leq 1)$.
- nozick picks p rationally and the being picks q arbitrarily.
- referee tosses the two coins and implements the agents' moves, as follows: nozick's action is **OP** or **BO** according as his coin falls Heads or Tails, and the being's action is **IN** or **NO** according as its coin falls Heads or Tails; these outcomes determine nozick's payoff, in accordance with the table above.
- all the above is common knowledge.

Recall that Nozick the philosopher had actually tried to restrict nozick the agent to choosing $p=0$ or 1, by decreeing that any attempt by nozick to pick p with $0<p<1$would automatically nail the being's action at **NO**. But this decree is untenable, because, as we shall see, restricting p to 0 or 1 is mathematically incompatible with a positive correlation, $r>0$. And since positive correlation is the crux of the story, we *must* allow both agents to randomize.

Apart from the fact that here the two interacting agents are allowed to randomize, the formulation above is precisely Newcomb's problem, as laid down by Robert Nozick in 1969: Which of two actions, of which one payoff-dominates the other, should you pick, when your partner's action is not independent of yours? This question is meaningful and well-stated regardless of how the two agents' actions might have come to be mutually dependent (correlated). Thus, providing a satisfactory answer to this question would "resolve" Newconb's paradox, **regardless** of how the agents' actions might have come to be mutually dependent. The aim in what follows is indeed to offer a satisfactory answer to the above question.

3. Analysis

Given the three parameters, p, q, and r, i.e. given the choices made by nozick and the being, and given the correlation coefficient, the joint distribution of the two Bernoullian random variables (the two agents' respective coins) is completely specified, as follows:

	IN	**NO**
OP	$pq+r\Delta$	$p(1-q)-r\Delta$
BO	$(1-p)q-r\Delta$	$(1-p)(1-q)+r\Delta$

where $\Delta=(pq(1-p)(1-q))^{1/2}$. (See Appendix, Sec.5, below.)

Using this distribution, we can calculate what nozick's expected payoff would be in this interaction, given the values of the three parameters. This leads to

$$E(nozick's\ payoff)=1-p+qM, \quad independently\ of\ r.$$

If nozick's expected payoff is given by this formula, independently of the correlation coefficient, does this not mean that he can proceed as though $r=0$? Shouldn't nozick simply go ahead and select the dominant action **BO**? Indeed, with expected payoff given by $1-p+qM$, the maximum is clearly at $p=0$, i.e. **BO** should be picked regardless of anything else. Well, this last conclusion ($p=0$) turns out to be wrong. The fact that expected payoff is given by $1-p+qM$ does indeed mean that in order to maximize, nozick needs to set p as low as possible. But with $r>0$, the lowest possible p is not 0.

The probabilities p and q (i.e., nozick's and the being's actions) are constrained by the need to ensure that the probability numbers appearing in the foregoing 2×2 table are nonnegative. In the case where $r\geq 0$, the offdiagonal numbers are not immediately guaranteed to be nonnegative, so we need to impose the two conditions

$$p(1-q)-r\Delta \geq 0 \quad and \quad q(1-p)-r\Delta \geq 0$$

where, as has already been noted, Δ is given by $\Delta = (pq(1-p)(1-q))^{1/2}$. Simplifying, these two conditions reduce to –

$$(\boldsymbol{A}) \quad \frac{p(1-q)}{q(1-p)} \geq r^2 \quad and \quad (\boldsymbol{B}) \quad \frac{q(1-p)}{p(1-q)} \geq r^2$$

Let a set S_r be defined as follows:

$$S_r = \{(p,q) \mid 0 \leq p, q \leq 1, (\boldsymbol{A}) \; and \; (\boldsymbol{B}) \; hold\}$$

S_r is the *feasible set* of probability pairs which nozick and the being can jointly select, when the correlation coefficient is r. Here are some of the properties of this feasible set, S_r:

(i) S_r *is convex*;

(ii) *if* $(p,q) \in S_r$ *then* $(q,p) \in S_r$

(iii) *if* $(p,q) \in S_r$ *then* $(1-p, 1-q) \in S_r$

(iv) *if* $r \geq r'$ *then* $S_r \subset S_{r'}$

(v) S_r *converges to the main diagonal as* $r \to 1$

(vi) S_r *converges to the unit square as* $r \to 0$

(vii) *for* $r > 0$ *and* $(p,q) \in S_r$ *if* $p = 0$ *then* $q = 0$

Graphically, S_r looks like this:

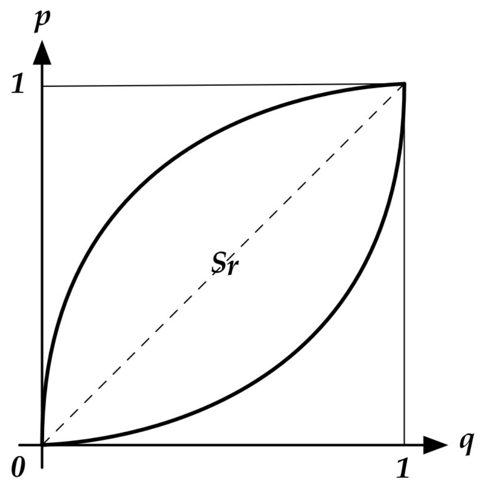

As we have already seen, in order to maximize his expected payoff, nozick needs to set the value of p (i.e., the parameter of his coin) as low as possible. That is, he must locate himself on the *lower boundary* of the set S_r. This lower boundary is given by the equation $p(1-q) = q(1-p)r^2$, so the solution to nozick's decision problem is to set p at $p*$, given by

$$p* = \frac{qr^2}{1-q(1-r^2)}$$

where q is the being's action (i.e., its probability choice). We see that, for $r > 0$, $p* \to 1$ as $q \to 1$ and $p* \to 0$ as $q \to 0$. In other words, when the probability picked by the being (for placing the money in the opaque box) is high, nozick in turn is very likely to pick just the *one* (opaque) box. And when the being chooses, with high probability, to leave the opaque box empty, nozick is very likely to pick *both* boxes. Note that, by definition, nozick's optimal choice, $p*$, is such that at $p*$ the equation $p(1-q) - r\Delta = 0$ holds. At $p*$, the event (**OP**, **NO**), of nozick picking the opaque box alone and finding it empty, *has probability zero*. His action, $p*$, protects nozick completely against the dreaded event of finding himself "holding the bag", with no money whatsoever.

All this is just a consequence of nozick's desire to keep away, as much as possible, from the dominated strategy, by setting its probability as low as is feasible. Note also that nozick's expected payoff is close to M when q is high and close to 1 when q is low. In other words, nozick's fleeing from the dominated strategy actually pushes the interaction towards the diagonal entries in his payoff matrix. Correlation allows nozick to be a one-boxer when one-boxing is advantageous and a two-boxer when two-boxing is advantageous. No conflicting principles of choice and no paradox.

4. Conclusion

The need to grapple with the issues brought up in Robert Nozick's 1969 paper on Newcomb's Problem gave rise to literally dozens of subsequent contributions. Much of this ensuing literature carried the discussion of Newcomb's Problem far afield, notably towards exploring the nature of causation and causality. For Nozick himself, however, Newcomb's Problem was a matter largely confined to Decision Theory and the philosophical elucidation of human choice among available alternatives, with expected utility maximization seemingly coming into conflict with one of its own axioms (monotonicity). He thought, I am sure, that the problem he was posing was hard enough as is, without inflicting upon his readers the further punishment of having to deal with randomized choices (or mixed strategies). So, he simply ruled out randomization altogether, right at the outset. This move turned out to be a major stumbling block. The key to understanding the workings of an interactive system where agents' actions are correlated lies precisely in allowing these agents to randomize. Once this is done, one finds, in Newcomb's story, that if the being wishes to coax its partner-in-interaction into picking the opaque box alone, all it needs to do is to see to it that the probability of the money actually going into that box is high. Correlation – even low correlation – will do the rest. In the limit, as the probability of the money going into the opaque box approaches 1, so does the probability of the chooser picking just that one box. And this is true without the so-called

Dominance Principle being in any way flouted or abandoned. The dominant strategy is in fact being selected, with as high a probability as possible.

* * *

This essay is dedicated to the memory of Robert Nozick. Eleven years after his untimely passing, many of those who knew him still miss him, as friend, communicator, and wise interlocutor. Right now, I find myself relishing the thought of how, if he were still with us, he would tear into this very piece.

5. Appendix: Correlated Coins

Let x and y be two Bernoulli random variables, with x taking the values 1 and 0 with probabilities p and $1-p$, and y taking the values 1 and 0 with probabilities q and $1-q$. As is well known (and easily seen), we have:

$$E(x) = p, \quad E(y) = q$$

$$\sigma_x = (p(1-p))^{1/2}, \quad \sigma_y = (q(1-q))^{1/2}$$

where the σ's are the *standard deviations* of the two variables. (In Sec. 3, above, the symbol Δ was used to denote the *product* of σ_x and σ_y.)

The coefficient of **correlation**, r, between x and y is defined by

$$r = \frac{\mathrm{cov}(x, y)}{\sigma_x \sigma_y}$$

with the **covariance,** $\mathrm{cov}(x, y)$, being defined as follows:

$$\begin{aligned}\operatorname{cov}(x,y) &= E((x-E(x))(y-E(y))) \\ &= E(xy)-pq.\end{aligned}$$

Let us write down the joint distribution of x and y:

	$y=1$	$y=0$
$x=1$	π_{11}	π_{10}
$x=0$	π_{01}	π_{00}

and, using this distribution, let us calculate $E(xy)$:

$$E(xy) = 1\cdot\pi_{11} + 0\cdot\pi_{10} + 0\cdot\pi_{01} + 0\cdot\pi_{00} = \pi_{11}.$$

But, from the definition of the covariance, we also have:

$$\begin{aligned}E(xy) &= \operatorname{cov}(x,y)+pq \\ &= r\sigma_x\sigma_y + pq,\end{aligned}$$

so, the upper-left entry, π_{11}, of the joint distribution matrix is given by $pq + r\Delta$, as stated in Sec. 3, above. Verifying the remaining entries is straightforward.

6. References

Campbell, R. and L. Sowden, eds., *Paradoxes of Rationality and Cooperation*, Vancouver, UBC Press, 1985.

Jeffrey, R.C., *The Logic of Decision*, New York, McGraw Hill, 1965, (also 2nd ed., 1983).

von Neumann, J., "Zur Theorie der Geselschaftsspiele", *Mathematische Annalen*, 100, 295-320, 1928 (translated: "On the Theory of Games of Strategy", in A.W. Tucker and R.D. Luce, eds., *Contributions to the Theory of Games*, 4, 13-42, NJ, Princeton University Press, 1959).

Nozick, R., “Newcomb’s Problem and Two Principles of Choice”, in Rescher N. *et al*, eds., *Essays in Honor of Carl G. Hempel*, Dordrecht, Reidel, 1969 (reprinted in Campbell and Sowden, above).

Savage, L.J., *The Foundations of Statistics*, New York, Wiley, 1954.