

Safeguards of a Disunified mind

Wlodek Rabinowicz

Subject:

Pragmatic arguments for rationality constraints on a decision maker's state of mind (on her beliefs and desires).

An argument of this kind appeals to the costs/benefits of a given constraint's violation/satisfaction.

The form of a pragmatic argument

The argument shows that a violator of a given constraint can be made to act to her guaranteed disadvantage.

Dramatically put: She can be exploited by a clever bookie, even if the latter lacks superior knowledge.

Frank Ramsey, “Truth and Probability” (1926: 1990), p. 78:

If anyone's mental condition violated these laws, ... [h]e could have a book made against him by a cunning bettor and would then stand to lose in any

Examples of pragmatic arguments

Synchronic Dutch Books – standard probability laws

Diachronic Dutch Books

– Reflection, Conditionalization

Money Pumps – acyclicity of preference

A common feature

Suppose an agent is (i) logically competent.
And (ii) prefers to be better off rather than worse off.

If such an agent violates a given constraint,
then she can be exploited *only if*
she addresses different choice problems separately
rather than jointly,
i.e., only if she is *disunified* in her decision-making.

(Note: Disunification in the *object* of decision-making,
not necessarily in the agent.)

Example: Violation of probability laws

Assumed: *Betting interpretation of probabilities.*

Basic idea: Probabilities are guides to action.

They are measures of betting dispositions/betting commitments.

Probability equals the agent's betting rate:

$P(X)$ = the betting rate for X

i.e., the ratio between the price and the stake in a fair bet on X

A bet is fair iff the agent is willing to take each of its sides.

Assumed: the stake-price ratio is constant for all fair bets on a given X

Note: Subjectivists interpret probabilities as *degrees of belief*.

Violation of the addition axiom for probabilities

A, B – logically incompatible propositions.

$$P(A) = \frac{1}{2} \quad P(B) = \frac{1}{2} \quad P(A \text{ or } B) = \frac{3}{4}$$

Dutch Book:

We *sell* to the agent a bet on A and another bet on B , each with the stake \$4 and the price \$2, and

we *buy* from her a bet on the disjunction A -or- B with the same stake \$4 and the price \$3.

Our guaranteed profit: \$1 (= price difference: $2 + 2 - 3$).

Objection (Schick 1986):

A logically competent agent,
who prefers to be better off rather than worse off,
is not going to accept all the three bets,
if she takes a decision *on the whole bet package*
rather than separately on each component.

She will realize that together they lead to a sure loss.

BUT: Even if a unified agent won't be exploited, doesn't the very possibility of a Dutch book show that there's something wrong with the agent?

She finds each bet in the package attractive, but assigns a negative value to the package as a whole.

Doesn't this mean that *she evaluates one and the same betting arrangement differently under different but logically equivalent descriptions?* (Skyrms 1980, p. 119)

NO. It only shows that her evaluations are not additive.

[In fact, not even weakly separable.]

But why should they be? (Schick 1986)

Complementarity, principle of organic unities.

Violations of preference acyclicity – Money Pump

The agent's preferences form a cycle: $x \prec y \prec z \prec x$.

Suppose she holds x and is invited to trade x for y , then y for z , and finally z for x .

I.e., each time she is offered something she prefers to what she holds at the time the offer is made.

Each trade costs her a small amount ϵ , which does not reverse her preference.

After three trades, she is back to x , minus 3ϵ .

Again, objection:

If the agent were to make just *one* decision on her whole course of action, she would not be pumped.

Cf. McClennen's idea of *resolute choice*.

Conclusion: Vulnerability to exploitation in the Money Pump is based on diachronic disunification.

Query: But is diachronic unification really *necessary* if the pump is to be avoided?

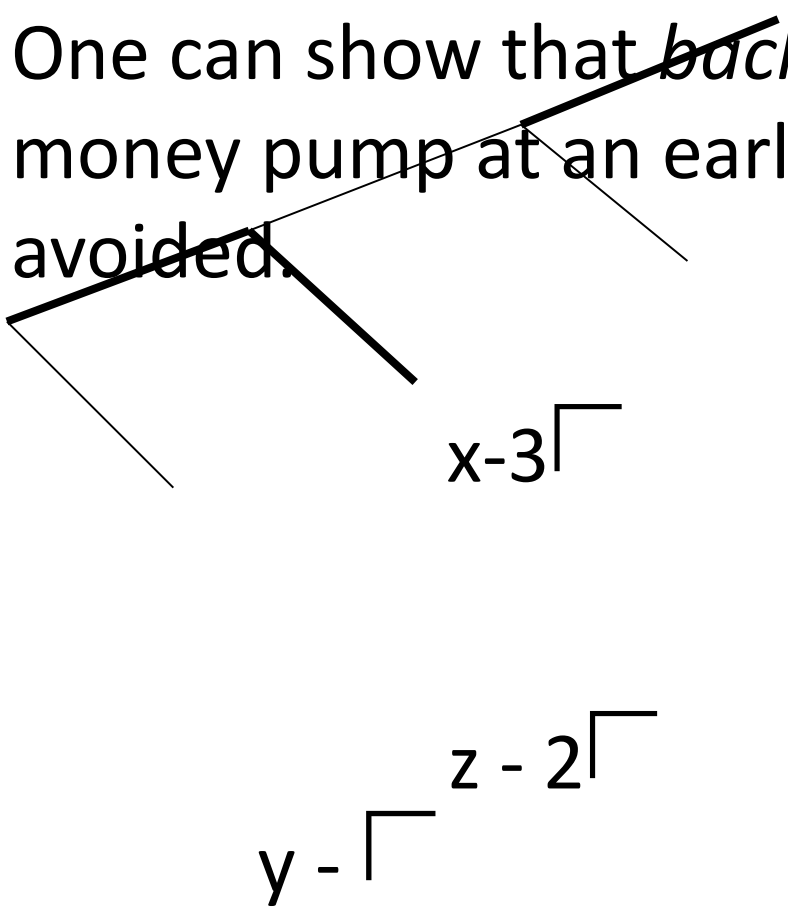
Isn't *foresight* enough?

(Schick 86, Schwartz 86, McClennen 90, Rabinowicz 95)

Symmetry in knowledge between the agent and the exploiter.

Advantages of foresight

One can show that *backward induction* stops the money pump at an early stage. Exploitation is avoided.



$x - 3\Gamma$

$z - 2\Gamma$

$y - \Gamma$

x

BUT:

Foresight is not enough if the exploiter is persistent.

(Rabinowicz 2000)

A persistent exploiter *comes back* with the offer the agent has rejected at the previous stage.


Thus, a refusal to accept an exchange does not terminate the interaction; it does not get one off the hook.

In this more complicated setting, with three stages, backward induction prescribes trading at each stage.

The series of trades leads the agent back to where she started, minus $3\sqrt{\quad}$.

So, the Money Pump does work, *despite* foresight.

$$x < y < z < x$$



Violations of *Reflection* - Diachronic Dutch book

Let P be the agent's probability at t and P' her probability at $t' \succ t$.

Reflection: $P(A/P'(A) = k) = k$,
provided that $P(P'(A) = k) > 0$.

Intuition: One's current probabilities conditional on one's future probabilities should reflect the latter. (Epistemic self-trust)

Let $E = P'(A) = k$

van Fraassen 1984: A violator of Reflection is vulnerable to a diachronic Dutch book

As before:

- (i) For exploitation, we need to assume diachronic disunification.
- (ii) Given diachronic disunification, foresight isn't enough to get the agent off the hook.
(pace Levi 1988 and Maher 1992)

Skyrms 1993:

Even with foresight, the Dutch book does work *if the exploiter is persistent*, in the sense that the potential bet offer at t' is not dependent on the acceptance of bets at t .

From the *earlier* perspective, the bet to be offered at t' is disadvantageous to the agent.

But she knows she is going to accept it at t' if it will be offered (i.e. if $P'(A) = k$).

For then that bet will be fair in terms of the agent's probabilities at t' .

She would prefer to refuse earlier bets, at t , if this would prevent the future bet offer.

But since this *wouldn't* prevent it (the bookie is persistent), she has no reason to abstain from the earlier bets. This makes the Dutch book work.

Isaac Levi's critique of diachronic exploitation set-ups

Levi (2002) : Diachronic set-ups, unlike the synchronic ones, are unfit for legitimate pragmatic arguments.

In a correct pragmatic argument, the exploited agent has at her disposal some option whose outcome dominates (= is better than, however things develop) the exploitation outcome. (Presence of a dominating option.)

This is the case in synchronic Dutch books.

Refusal to take any bets at all and thus to avoid any losses is an option available to the agent if the bets are offered at the same time. If she nevertheless gets exploited, she must be irrational.

But in diachronic set-ups, there is no feasible option that dominates the exploitation outcome.

Examples

In Money Pump, refusing from the outset all the future trades is not a feasible option (in the absence of resoluteness).

Similarly in the case of the violation of Reflection:
Refusing the two bets offered now, at t , *and* the one that might be offered in the future, at t' , is not a feasible option at t (again, in the absence of resoluteness).

Reply to Levi's critique: We could change the examples.

Suppose we consider cases in which a wholesale choice *is* available to the agent in the diachronic set-up.

I.e., she is able to make a resolute choice.

Though she *actually* does not and instead decides in a disunified way.

This puts the synchronic and the diachronic arguments on the same footing.

In both cases, the exploitation takes place only because the agent makes a decision on each component in the package separately rather than jointly.

In both cases, making a joint decision is a feasible option.

There is still this *disanalogy*:

In the diachronic case, practical deliberation on earlier transactions (bets/trades) is allowed to take into consideration predictions about later transactions.

In the synchronic case, each transaction is decided without consideration given to other transactions.

But does this disanalogy show that diachronic arguments are less legitimate?

If anything, it should rather be the opposite.

If a violator of a constraint can be exploited even if she considers her other transactions, as in the diachronic case, then this seems to be a more serious problem, not a less serious one.

Question: Could disunified decision making *as such* make one vulnerable to foreseeable exploitation, *even if one does not violate any (other) rationality constraints?*

If so, then it wouldn't be appropriate to view the standard pragmatic arguments as ways to defend different constraints on beliefs or desires.

Instead, they would more naturally be interpretable as different arguments in favour of unification in decision-making.

To consider this issue, let's dramatize the situation:
Let's represent disunification as a situation involving a *group* of agents making independent decisions on different issues, while sharing goals and assets.

Group members: Bayesians with common priors and full trust in each other's epistemic capacities.

⤴ Common knowledge in the group.

*Can a bookie exploit such a group,
even if he doesn't know more than any group
member?*

THE STORY OF THE HATS

Common knowledge in a group of three persons:

- Each person will be given a hat to put on in the dark (Stage 1).
- Each hat's colour, either black or white, will be decided
in secret, independently and at random.
- After the lights are turned on (Stage 2), each person will
see the hats of the other two persons, but not her own.

The three persons are Bayesians with common priors.

Consider the proposition:

(A) The three hats are not of the same colour.

Stage 1: Each person assigns probability $\frac{3}{4}$ to A.

[A is true in 6 out of 8 possible cases.]

Stage 2:

– If A is *true*, two persons assign probability 1 to A, while one person assigns $\frac{1}{2}$.

– If A is *false*, each person assigns probability $\frac{1}{2}$ to A.

NOTE: Whether A is true or not,

Stage 1: The bookie offers to *sell*, to the first comer, a bet on A with a price \$3 and a stake \$4.

For every i in the group, $P_{i,1}(A) = \frac{3}{4}$. The bet is fair for every i .

Stage 2: The bookie offers to *buy*, from the first comer, a bet on A with a price \$2 and a stake \$4.

For at least one i , $P_{i,2}(A) = \frac{1}{2}$. For that i , the bet is fair.

Whether A is true or not, the bookie's net profit is \$1 (3 - 2).

Due to the foreseeable changes in probabilities, the bookie who sells a bet on A can count on buying it back later, at a lower price.

👤 A successful diachronic Dutch book against the group

BUT: If your probability for A at stage 2 is $\frac{1}{2}$, should you declare yourself willing to sell a bet on A at the rate $\frac{1}{2}$?

Either (i) A is true and you are the only person with this probability,

or (ii) A is false and the other two persons also assign probability $\frac{1}{2}$ to A .

Selling the bet on A is *disadvantageous* to the group in case (i) and *advantageous* only in case (ii).

Still, in case (ii), the other group members should also be willing to sell. So why not defer to them and keep silent?

But, on the other hand, what is their reason likewise?

Question: Is there a symmetrical Nash equilibrium for this game among group members? Yes: Everyone keeps silent.

Which means that the Dutch book argument is spurious!

Levi (2006): Disunification *is* always irrational if it is not inescapable, i.e. if a decision on the whole package of opportunities is available.

Why? Because under such circumstance disunification involves an omission to consider an available option.

Which is irrational.

If Levi is right then pragmatic arguments no longer can provide support for the different constraints on beliefs and preferences they purport to target.

Instead they are all arguments in favour of unification.

But this requirement to consider every available option is too stringent, in my view.

There are just *too many* available options to consider.

We *can* require that the set of options the agent considers is not too meagre – that it obeys *some* richness conditions.

One such plausible condition is the options considered should be *jointly exhaustive*.

Another condition that seems plausible is *closure under dominance*:

If an option considered is dominated by some available options,
then these dominating options, or at least some of them,
should also be considered.

But: If the agent is disunified in dealing with a package of opportunities, then the option to refuse all the opportunities in

Suggestion:

The real purpose of pragmatic arguments is to identify constraints we need to satisfy if we want to engage into disunified decision-making without making ourselves vulnerable to exploitation.

I.e., their purpose is to identify *safeguards of a disunified mind*.

Thus, such arguments do not establish the inherent rationality of the constraints.

They result in a *conditional* advice:

To put it differently:

Pragmatic arguments identify constraints that a disunified decision maker needs to satisfy in order to be *practically coherent* in his actions.

Cf. **Ramsey** (1926: 1990, p. 79, my italics):

“Having degrees of belief obeying the laws of probability implies a further measure of *consistency*, namely *consistency between the odds* acceptable on different propositions as shall prevent a book being made against you.”

Disunified decision-making is of great practical value.

This applies, in particular, to decision-making *over time*.

Diachronic unification is often difficult (requires foresight, resoluteness, and/or costly pre-commitments) and sometimes it is impossible.

Synchronic unification, on the other hand, is very easy.

Which makes diachronic pragmatic arguments more compelling than the synchronic ones, contrary to the received view.