A Theory of Rational Action

John L. Pollock
Department of Philosophy
University of Arizona
Tucson, Arizona 85721
pollock@arizona.edu
http://www.u.arizona.edu/~pollock

(This work was supported by NSF grant no. IRI-IIS-0080888.)

A theory of rational action is a theory of rational choice.

It is a theory of how an agent should, rationally, go about deciding what actions to perform at any given time.

It is assumed that these decisions must be made in the face of uncertainty regarding both the agent's initial situation and the consequences of his actions.

Most discussions of rational choice proceed against the background of classical decision theory, which is generally assumed uncritically.

This paper argues that although classical decision theory is based upon some sound intuitions, the details are badly flawed.

It is a difficult task to construct a modified theory that gets the details right, and the resulting theory differs from classical decision theory in major respects.

1. Classical Decision Theory

set A of *alternative actions*

set O of pairwise exclusive and jointly exhaustive outcomes

$\mathsf{PROB}(O/A)$

utility-measure $\mathbf{U}(O)$

$$\mathbf{EV}(A) = \sum_{O \in \mathrm{O}} \mathbf{U}(O) \cdot \mathsf{PROB}(O/A).$$

2. Action Omnipotence

Classical decision theory assumes that actions can be infallibly performed.

Choosing between running in the Boston Marathon or staying home and watching the race on TV.

I will only enjoy running in the race if I win it, but that would be much better than watching the race on TV.

I could win the race by running 100 mph for the duration of the race.

2.1 High-Level Actions

To run 100 mph for the duration of the race I must perform a whole sequence of (simpler) actions over an extended period of time.

Actions exhibit a continuous range of abstractness:

wiggle your finger
walk across the room
make a cup of coffee
vacation in Afghanistan
save the world.

2.2 Basic Actions

Restrict the dictates of decision theory to low-level actions?

Most actions are performed *by* performing other lower level actions.

> *Basic actions* are actions that can be performed without performing them by performing lower level actions.

Even basic actions are not infallibly performable.

> your finger has been injected with a muscle paralyzer

2.3 Trying

At one time I thought that although we cannot always perform an action, we can always try, and so classical decision theory should be restricted to tryings.

Try to paint a wooden block.

The ability to try is influenced by your beliefs.

2.4 Expected-Utility

The dictates of classical decision theory would only be reasonable for actions that are infallibly performable, but there do not appear to be any.

Reformulate the theory so that it applies to ordinary actions but takes account of the failure of action omnipotence.

Action omnipotence fails in two ways:

- we may fail to perform an action when we try
- we may not even be able to try.

Discount the expected-value of performing an action by the probability that we will be able to perform it if we try?

$$\mathbf{EV}(A)\cdot \mathbf{PROB}(A \,/\, try\text{-}A)$$

(1) the values of goals
(2) execution costs
(3) side-effects

Assess actions in terms of the expected-value of trying to perform them?

EV(*try-A*)

If we know that we cannot try to perform one of them, then it should not be a contender in the choice.

We may be uncertain whether we will be able to try

war game

Discount the expected-value of trying to perform an action by the probability that, under the present circumstances, we can try to perform it?

$$\mathbf{EV}(\textit{try-A})\cdot\mathrm{PROB}(A \,/\, \textit{can-try-A})$$

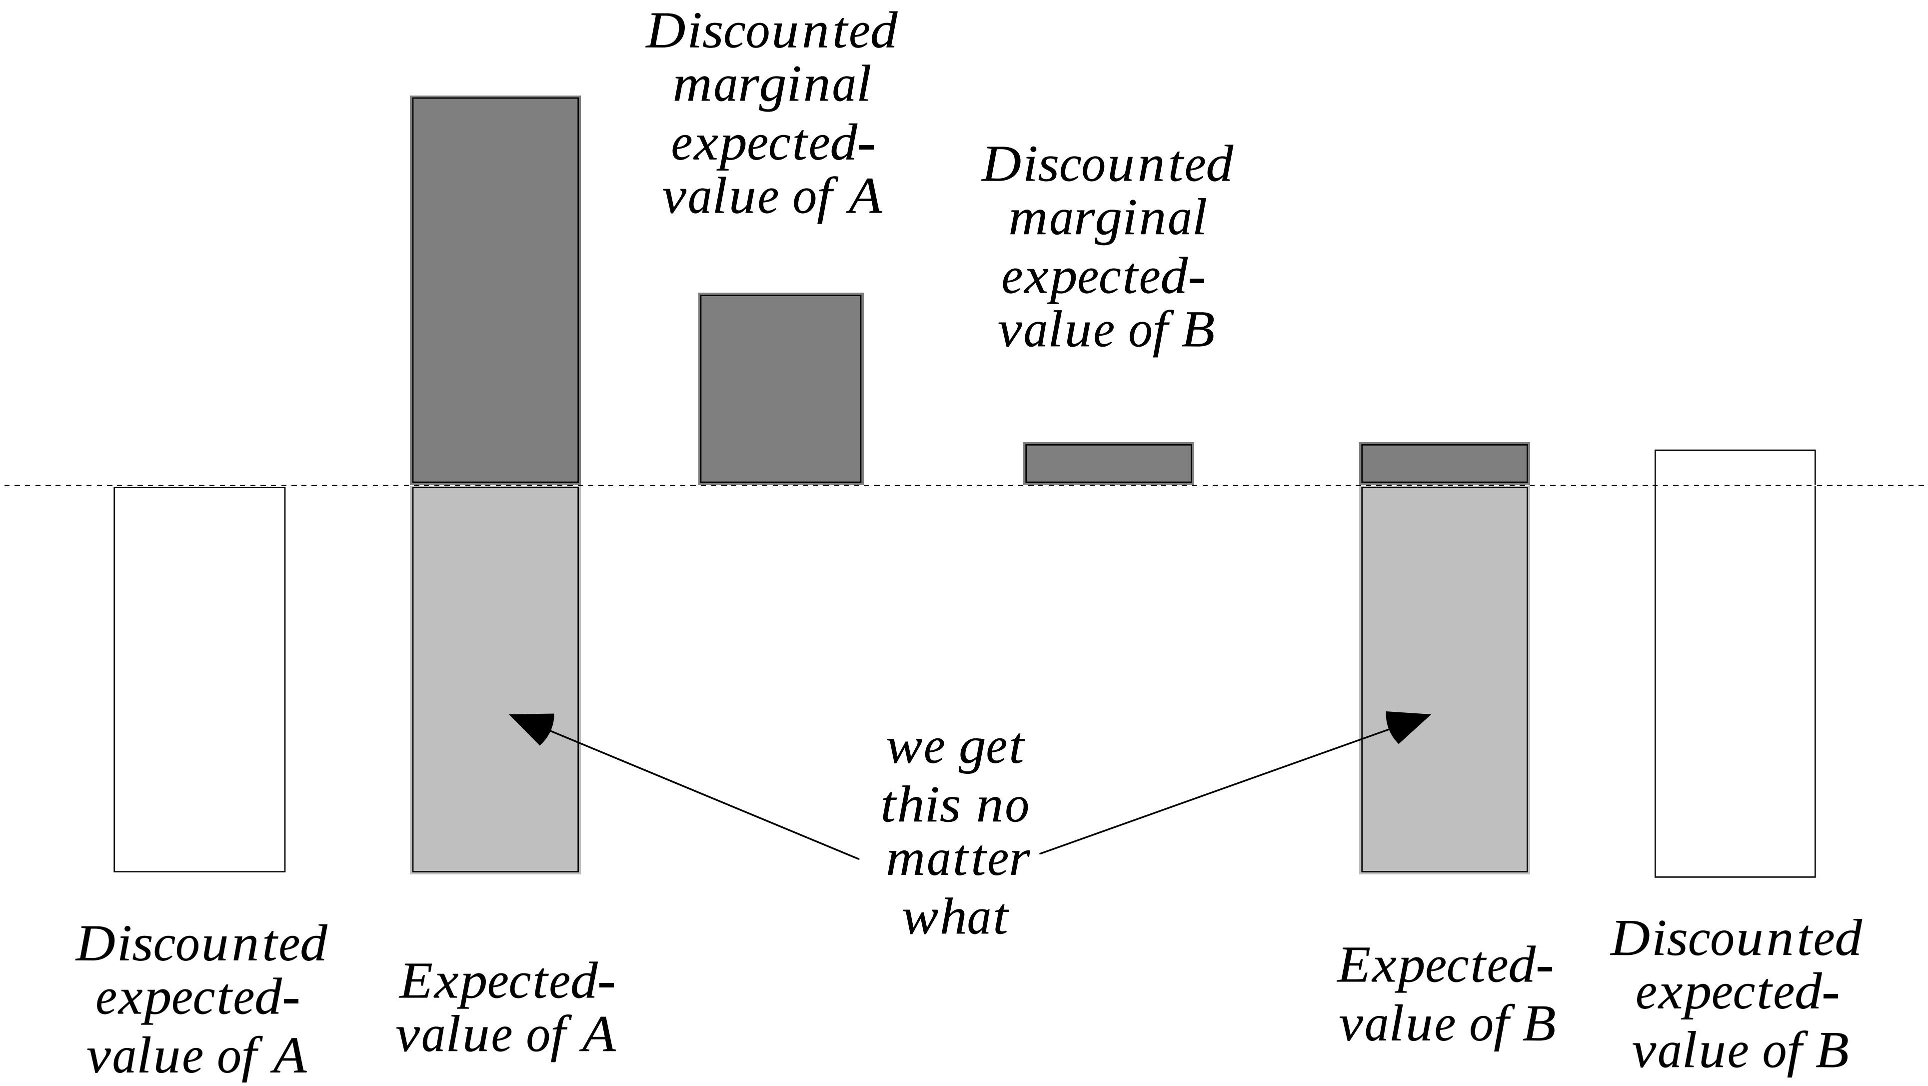
Discounted marginal expected-value of A
Discounted marginal expected-value of B
we get this no matter what
Discounted expected-value of A
Expected-value of A
Expected-value of B
Discounted expected-value of B

We should be comparing the *extra amounts* we will get, over and above the ten dollars we will get no matter what , and discounting those extra amounts by the probabilities of being able to try to perform the actions.

marginal expected-value of trying to perform A = the difference between the expected-value of trying to perform *A* and the expected-value of not trying to perform any of the alternative actions (*nil* — the "null-action").

$$\mathsf{PROB}(\mathit{can\text{-}try\text{-}A})\cdot[\mathbf{EV}(\mathit{try\text{-}A}) - \mathrm{EV}(\mathit{can\text{-}try\text{-}A})]$$

We can have a case in which all of the value attached to *try-A* derives from its being the case that the agent *can-try-A*.

There is no point to actually trying *A*.

This is captured by considering the difference between the expected-value of trying to perform *A* and the expected-value of not trying to perform any of the alternative actions *given that the agent can try to perform A*.

$$\textbf{expected-utility}(A) = \mathrm{PROB}(\textit{can-try-A})\cdot[\mathbf{EV}(\textit{try-A}) - \mathrm{EV}(\textit{can-try-A})]$$

Ad hoc?

Equivalent to comparing the expected-values of “conditional policies” of the form *try to do A if you can try to do A*.

3. Conditional Policies

conditional decisions about what to do if some condition *P* turns out to be true.

EV(*A if P*) = **PROB**(*P*)·**EV**(*A* / *P*) + **PROB**(~*P*)·**EV**(*nil* / ~*P*).

Define **MEV**(*A if P*) = **EV**(*A if P*) – **EV**(*nil if P*).

MEV(*A* / *P*) = **EV**(*A* / *P*) – **EV**(*nil* / *P*).

The conditional policy *nil if P* prescribes doing *nil* if *P* and *nil* if ~*P*, so it is equivalent to *nil* simpliciter. Thus we could just as well have defined:

MEV(*A if P*) = **EV**(*A if P*) – **EV**(*nil*).

TH 1: MEV(*A if P*) > **MEV**(*B if Q*) iff **EV**(*A if P*) > **EV**(*B if Q*).

TH 2: $\mathbf{MEV}(A \textit{ if } P) = \textsf{PROB}(P)\cdot\mathbf{MEV}(A / P)$

Proof:

$$\mathbf{MEV}(A \textit{ if } P)$$

$$= \mathbf{EV}(A \textit{ if } P) - \mathbf{EV}(nil \textit{ if } P)$$

$$= \textsf{PROB}(P)\cdot\mathbf{EV}(A / P) + \textsf{PROB}(\sim P)\cdot\mathbf{EV}(nil / \sim P)$$

$$- \textsf{PROB}(P)\cdot\mathbf{EV}(nil / P) - \textsf{PROB}(\sim P)\cdot\mathbf{EV}(nil / \sim P)$$

$$= \textsf{PROB}(P)\cdot\mathbf{EV}(A / P) - \textsf{PROB}(P)\cdot\mathbf{EV}(nil / P)$$

$$= \textsf{PROB}(P)\cdot[\mathbf{EV}(A / P) - \mathbf{EV}(nil / P)]$$

$$= \textsf{PROB}(P)\cdot\mathbf{MEV}(A / P). \quad \blacksquare$$

We have:

$$\textbf{expected-utility}(A) = \textsf{PROB}(\textit{can-try-A})\cdot[\textbf{EV}(\textit{try-A}) - \textbf{EV}(\textit{can-try-A})]$$
$$= \textsf{PROB}(\textit{can-try-A})\cdot\textbf{MEV}(\textit{try-A} \;/\; \textit{can-try-A})$$
$$= \textbf{MEV}(\textit{try-A if can-try-A}).$$

Thus

TH 3:

$$\textbf{expected-utility}(A) > \textbf{expected-utility}(B)$$
$$\text{iff}$$
$$\textbf{EV}(\textit{try-A if can-try-A}) > \textbf{EV}(\textit{try-B if can-try-B}).$$

4. Plan-Based Decision Theory

Classical decision theory has us focus on actions individually and choose them one at a time.

Transport a ton of silver and a ton of gold from one location to another in a one-ton truck.

> transport the gold on a single trip
>
> transport the silver on a single trip
>
> transport both on a single trip

Imagine the probabilities and utilities to be such that the action with the highest expected-utility is that of transporting both on a single trip, even though that risks damaging the springs.

If I have time to make two trips, that might be the better choice.

Sometimes decision-theoretic choices must be between groups of actions, and the performance of a single action becomes rational only because it is part of a group of actions whose choice is dictated by practical rationality.

Groups of actions are *plans*.

How should we choose plans?

Obvious proposal — apply classical decision theory to plans.

Define the *expected-utility of a plan* to be the marginal expected-value of trying to execute the plan if you can try.

Competing plans should be plans that we must choose between, rather than adopting both.

> *strong competition* — impossible to execute both

Deciding what to cook for dinner — roast chicken or barbecue lamb chops.

Two plans *compete weakly* iff the plan that results from merging the two plans into a single plan has a lower expected-utility than one of the original plans.

> It is rational to adopt (decide to execute) a plan iff it has no weak competitor with a higher expected-utility.

5. A Fundamental Problem for Plan-Based Decision Theory

Does there exist (in logical space) a competing plan with a higher expected-utility? — Almost always.

EU(roast chicken) > EU(barbecue lamb chops)

EU(barbecue lamb chops + go to movie) > EU(roast chicken)

EU(roast chicken + go to movie) > EU(barbecue lamb chops + go to movie)

... *and so on.*

Given two competing plans P_1 and P_2, if the expected-utility of P_1 is greater than that of P_2, the comparison can generally be reversed by finding another plan P_3 that pursues unrelated goals and then merging P_2 and P_3 to form P_2+P_3.

If P_3 is well chosen, this will have the result that P_2+P_3 still competes with P_1 and the expected-utility of P_2+P_3 is higher than the expected-utility of P_1.

If this is always possible, then plan-based decision theory implies that it is not rational to adopt any plan.

It might be objected that P_2+P_3 is not an appropriate object of decision-theoretic choice, because it merges two unrelated plans.

Recall the example of transporting a ton of gold and a ton of silver:

> The plan to transport the gold on one trip and the silver on another trip is constructed by merging two unrelated plans for achieving unrelated goals.

6. Maximal Plans

A way out?

maximal plans — plans prescribing what the agent should do for all the rest of its existence.

Proposal:

> A maximal plan is rationally adoptable iff no other maximal plan has a higher expected-utility.

There may be no optimal master plans.

In a domain of real-world complexity, it is computationally impossible for agents with realistic resource constraints to construct maximal plans.

7. Real Rationality and Ideal Rationality

Classical decision theory tries to do two things:

> characterize what is the "objectively best thing to do".
>
> tells us that rationality requires an agent to choose what he rationally believes to be an optimal action out of all the alternatives available.

The prescription to choose an optimal action can seem sensible as long as it is presupposed that the choice of which action to perform is made from a manageably small set of alternatives.

Choices must be made between plans.

Plans are mathematical or logical constructions of unbounded complexity.

In the real world there will be infinitely many competitors for any given plan.

Choosing between plans would then involve evaluating and comparing infinitely many plans — *this is impossible.*

In the real world, for every competitor there is likely to be another with a higher expected-utility, so optimality makes no sense.

Classical decision theory envisages a kind of “ideal rationality” where an agent can survey all possible courses of action and choose an optimal one.

That is a computationally impossible ideal.

Real rationality — the rules governing rational choice in real agents operating in complex environments — must set different standards.

In a general setting there may be no optimal plans, but even if there are, rationality cannot require selecting optimal courses of action.

8. Locally Global Planning

Optimality is not achievable.

The basic insight of classical decision theory — what makes a plan "good" is that it will, with various probabilities, bring about various desirable states, and the cost of doing this will be less than the value of what is achieved.

A good plan is one with a positive expected-utility.

Rationality must require agents to select “good” plans.

Given more time to reason, good plans might be supplanted by better plans.

There will be no point at which an agent has exhausted all possibilities in searching for plans.

The agent must take action.

His decisions about how to act must be directed by the best plans found to date — not by the best possible plans that *could* be found (whatever that might mean).

Planning in rational agents is an incremental process.

The agent constructs and adopts plans, but the search never ends.

New plans may conflict with previously adopted plans, leading either to the rejection of the new plans or the withdrawal of the previously adopted plans.

The agent's plans "evolve" as time passes and more reasoning is performed.

When the time comes to act, the agent directs his actions on the basis of the present set of adopted plans.

How are plans adopted incrementally?

Suppose the agent has no background of adopted plans:

A new plan is constructed.

If this is the only plan the agent has constructed, there is only one other option — do nothing.

In this limiting case, we can evaluate the plan by simply comparing it with doing nothing.

This is the same thing as asking whether its expected-utility is positive.

Suppose the agent has already adopted a number of other plans:

A new plan cannot be evaluated in isolation from the previously adopted plans.

The new plan can affect the value of the old plans.

The decision whether to adopt a new plan must take account of both the effect of previously adopted plans on the new plan, and the effect of the new plan on previously adopted plans.

The agent’s *master-plan* — the result of merging all of the agent’s adopted plans into a single plan.

The expected-utility of the master-plan is our expectation of how much better the world will be if we adopt that as our master-plan.

One master-plan is better than another iff it has a higher expected-utility.

Equivalently, rationality dictates that if an agent is choosing between two master-plans, he should choose the one with the higher expected-utility.

If a master-plan has a positive expected-utility it is better than doing nothing.

This is a reason for adopting it.

But only a defeasible reason — if a different master-plan with a higher expected-utility is found later, then it is a better master-plan and should be adopted in place of the first master-plan.

A general theory of rational choice:

An intention to perform an action is rationally acceptable iff it is part of a rationally adoptable master-plan.

The dynamics of rational choice all concern the selection of master-plans.

A master-plan is defeasibly adoptable iff it has a positive expected-utility, and it remains adoptable iff no master-plan with a higher expected-utility has been found.

Rational deliberation is aimed at finding better and better master-plans.

A master-plan is a *global plan* for achieving all of the agent's goals simultaneously (insofar as that is possible).

The only way resource-bounded agents can construct and improve upon master-plans reflecting the complexity of the real world is by constructing them incrementally.

> When trying to improve our master-plan, rather than throwing it out and starting over from scratch, what we must do is try to improve it piecemeal, leaving the bulk of it intact at any given time.

Normal planning processes produce *local plans* — small plans of limited scope, aiming at isolable goals.

> Local plans represent the building blocks for master-plans.

How can we evaluate a newly constructed local plan, given that we must take account both of its effect on the agent's other plans, and the effect of the agent's other plans on the new plan?

> What we want to know is whether we can improve the master-plan by adding the local plan to it.

Define the *marginal expected-utility* of the local plan *P* to be the difference its addition makes to the master-plan *M*:

MEU(*P*,*M*) = **expected-utility**(*M*+*P*) – **expected-utility**(*M*).

<u>First approximation</u>:

> It is rational for an agent to adopt a local plan iff its marginal expected-utility is positive.

Two problems with the first approximation:

Adding the new plan may only increase the expected-utility of the master-plan if we simultaneously delete conflicting plans.

Plans may have to be added in groups rather than individually.

Where M is a master-plan and C a change, let $M\Delta C$ be the result of making the change to M.

Define the *marginal expected-utility of a change C* to be the difference it makes to the expected-utility of the master-plan:

$$\textbf{MEU}(C,M) = \textbf{expected-utility}(M\Delta C) - \textbf{expected-utility}(M).$$

<u>The principle of locally global planning</u>:

> It is rational for an agent to make a change C to the master-plan M iff the marginal expected-utility of C is positive, i.e., iff $\textbf{expected-utility}(M\Delta C) > \textbf{expected-utility}(M)$.

This is my proposed replacement for classical decision theory.

9. Conclusions

Classical decision theory must be rejected for two reasons.

It would only be reasonable if action omnipotence held.

Action omnipotence fails in two ways — we cannot always perform actions when we try, and sometimes we cannot even try.

This first difficulty can be overcome by evaluating actions in terms of their expected-utilities rather than their expected-values.

Conclusions

The second difficulty is that in general actions cannot be chosen in isolation. They must be chosen as parts of plans.

We cannot save classical decision theory by adopting a plan-based decision theory according to which it is rational to adopt a plan iff it is an optimal plan from a set of alternatives:

> There is no apparent way to define “alternative” for local plans so that rational choice consists of choosing an optimal plan from a bounded list of alternatives.
>
> Even if there were optimal plans, real agents could not be required to find them, because that would require surveying and comparing infinitely many plans.

Conclusions

The only appropriate objects for evaluation in terms of expected-utilities are master-plans.

The objective of rational decision making cannot be the construction of optimal master-plans:

> There is no reason to think that there will exist optimal master-plans.
>
> Even if there were optimal master-plans, the search for them would be computationally intractable. Real agents would not be able to find them.

The objective of rational deliberation must be to find a good master-plan, and to be on the continual lookout for ways of making the master-plan better.

Conclusions

Real agents will not be able to construct master-plans as the result of single planning exercises. The plans are too complex for that.

The master-plan must instead be constructed incrementally, by engaging in local planning and then merging the local plans into the master-plan.

THE END