Title: Decisions, Decisions, Decisions: Thoughts about actions in an Everett World

Date: Sep 23, 2007 11:40 AM

URL: http://pirsa.org/07090075

Abstract: The most common objection to the Everett view of QM is that it \'cannot make sense of probability\'. The \'Oxford project\' of writers such as Deutsch, Wallace, Saunders and Greaves seeks to meet this objection by showing that the Everett view allows some suitable analogue of decision under uncertainty, and that probability (or some suitable analogue of probability) can be understood on that basis. As a pragmatist, I\'m very sympathetic to the idea that probability in general needs to be understood in terms of its links with decision; but I\'m sceptical about whether the Everett picture provides a suitable analogue of decision under uncertainty. In this talk I\'ll try to justify my scepticism.

Decisions, Decisions, Decisions

Huw Price

Centre for Time
University of Sydney

23 September 2007

Decisions, Decisions, Decisions

Thoughts About Actions in the Everett World

Huw Price

Centre for Time
University of Sydney

23 September 2007

The right road for quantum theory?

The right road for quantum theory?

'Classical probability is mysterious, too'

'Classical probability is mysterious, too'

- Papineau on the 'Decision-Theoretic Link':

'Classical probability is mysterious, too'

- Papineau on the 'Decision-Theoretic Link':

> '[W]hy are rational agents well advised to choose actions that make their desired results objectively probable? ... [T]here is no good answer to this question ... . [M]any philosophers in this area now simply take it to be a primitive fact that you ought to weight future possibilities according to known objective probabilities in making rational decisions. ... It is not just that philosophers can't agree on the right justification; many have concluded that there simply isn't one.' (Papineau 1996, 238)

The pragmatist solution

The pragmatist solution

- Emphasise the practical . . .

The pragmatist solution

- Emphasise the practical ...

> 'The [Principal Principle][1] is intended by Lewis to provide a *functional definition* of chance ... i.e. it *defines* chance by what it does ... "chance" is that property of the physical world – *whatever that property is* – that fits the "chance" role of the [Principal Principle].' (Wallace)

[1] "If I know that the chance of an outcome is p then I am rationally required to have credence p in that outcome."

The pragmatist solution

- Emphasise the practical ...

> 'The [Principal Principle][1] is intended by Lewis to provide a *functional definition* of chance ... i.e. it *defines* chance by what it does ... "chance" is that property of the physical world – *whatever that property is* – that fits the "chance" role of the [Principal Principle].' (Wallace)

- And don't worry about the metaphysics.

[1] "If I know that the chance of an outcome is p then I am rationally required to have credence p in that outcome."

- Consider people who mark in blue on their maps the places where they find stuff they can drink and wash in.

- There might be other mysteries in the vicinity – e.g., about why blue lines are correlated with contour lines in a distinctive way,

- There might be other mysteries in the vicinity – e.g., about why blue lines are correlated with contour lines in a distinctive way,

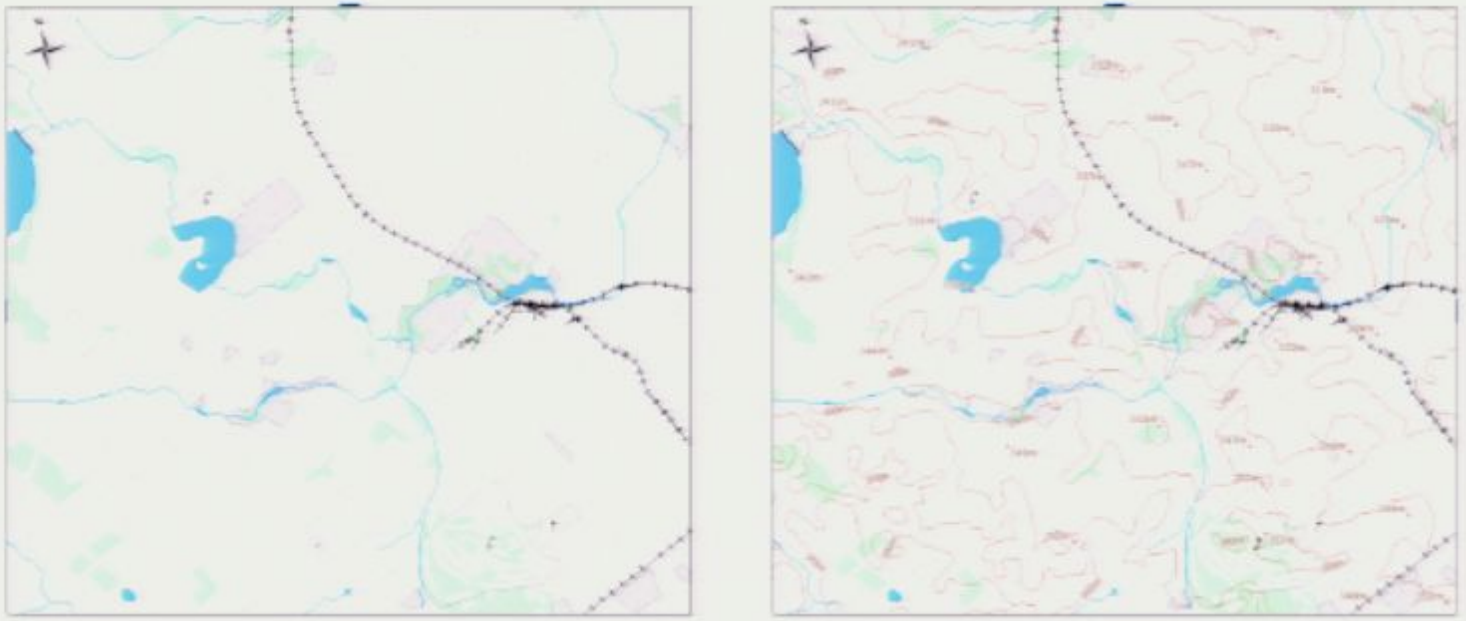

- There might be other mysteries in the vicinity – e.g., about why blue lines are correlated with contour lines in a distinctive way,

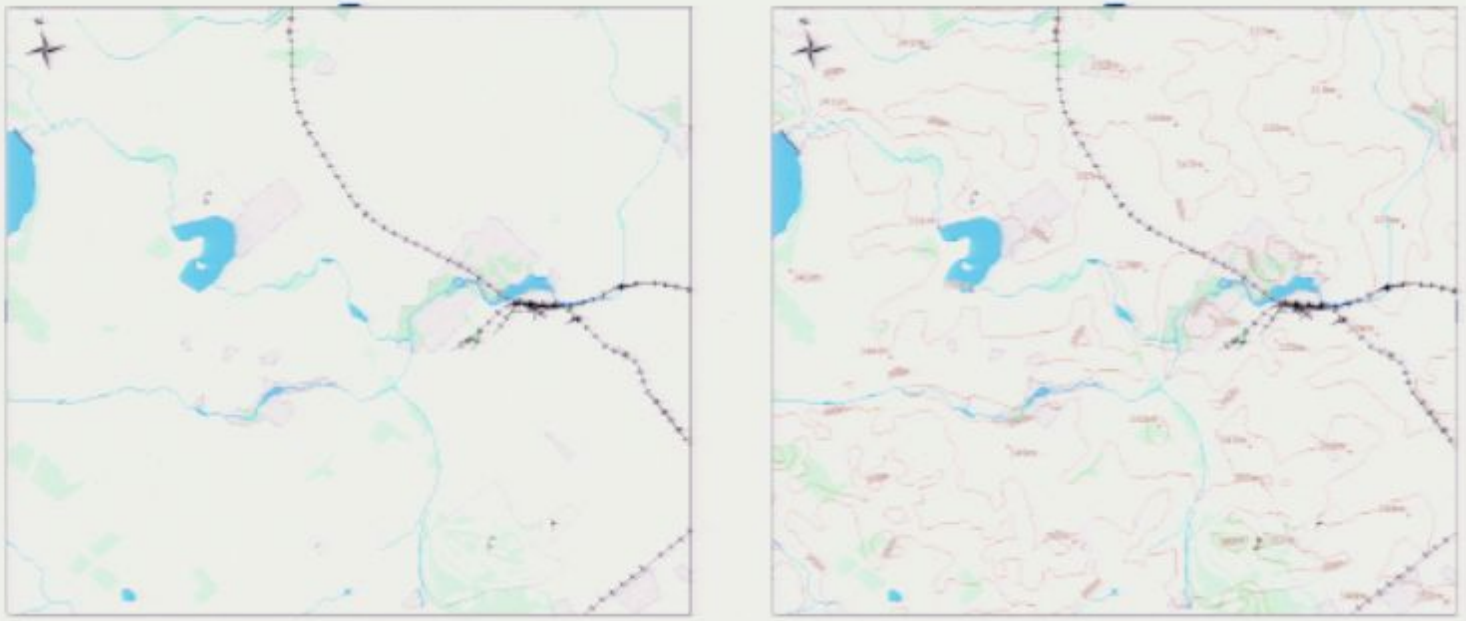

- There might be other mysteries in the vicinity – e.g., about why blue lines are correlated with contour lines in a distinctive way, or whether there is any unified **physical** account of the stuff we drink and wash in –

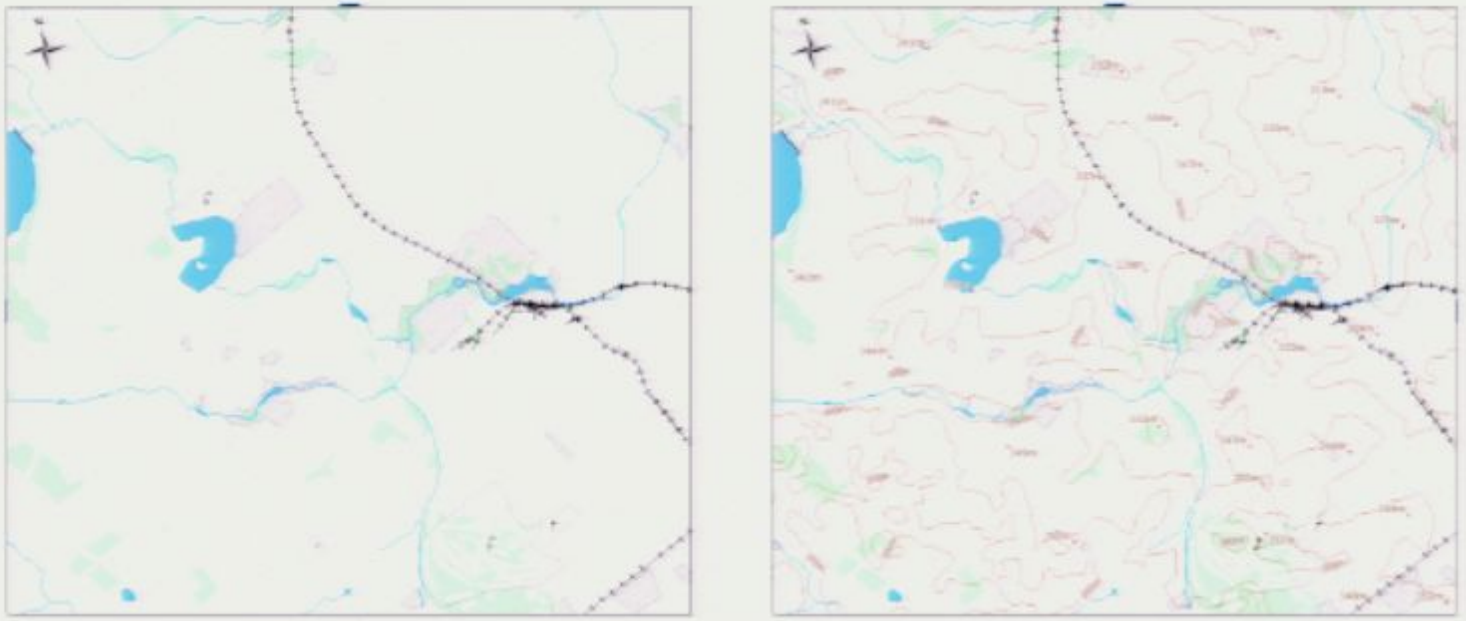

- There might be other mysteries in the vicinity – e.g., about why blue lines are correlated with contour lines in a distinctive way, or whether there is any unified **physical** account of the stuff we drink and wash in – but these are not the practical puzzle about why places marked with blue lines are good for washing and drinking.

Applying this to probability:

Applying this to probability:

- If we **start** with idea that a map of probability is a guide to decision under uncertainty, it isn't a mystery why the map can be used for exactly that purpose.

Applying this to probability:

- If we **start** with idea that a map of probability is a guide to decision under uncertainty, it isn't a mystery why the map can be used for exactly that purpose.
- There might be other mysteries in the vicinity – e.g., about whether, or why, decision-theoretic 'blue lines' are correlated with something else on our maps – but there's no mystery about why probability properly applies [illegible]

Applying this to probability:

- If we **start** with idea that a map of probability is a guide to decision under uncertainty, it isn't a mystery why the map can be used for exactly that purpose.
- There might be other mysteries in the vicinity – e.g., about whether, or why, decision-theoretic 'blue lines' are correlated with something else on our maps – but these are not the practical puzzle about why probability properly guides action.

Applying this to probability:

- If we **start** with idea that a map of probability is a guide to decision under uncertainty, it isn't a mystery why the map can be used for exactly that purpose.
- There might be other mysteries in the vicinity – e.g., about whether, or why, decision-theoretic 'blue lines' are correlated with something else on our maps – but these are not the practical puzzle about why probability properly guides action.

Applying this to probability:

- If we **start** with idea that a map of probability is a guide to decision under uncertainty, it isn't a mystery why the map can be used for exactly that purpose.
- There might be other mysteries in the vicinity – e.g., about whether, or why, decision-theoretic 'blue lines' are correlated with something else on our maps – but these are not the practical puzzle about why probability properly guides action.
- (Would it matter is there wasn't any **unified** story about 'something else' that correlated with decision theoretic probability? If so, why?)

Lessons

Lessons

- We shouldn't lose sight of the fact that probability **begins** with decision under uncertainty.

Lessons

- We shouldn't lose sight of the fact that probability **begins** with decision under uncertainty.
- As long as we keep that in mind, there isn't any mystery about Papineau's Decision-Theoretic Link.

Lessons

- We shouldn't lose sight of the fact that probability **begins** with decision under uncertainty.
- As long as we keep that in mind, there isn't any mystery about Papineau's Decision-Theoretic Link.
- The crucial issue for the coherence of Everettian probability is whether decision under uncertainty – or some suitable analogue – does make sense in the many worlds case.

My project

My project

1. Begin by applauding the Deutsch-Wallace ("Oxford") approach to Everettian probability for beginning with the issue of rational decision

My project

1. Begin by applauding the Deutsch-Wallace ("Oxford") approach to Everettian probability for beginning with the issue of rational decision (asking it to allow its one-world rivals to do the same).
2. Raise some difficulties for the view that rational action in an Everett world should conform to an MEU model *at all* – I've shown that its [illegible]

My project

1. Begin by applauding the Deutsch-Wallace ("Oxford") approach to Everettian probability for beginning with the issue of rational decision (asking it to allow its one-world rivals to do the same).
2. Raise some difficulties for the view that rational action in an Everett world should conform to an MEU model *at all*

My project

1. Begin by applauding the Deutsch-Wallace ("Oxford") approach to Everettian probability for beginning with the issue of rational decision (asking it to allow its one-world rivals to do the same).
2. Raise some difficulties for the view that rational action in an Everett world should conform to an MEU model *at all* – i.e., should aim to maximise a weighted sum of in-branch utilities — whether via the Born rule or otherwise.

My project

1. Begin by applauding the Deutsch-Wallace ("Oxford") approach to Everettian probability for beginning with the issue of rational decision (asking it to allow its one-world rivals to do the same).
2. Raise some difficulties for the view that rational action in an Everett world should conform to an MEU model *at all* – i.e., should aim to maximise a weighted sum of in-branch utilities — whether via the Born rule or otherwise.

Thinking about preferences

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference – unfamiliar, but surely not irrational

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference – unfamiliar, but surely not irrational – viz., a 'global' preference about the future QM state.

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference – unfamiliar, but surely not irrational – viz., a 'global' preference about the future QM state.
- For an agent whose only preferences are of this kind, the choice between QM games isn't a choice between two weighted sets of alternative '**possible**' outcomes

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference – unfamiliar, but surely not irrational – viz., a 'global' preference about the future QM state.
- For an agent whose only preferences are of this kind, the choice between QM games isn't a choice between two weighted sets of alternative '**possible**' outcomes – it is simply a choice between two **certain** outcomes (i.e., the two different states which result from the two games in question).

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference – unfamiliar, but surely not irrational – viz., a 'global' preference about the future QM state.
- For an agent whose only preferences are of this kind, the choice between QM games isn't a choice between two weighted sets of alternative '**possible**' outcomes – it is simply a choice between two **certain** outcomes (i.e., the two different states which result from the two games in question).
- **Conclusion:** It can't be true *in general* that rational choice of QM games conforms to a MEU model.

Thinking about preferences

- A preference is a desire *that* some state of affairs should obtain – *that* some proposition should turn out to be true.
- The Everett picture provides the ontological basis for a new kind of preference – unfamiliar, but surely not irrational – viz., a 'global' preference about the future QM state.
- For an agent whose only preferences are of this kind, the choice between QM games isn't a choice between two weighted sets of alternative '**possible**' outcomes – it is simply a choice between two **certain** outcomes (i.e., the two different states which result from the two games in question).
- **Conclusion:** It can't be true *in general* that rational choice of QM games conforms to a MEU model. At most, the MEU model applies (non-trivially) only to the component of choice which is not guided by such global preferences.

The global challenge

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, where this preference is not explained in terms of co-branch utilities.

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, *where this preference is not explained in terms of in-branch utilities.*

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, *where this preference is not explained in terms of in-branch utilities.*
- On the contrary, the in-branch utilities would reflect the global preferences –

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, *where this preference is not explained in terms of in-branch utilities.*
- On the contrary, the in-branch utilities would reflect the global preferences – her descendants would be happy or unhappy, according to whether they believed her preferences had been met at the global level.

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, *where this preference is not explained in terms of in-branch utilities.*
- On the contrary, the in-branch utilities would reflect the global preferences – her descendants would be happy or unhappy, according to whether they believed her preferences had been met at the global level.
- In this case, there's no space of differentially-valued outcomes, to which analogues of probabilistic weights can attach (non-trivially).

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, *where this preference is not explained in terms of in-branch utilities.*
- On the contrary, the in-branch utilities would reflect the global preferences – her descendants would be happy or unhappy, according to whether they believed her preferences had been met at the global level.
- In this case, there's no space of differentially-valued outcomes, to which analogues of probabilistic weights can attach (non-trivially).
- So any argument for the rationality of a Born-weighted MEU model needs to come with a rider:

The global challenge

- In other words, an Everettian agent might attach values to the state itself – i.e., might prefer one game to another simply because she prefers one future evolution of the state function to another, *where this preference is not explained in terms of in-branch utilities.*
- On the contrary, the in-branch utilities would reflect the global preferences – her descendants would be happy or unhappy, according to whether they believed her preferences had been met at the global level.
- In this case, there's no space of differentially-valued outcomes, to which analogues of probabilistic weights can attach (non-trivially).
- So any argument for the rationality of a Born-weighted MEU model needs to come with a rider: e.g., that it only applies to actions *in so far as they are determined by non-global preferences.*

Three reasons why this might matter

Three reasons why this might matter

1. Because the DW argument appeals to global preferences at crucial points, in claiming that rational agents should be indifferent between games which give rise to the same QM state

Three reasons why this might matter

1. Because the DW argument appeals to global preferences at crucial points, in claiming that rational agents should be indifferent between games which give rise to the same QM state – if this isn't a *direct* appeal to global preferences (which the rider would exclude), it looks suspiciously like question begging.

Three reasons why this might matter

1. Because the DW argument appeals to global preferences at crucial points, in claiming that rational agents should be indifferent between games which give rise to the same QM state – if this isn't a *direct* appeal to global preferences (which the rider would exclude), it looks suspiciously like question begging.
2. Because reflective Everettian agents are unlikely to be indifferent *in practice* to the welfare of their co-descendants (and others!) in other branches, so that the rider is unrealistic *de facto*.

Three reasons why this might matter

1. Because the DW argument appeals to global preferences at crucial points, in claiming that rational agents should be indifferent between games which give rise to the same QM state – if this isn't a *direct* appeal to global preferences (which the rider would exclude), it looks suspiciously like question begging.
2. Because reflective Everettian agents are unlikely to be indifferent *in practice* to the welfare of their co-descendants (and others!) in other branches, so that the rider is unrealistic *de facto*.
3. Because reflective Everettian agents *should* care about their co-descendants in a way which may be in tension with the recommendations of the MEU model, when the rider is assumed.

Three reasons why this might matter

1. Because the DW argument appeals to global preferences at crucial points, in claiming that rational agents should be indifferent between games which give rise to the same QM state – if this isn't a *direct* appeal to global preferences (which the rider would exclude), it looks suspiciously like question begging.
2. Because reflective Everettian agents are unlikely to be indifferent *in practice* to the welfare of their co-descendants (and others!) in other branches, so that the rider is unrealistic *de facto*.
3. Because reflective Everettian agents *should* care about their co-descendants in a way which may be in tension with the recommendations of the MEU model, when the rider is assumed.

Plan: Say no more about (1), but something about (2) and (3).

Caring about one's alter egos

Caring about one's alter egos

- Should in-branch 'net satisfaction' be expected to coincide in the Everett and one-world cases?
- Or should the value Everettian agents ascribe to branches *also* depend on how they believe their co-descendants are faring in other branches?

Caring about one's alter egos

- Should in-branch 'net satisfaction' be expected to coincide in the Everett and one-world cases?
- Or should the value Everettian agents ascribe to branches *also* depend on how they believe their co-descendants are faring in other branches?

A fundamental disanalogy?

A fundamental disanalogy?

- If I win a risky bet, I don't care about the misfortune of my unlucky possible twin (who loses)

A fundamental disanalogy?

- If I win a risky bet, I don't care about the misfortune of my unlucky possible twin (who loses) ... because **he doesn't exist.**

A fundamental disanalogy?

- If I win a risky bet, I don't care about the misfortune of my unlucky possible twin (who loses) ... because **he doesn't exist.**
- But in the Everett case my good fortune comes at a cost to someone very close to me.

A fundamental disanalogy?

- If I win a risky bet, I don't care about the misfortune of my unlucky possible twin (who loses) ... because **he doesn't exist.**
- But in the Everett case my good fortune comes at a cost to someone very close to me. Mightn't this make a difference to how much I enjoy it?

An example:

An example:

An example:

- Suppose a shark swallows my right leg (here at Bondi Beach).

An example:

- Suppose a shark swallows my right leg (here at Bondi Beach). I'm offered a procedure which clones (and reflects) my surviving leg, to make a replacement.
- There's a small risk that I might lose both legs, but if this risk is sufficiently small, it will be a risk worth taking.

An example:

- Suppose a shark swallows my right leg (here at Bondi Beach). I'm offered a procedure which clones (and reflects) my surviving leg, to make a replacement.
- There's a small risk that I might lose both legs, but if this risk is sufficiently small, it will be a risk worth taking.

An example:

- Now suppose I learn that the procedure actually clones all of me, and normally leaves my duplicate with no legs – that's where the spare leg comes from.

An example:

- Now suppose I learn that the procedure actually clones all of me, and normally leaves my duplicate with no legs – that's where the spare leg comes from. (He only gets the legs when the procedure 'fails'.)
- Doesn't this make a difference to my subjective utility of getting my leg back?

An example:

- Now suppose I learn that the procedure actually clones all of me, and normally leaves my duplicate with no legs – that's where the spare leg comes from. (He only gets the legs when the procedure 'fails'.)
- Doesn't this make a difference to my subjective utility of getting my leg back?

A reply?

A reply?

- Why not meet this challenge by emphasising the agent's point of view before the bet?

A reply?

- Why not meet this challenge by emphasising the agent's point of view before the bet? Isn't is my earlier self's responsibility to take this into account (thus relieving me of any obligation to feel guilty about my good fortune)?

A reply?

- Why not meet this challenge by emphasising the agent's point of view before the bet? Isn't is my earlier self's responsibility to take this into account (thus relieving me of any obligation to feel guilty about my good fortune)?

Answer

- This just pushes the problem somewhere else

A reply?

- Why not meet this challenge by emphasising the agent's point of view before the bet? Isn't is my earlier self's responsibility to take this into account (thus relieving me of any obligation to feel guilty about my good fortune)?

Answer

- This just pushes the problem somewhere else – shouldn't my earlier self's policy reflect the same kind of disquiet?
- And doesn't this recommend a policy according to which the value of a bet isn't always the weighted sum of one-world utilities (even in the absence of explicit global preferences)?

A reply?

- Why not meet this challenge by emphasising the agent's point of view before the bet? Isn't is my earlier self's responsibility to take this into account (thus relieving me of any obligation to feel guilty about my good fortune)?

Answer

- This just pushes the problem somewhere else – shouldn't my earlier self's policy reflect the same kind of disquiet?
- And doesn't this recommend a policy according to which the value of a bet isn't always the weighted sum of one-world utilities (even in the absence of explicit global preferences)?

The basic insight

The basic insight

- Decision in the Everett world concerns the welfare of a **group** of future individuals – all one's future descendants.

The basic insight

- Decision in the Everett world concerns the welfare of a **group** of future individuals – all one's future descendants.
- This is highlighted by Greaves' 'caring measure' approach, but true for the 'subjective uncertainty' approach as well.

The basic insight

- Decision in the Everett world concerns the welfare of a **group** of future individuals – all one's future descendants.
- This is highlighted by Greaves' 'caring measure' approach, but true for the 'subjective uncertainty' approach as well.[2]

[2] My disquiet in the Bondi example isn't offset even by certainty that it's me who gets the legs.

The basic insight

- Decision in the Everett world concerns the welfare of a **group** of future individuals – all one's future descendants.
- This is highlighted by Greaves' 'caring measure' approach, but true for the 'subjective uncertainty' approach as well.[2]
- **Challenge:** Isn't rational decision in such a context fundamentally different from any weighted sum model?

[2] My disquiet in the Bondi example isn't offset even by certainty that it's me who gets the legs.

Distributive justice

Distributive justice

- A weighted sum allocation of goods to a group always permits a large cost to one individual to be offset by small gains to others.
- A principle of maximising such a sum may thus conflict with plausible principles of justice:

Distributive justice

- A weighted sum allocation of goods to a group always permits a large cost to one individual to be offset by small gains to others.
- A principle of maximising such a sum may thus conflict with plausible principles of justice:

> "Pleasure for one should not (knowingly) be gained at the cost of pain for another."

Distributive justice

- A weighted sum allocation of goods to a group always permits a large cost to one individual to be offset by small gains to others.
- A principle of maximising such a sum may thus conflict with plausible principles of justice:

> *"Pleasure for one should not (knowingly) be gained at the cost of pain for another."*

Some attractive decision principles for Everettian agents

Some attractive decision principles for Everettian agents

1. Try to establish a baseline, below which you don't knowingly allow your descendants to fall (at least not simply for the sake of modest advantage to others).

Some attractive decision principles for Everettian agents

1. Try to establish a baseline, below which you don't knowingly allow your descendants to fall (at least not simply for the sake of modest advantage to others).
2. Encourage highflyers ... but not if it means forcing another descendant to fall below the baseline.

Some attractive decision principles for Everettian agents

1. Try to establish a baseline, below which you don't knowingly allow your descendants to fall (at least not simply for the sake of modest advantage to others).
2. Encourage highflyers . . . but not if it means forcing another descendant to fall below the baseline. ('Keep the peaks narrow to keep the baseline high.')
3. Trade off quantity for quality?

Some attractive decision principles for Everettian agents

1. Try to establish a baseline, below which you don't knowingly allow your descendants to fall (at least not simply for the sake of modest advantage to others).
2. Encourage highflyers ... but not if it means forcing another descendant to fall below the baseline. ('Keep the peaks narrow to keep the baseline high.')
3. Trade off quantity for quality?

Some attractive decision principles for Everettian agents

1. Try to establish a baseline, below which you don't knowingly allow your descendants to fall (at least not simply for the sake of modest advantage to others).
2. Encourage highflyers ... but not if it means forcing another descendant to fall below the baseline. ('Keep the peaks narrow to keep the baseline high.')
3. Trade off quantity for quality? (Quantum Russian roulette, ordinary lotteries.)

The challenge

- It is (at best) an **open question** whether an Everettian agent should assess the value of a quantum bet as a weighted sum of the 'one-world' utilities of the in-branch outcomes

The challenge

- It is (at best) an **open question** whether an Everettian agent should assess the value of a quantum bet as a weighted sum of the 'one-world' utilities of the in-branch outcomes ... *even* if we grant that those utilities (rather than global preferences about the state) are to be the primary drivers of the decisions in question.

The challenge

- It is (at best) an **open question** whether an Everettian agent should assess the value of a quantum bet as a weighted sum of the 'one-world' utilities of the in-branch outcomes ... *even* if we grant that those utilities (rather than global preferences about the state) are to be the primary drivers of the decisions in question.
- The Oxford model of rational action seems to need to be qualified in at least two ways:
  1. Restrictions on 'out of branch' preferences.
  2. Restrictions to a sufficiently narrow range of utilities to avoid issues of justice, uniformness, etc.

The challenge

- It is (at best) an **open question** whether an Everettian agent should assess the value of a quantum bet as a weighted sum of the 'one-world' utilities of the in-branch outcomes ... *even* if we grant that those utilities (rather than global preferences about the state) are to be the primary drivers of the decisions in question.
- The Oxford model of rational action seems to need to be qualified in at least two ways:
  1. Restrictions on 'out of branch' preferences.
  2. Restrictions to a sufficiently narrow range of utilities to avoid issues of justice, unfairness, etc.

Objection

Objection

- These issues of 'justice' are no different from the one-world case, in which betting produces losers as well as winners.

Objection

- These issues of 'justice' are no different from the one-world case, in which betting produces losers as well as winners.

Answer

- In the one-world case the losses and gains are born by the same individual. In the Everett world there's a new sense in which this isn't so – that's why it is like the distributive justice case.

Objection

- These issues of 'justice' are no different from the one-world case, in which betting produces losers as well as winners.

Answer

- In the one-world case the losses and gains are born by the same individual. In the Everett world there's a new sense in which this isn't so – that's why it is like the distributive justice case.
- And repetition doesn't change this, because Everett implies that for any finite number of repetitions, there are individuals who always lose.

Objection

Objection

- You've ignored the weights

Objection

- You've ignored the weights – good betting practice ensures that the losers are on low-weight branches.

Objection

- You've ignored the weights – good betting practice ensures that the losers are on low-weight branches.

Answer

- This brings me to another problem ...

A fundamental issue

- What entitles us to give less consideration to low-weight futures and descendants?

A fundamental issue

- What entitles us to give less consideration to low-weight futures and descendants?

The one-world solution

The one-world solution

- In the limit, no matter how fine my 'epistemic grain', the low-weight alternatives simply disappear off my epistemic map.

The one-world solution

- In the limit, no matter how fine my 'epistemic grain', the low-weight alternatives simply disappear off my epistemic map.
- That is, any graining has a resolution limit, beneath which low probability equates to zero, for the purposes at hand.

The one-world solution

- In the limit, no matter how fine my 'epistemic grain', the low-weight alternatives simply disappear off my epistemic map.
- That is, any graining has a resolution limit, beneath which low probability equates to zero, for the purposes at hand.
- So for the purposes of my view of the possibilities, it is just as if I believe that the low-weight alternatives *don't exist*

The one-world solution

- In the limit, no matter how fine my 'epistemic grain', the low-weight alternatives simply disappear off my epistemic map.
- That is, any graining has a resolution limit, beneath which low probability equates to zero, for the purposes at hand.
- So for the purposes of my view of the possibilities, it is just as if I believe that the low-weight alternatives *don't exist* – and it isn't a puzzle why I needn't care about the welfare of non-existents.

The one-world solution

- In the limit, no matter how fine my 'epistemic grain', the low-weight alternatives simply disappear off my epistemic map.
- That is, any graining has a resolution limit, beneath which low probability equates to zero, for the purposes at hand.
- So for the purposes of my view of the possibilities, it is just as if I believe that the low-weight alternatives *don't exist* – and it isn't a puzzle why I needn't care about the welfare of non-existents.
- (Call this the credence-existence link.)

This doesn't work in the Everett case:

The one-world solution

- In the limit, no matter how fine my 'epistemic grain', the low-weight alternatives simply disappear off my epistemic map.
- That is, any graining has a resolution limit, beneath which low probability equates to zero, for the purposes at hand.
- So for the purposes of my view of the possibilities, it is just as if I believe that the low-weight alternatives *don't exist* – and it isn't a puzzle why I needn't care about the welfare of non-existents.
- (Call this the credence-existence link.)

This doesn't work in the Everett case: if I believe that well-weighted pleasure comes at the expense of low-weighted pain, I believe that it's *real* pain, despite its weight.

Appeal to frequencies?

Appeal to frequencies?

- Everettians might have an answer to this challenge, if weight could be construed in terms of frequency,

Appeal to frequencies?

- Everettians might have an answer to this challenge, if weight could be construed in terms of frequency, and we had a suitable principle of indifference, allowing us to treat everyone equally.

Appeal to frequencies?

- Everettians might have an answer to this challenge, if weight could be construed in terms of frequency, and we had a suitable principle of indifference, allowing us to treat everyone equally. But:
  - It is not clear that it would be a good enough answer, because the moral principle that one shouldn't harm X to benefit Y is arguably insensitive to the number of Ys.

Appeal to frequencies?

- Everettians might have an answer to this challenge, if weight could be construed in terms of frequency, and we had a suitable principle of indifference, allowing us to treat everyone equally. But:
  - It is not clear that it would be a good enough answer, because the moral principle that one shouldn't harm X to benefit Y is arguably insensitive to the number of Ys.
  - As the Oxford school have emphasised, it is doubtful both whether we have well-defined frequencies in the Everett framework, and whether *mere* frequency could be a relevant consideration – branching seems far too 'contingent' for that.

Isn't sizism unavoidable?

Isn't sizism unavoidable?

- **Claim:** There are too many low-weighted branches

Isn't sizism unavoidable?

- **Claim:** There are too many low-weighted branches – we can't consider them all, so a 'pragmatic limit' requires us to make a cut-off higher up.

Reply I

Reply I

- The fact that we can't (practically) consider all branches is no *justification* for ignoring a salient branch of which we are already well aware (such as the branch containing my legless twin).

Reply I

- The fact that we can't (practically) consider all branches is no *justification* for ignoring a salient branch of which we are already well aware (such as the branch containing my legless twin).
- What we need is some explanation why the low-weighted branches (in particular) 'disappear' for the purposes of deliberation

Reply I

- The fact that we can't (practically) consider all branches is no *justification* for ignoring a salient branch of which we are already well aware (such as the branch containing my legless twin).
- What we need is some explanation why the low-weighted branches (in particular) 'disappear' for the purposes of deliberation – that's what the credence-existence link gives us in the one-world case.

Reply II

- And besides, the problem also occurs 'above the grain' –

Reply II

- And besides, the problem also occurs 'above the grain' – what entitles us to privilege heavily-weighted branches over moderately-weighted branches?

Reply II

- And besides, the problem also occurs 'above the grain' – what entitles us to privilege heavily-weighted branches over moderately-weighted branches?
- (Again, the answer 'Because there are more of them' seems off the table.)

The basic problem

- It hasn't been shown that a weighted sum model of any kind – whether with Born weights or not – is the appropriate way to formalise rational choice between QM games, in an Everettian world.

The basic problem

- It hasn't been shown that a weighted sum model of any kind – whether with Born weights or not – is the appropriate way to formalise rational choice between QM games, in an Everettian world.

The sub-problems

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it).

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it). This is a reason to think that if the Born weights were (quasi-)probabilities, the expected value of an Everettian game would in general **not** equal that of the corresponding one-world game.

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it). This is a reason to think that if the Born weights were (quasi-)probabilities, the expected value of an Everettian game would in general **not** equal that of the corresponding one-world game.
3. In general, rational decision on behalf of a group may be expected **not** to go by a weighted sum of individual utilities, when the welfare of some comes at severe cost to others.

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it). This is a reason to think that if the Born weights were (quasi-)probabilities, the expected value of an Everettian game would in general **not** equal that of the corresponding one-world game.
3. In general, rational decision on behalf of a group may be expected **not** to go by a weighted sum of individual utilities, when the welfare of some comes at severe cost to others.
4. Without the credence-existence link, the Everettian seems to lack any rationale for "sizism"

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it). This is a reason to think that if the Born weights were (quasi-)probabilities, the expected value of an Everettian game would in general **not** equal that of the corresponding one-world game.
3. In general, rational decision on behalf of a group may be expected **not** to go by a weighted sum of individual utilities, when the welfare of some comes at severe cost to others.
4. Without the credence-existence link, the Everettian seems to lack any rationale for "sizism" – i.e., for giving preference to higher-weighted individuals.

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it). This is a reason to think that if the Born weights were (quasi-)probabilities, the expected value of an Everettian game would in general **not** equal that of the corresponding one-world game.
3. In general, rational decision on behalf of a group may be expected **not** to go by a weighted sum of individual utilities, when the welfare of some comes at severe cost to others.
4. Without the credence-existence link, the Everettian seems to lack any rationale for "sizism" – i.e., for giving preference to higher-weighted individuals. (Without such a justification, a proof that there is no rational MEU alternative to a Born-weights-based model might be seem as a reason for abandoning the MEU model altogether.)

A new argument for fatalism?

A new argument for fatalism?

- This whole issue takes for granted that free choice makes sense in the Everett world, as much as it does in the one-world case.

A new argument for fatalism?

- This whole issue takes for granted that free choice makes sense in the Everett world, as much as it does in the one-world case. But there are reasons to doubt that, too.

A new argument for fatalism?

- This whole issue takes for granted that free choice makes sense in the Everett world, as much as it does in the one-world case. But there are reasons to doubt that, too.
- Whenever I form an intention to ϕ, there's a small 'probability' that my intention will be thwarted.

A new argument for fatalism?

- This whole issue takes for granted that free choice makes sense in the Everett world, as much as it does in the one-world case. But there are reasons to doubt that, too.
- Whenever I form an intention to ϕ, there's a small 'probability' that my intention will be thwarted.
- Hence in an Everett world, I know that my 'decision' to ϕ produces some descendants who ϕ and some descendants who don't ϕ.

A new argument for fatalism?

- This whole issue takes for granted that free choice makes sense in the Everett world, as much as it does in the one-world case. But there are reasons to doubt that, too.
- Whenever I form an intention to ϕ, there's a small 'probability' that my intention will be thwarted.
- Hence in an Everett world, I know that my 'decision' to ϕ produces some descendants who ϕ and some descendants who don't ϕ.
- What sense, then – at least *in advance* of a justification for treating weights as probabilities – can we make of the notion of an effective choice between games?

A new argument for fatalism?

- This whole issue takes for granted that free choice makes sense in the Everett world, as much as it does in the one-world case. But there are reasons to doubt that, too.
- Whenever I form an intention to ϕ, there's a small 'probability' that my intention will be thwarted.
- Hence in an Everett world, I know that my 'decision' to ϕ produces some descendants who ϕ and some descendants who don't ϕ.
- What sense, then – at least *in advance* of a justification for treating weights as probabilities – can we make of the notion of an effective choice between games? (The problem isn't that my choice is *determined*, but that I can't make an effective choice *between alternatives*.)

Confronting branching

The global viewpoint

The End

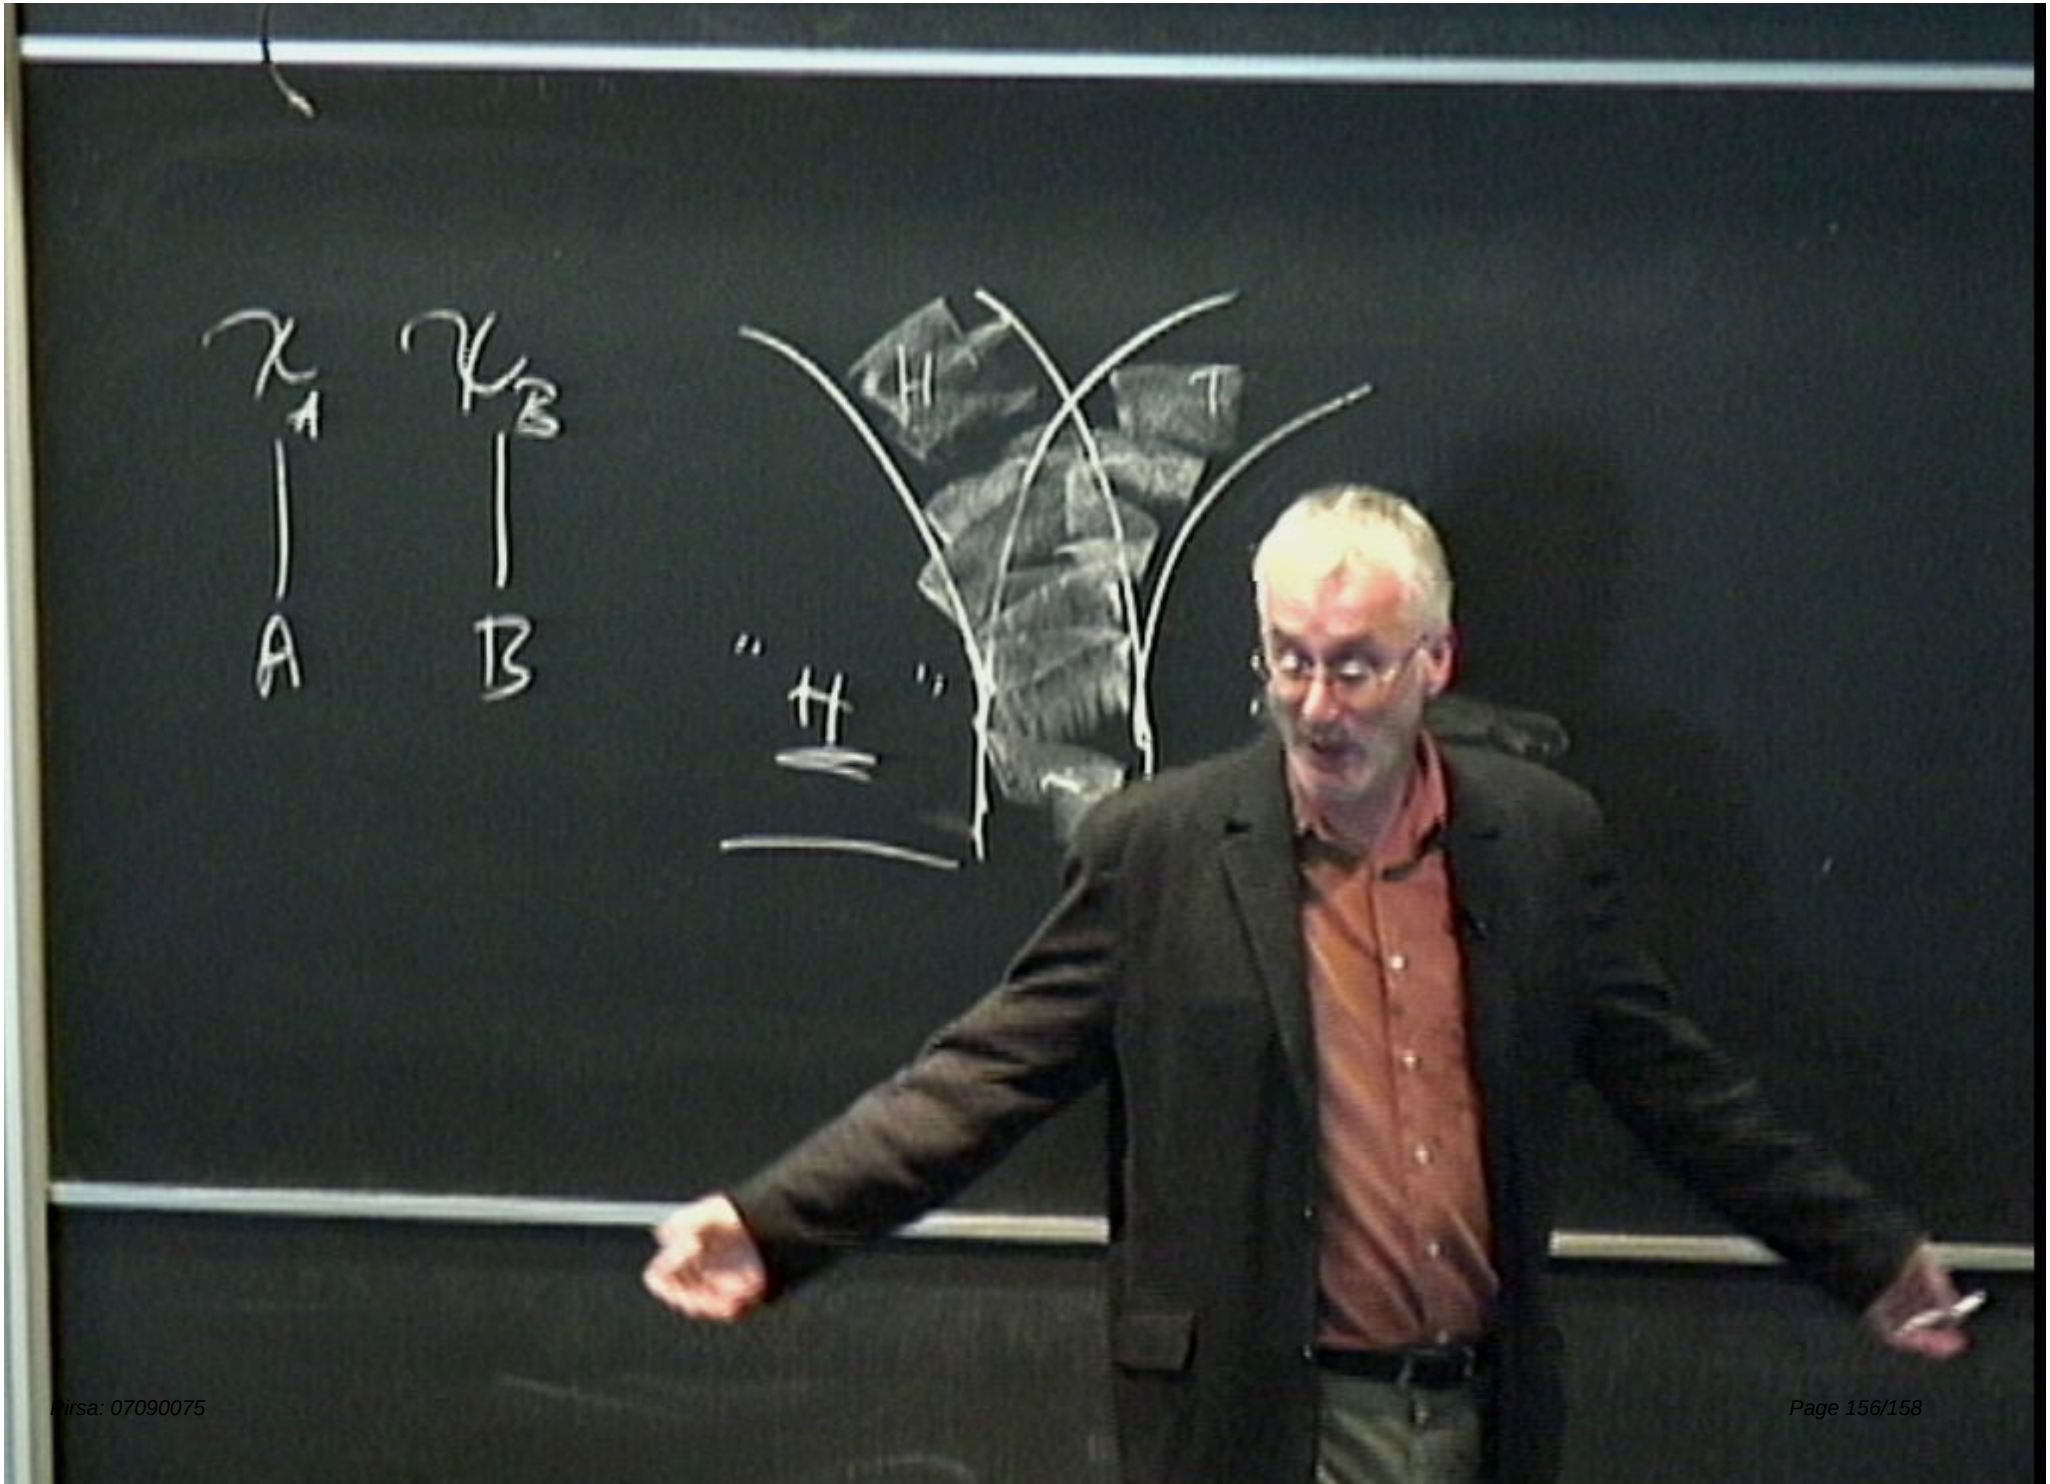
A
B
H
T
"H"

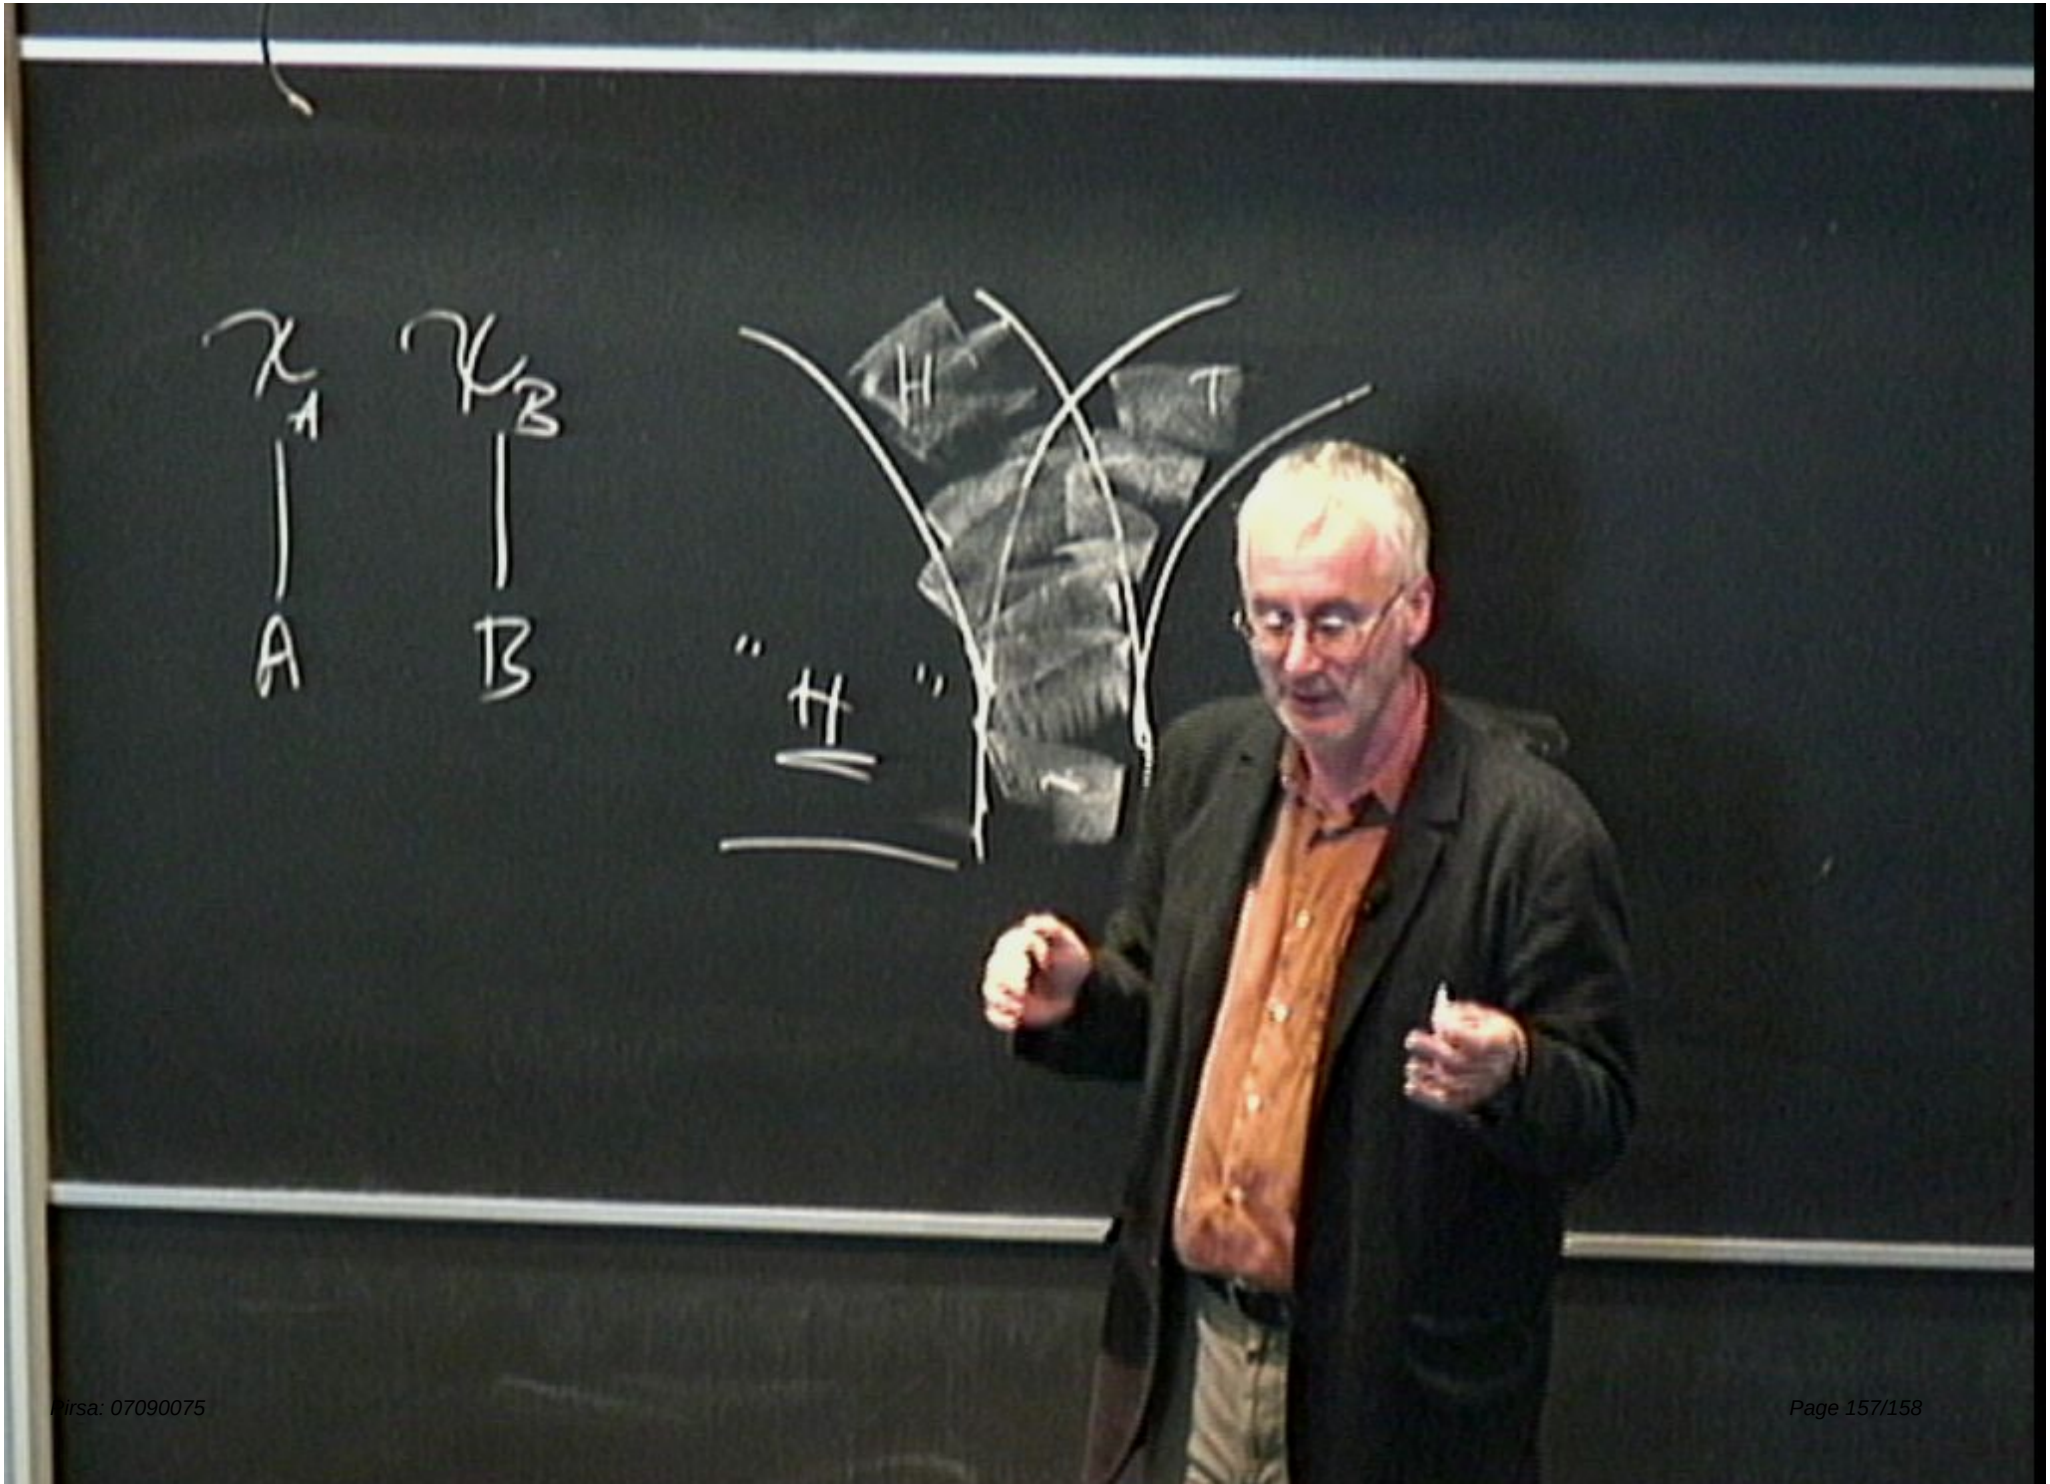
A
B
H
T
"H"

The sub-problems

1. An agent might simply value the global state, bypassing MEU altogether.
2. There's no reason to think that in-branch agents are indifferent to the rest of the Everett 'tree' (either in general, or specifically to the fate of their co-descendants within it). This is a reason to think that if the Born weights were (quasi-)probabilities, the expected value of an Everettian game would in general **not** equal that of the corresponding one-world game.
3. In general, rational decision on behalf of a group may be expected **not** to go by a weighted sum of individual utilities, when the welfare of some comes at severe cost to others.
4. Without the credence-existence link, the Everettian seems to lack any rationale for "sizism" – i.e., for giving preference to higher-weighted individuals. (Without such a justification, a proof that there is no rational MEU alternative to a Born-weights-based model might be seem as a reason for abandoning the MEU model altogether.)