THE ROLE OF EVIDENCE IN AN ACCURACY-CENTERED EPISTEMOLOGY FOR CREDENCES

FEW 2013
May 9, 2013
New Brunswick, New Jersey

Jim Joyce
Department of Philosophy
University of Michigan
jjoyce@umich.edu

TWO APPROACHS TO EPISTEMOLOGY

Accuracy-centered. The cardinal epistemic good is the holding of beliefs that accurately reflect the world's state. Believers have a duty to rationally pursue doxastic accuracy. (Compare A. Goldman's "veritistic value".)

Evidence-centered. Believers have an epistemic duty to hold beliefs that are *well-justified* in light of their evidence.

My Aim: To paint a compelling picture of accuracy-centered epistemology in which evidential considerations play a central role, and to refute a recent objection to the accuracy-centered approach.

Basic Line: The rational pursuit of accuracy never requires us to invest more confidence in any proposition than our evidence warrants. Honoring our duty to hold well justified beliefs never interferes with the rational pursuit of accuracy.

TWO APPROACHS TO EPISTEMOLOGY

Accuracy-centered. The cardinal epistemic good is the holding of beliefs that accurately reflect the world's state. Believers have a duty to rationally pursue doxastic accuracy. (Compare A. Goldman's "veritistic value".)

Evidence-centered. Believers have an epistemic duty to hold beliefs that are *well-justified* in light of their evidence.

My Aim: To paint a compelling picture of accuracy-centered epistemology in which evidential considerations play a central role, and to refute a recent objection to the accuracy-centered approach.

Basic Line: The rational pursuit of accuracy never requires us to invest more confidence in any proposition than our evidence warrants. Honoring our duty to hold well justified beliefs never interferes with the rational pursuit of accuracy.

TWO APPROACHS TO EPISTEMOLOGY

Accuracy-centered. The cardinal epistemic good is the holding of beliefs that accurately reflect the world's state. Believers have a duty to rationally pursue doxastic accuracy. (Compare A. Goldman's "veritistic value".)

Evidence-centered. Believers have an epistemic duty to hold beliefs that are *well-justified* in light of their evidence.

My Aim: To paint a compelling picture of accuracy-centered epistemology in which evidential considerations play a central role, and to refute a recent objection to the accuracy-centered approach.

Basic Line: The rational pursuit of accuracy never requires us to invest more confidence in any proposition than our evidence warrants. Honoring our duty to hold well justified beliefs never interferes with the rational pursuit of accuracy.

TWO APPROACHS TO EPISTEMOLOGY

Accuracy-centered. The cardinal epistemic good is the holding of beliefs that accurately reflect the world's state. Believers have a duty to rationally pursue doxastic accuracy. (Compare A. Goldman's "veritistic value".)

Evidence-centered. Believers have an epistemic duty to hold beliefs that are *well-justified* in light of their evidence.

My Aim: To paint a compelling picture of accuracy-centered epistemology in which evidential considerations play a central role, and to refute a recent objection to the accuracy-centered approach.

Basic Line: The rational pursuit of accuracy never requires us to invest more confidence in any proposition than our evidence warrants. Honoring our duty to hold well justified beliefs never interferes with the rational pursuit of accuracy.

ACCURACY-CENTERED EPISTEMOLOGY FOR CREDENCES

- An **inaccuracy score** $\mathfrak{I}$ associates each credal state ***b*** and world ω with a non-negative real number, $\mathfrak{I}$(***b***, ω), which measures ***b*'**s overall inaccuracy when ω is actual (where 0 = perfection).

  Truth-Directedness. Moving credences closer to truth-values improves accuracy.

  Extensionality. The inaccuracy of ***b*** at ω is solely a function of the credences ***b*** assigns and the truth-values ω assigns.

  Continuity. Inaccuracy scores are continuous (for each ω).

  Propriety. If ***b*** is a probability then ***b*** uniquely minimizes expected inaccuracy when expectations are calculated using ***b***.

- A score that meets these conditions captures ***a consistent way of valuing closeness to the truth***.

What kinds of values? Epistemic values!

Accuracy-Centered Epistemology for Credences

- An **inaccuracy score** $\mathfrak{I}$ associates each credal state $\boldsymbol{b}$ and world ω with a non-negative real number, $\mathfrak{I}(\boldsymbol{b}, \omega)$, which measures $\boldsymbol{b}$**'**s overall inaccuracy when ω is actual (where 0 = perfection).

  Truth-Directedness. Moving credences closer to truth-values improves accuracy.

  Extensionality. The inaccuracy of $\boldsymbol{b}$ at ω is solely a function of the credences $\boldsymbol{b}$ assigns and the truth-values ω assigns.

  Continuity. Inaccuracy scores are continuous (for each ω).

  Propriety. If $\boldsymbol{b}$ is a probability then $\boldsymbol{b}$ uniquely minimizes expected inaccuracy when expectations are calculated using $\boldsymbol{b}$.

- A score that meets these conditions captures ***a consistent way of valuing closeness to the truth***.

  What kinds of values? Epistemic values!

ACCURACY-CENTERED EPISTEMOLOGY FOR CREDENCES

- An **inaccuracy score** $\mathfrak{I}$ associates each credal state ***b*** and world ω with a non-negative real number, $\mathfrak{I}$(***b***, ω), which measures ***b'***s overall inaccuracy when ω is actual (where 0 = perfection).

  Truth-Directedness. Moving credences closer to truth-values improves accuracy.

  Extensionality. The inaccuracy of ***b*** at ω is solely a function of the credences ***b*** assigns and the truth-values ω assigns.

  Continuity. Inaccuracy scores are continuous (for each ω).

  Propriety. If ***b*** is a probability then ***b*** uniquely minimizes expected inaccuracy when expectations are calculated using ***b***.

- A score that meets these conditions captures ***a consistent way of valuing closeness to the truth***.

What kinds of values? Epistemic values!

ACCURACY-CENTERED EPISTEMOLOGY FOR CREDENCES

- An **inaccuracy score** $\mathfrak{I}$ associates each credal state ***b*** and world ω with a non-negative real number, $\mathfrak{I}(\boldsymbol{b}, \omega)$, which measures ***b***'s overall inaccuracy when ω is actual (where 0 = perfection).

  Truth-Directedness. Moving credences closer to truth-values improves accuracy.

  Extensionality. The inaccuracy of ***b*** at ω is solely a function of the credences ***b*** assigns and the truth-values ω assigns.

  Continuity. Inaccuracy scores are continuous (for each ω).

  Propriety. If ***b*** is a probability then ***b*** uniquely minimizes expected inaccuracy when expectations are calculated using ***b***.

- A score that meets these conditions captures ***a consistent way of valuing closeness to the truth***.

What kinds of values? Epistemic values!

USEFUL EXAMPLES OF ACCURACY SCORES

$\langle h, t \rangle$ is the state in which a believer assigns credence h to H and t to $\sim H$.

- *Brier*: $\mathscr{B}_1(h, t) = \frac{1}{2}\,[(1 - h)^2 + t^2]$

  $\mathscr{B}_0(h, t) = \frac{1}{2}\,[h^2 + (1 - t)^2]$

- *Square Root**: $\mathbb{S}_1(h, t) = \frac{1}{2}\,[(1 - h)^{1/2} + t^{1/2}]$

  $\mathbb{S}_0(h, t) = \frac{1}{2}\,[h^{1/2} + (1 - t)^{1/2}]$

*not proper

ACCURACY FOR CREDENCES

- Believers have an epistemic duty to hold credences that, by their best estimates in light of their evidence, are likely to strike the optimal balance between the good of being confident in truths and the evil of being confident in falsehoods (where the magnitudes of the goods and evils are measured by an appropriate scoring rule).
    - Believers often must take 'epistemic gambles' by holding credences that are certain to be less than perfectly accurate, as a way of hedging against even greater inaccuracy.

A Fundamental Accuracy Norm

Accuracy-Nondominance (**AN**). It is impermissible, ***whatever one's evidence***, to hold accuracy-dominated credences.

- Def. ***b dominates c*** iff $\mathfrak{I}(\boldsymbol{c}, \omega) > \mathfrak{I}(\boldsymbol{b}, \omega)$ for all worlds ω.

ACCURACY FOR CREDENCES

- Believers have an epistemic duty to hold credences that, by their best estimates in light of their evidence, are likely to strike the optimal balance between the good of being confident in truths and the evil of being confident in falsehoods (where the magnitudes of the goods and evils are measured by an appropriate scoring rule).
  - Believers often must take ‘epistemic gambles’ by holding credences that are certain to be less than perfectly accurate, as a way of hedging against even greater inaccuracy.

A Fundamental Accuracy Norm

Accuracy-Nondominance (**AN**). It is impermissible, ***whatever one's evidence***, to hold accuracy-dominated credences.

- Def. ***b dominates c*** iff $\mathfrak{I}(\boldsymbol{c}, \omega) > \mathfrak{I}(\boldsymbol{b}, \omega)$ for all worlds ω.

ACCURACY FOR CREDENCES

- Believers have an epistemic duty to hold credences that, by their best estimates in light of their evidence, are likely to strike the optimal balance between the good of being confident in truths and the evil of being confident in falsehoods (where the magnitudes of the goods and evils are measured by an appropriate scoring rule).
    - Believers often must take 'epistemic gambles' by holding credences that are certain to be less than perfectly accurate, as a way of hedging against even greater inaccuracy.

A Fundamental Accuracy Norm

Accuracy-Nondominance (**AN**). It is impermissible, ***whatever one's evidence***, to hold accuracy-dominated credences.

- Def. ***b dominates c*** iff $\mathfrak{I}(\boldsymbol{c}, \omega) > \mathfrak{I}(\boldsymbol{b}, \omega)$ for all worlds ω.

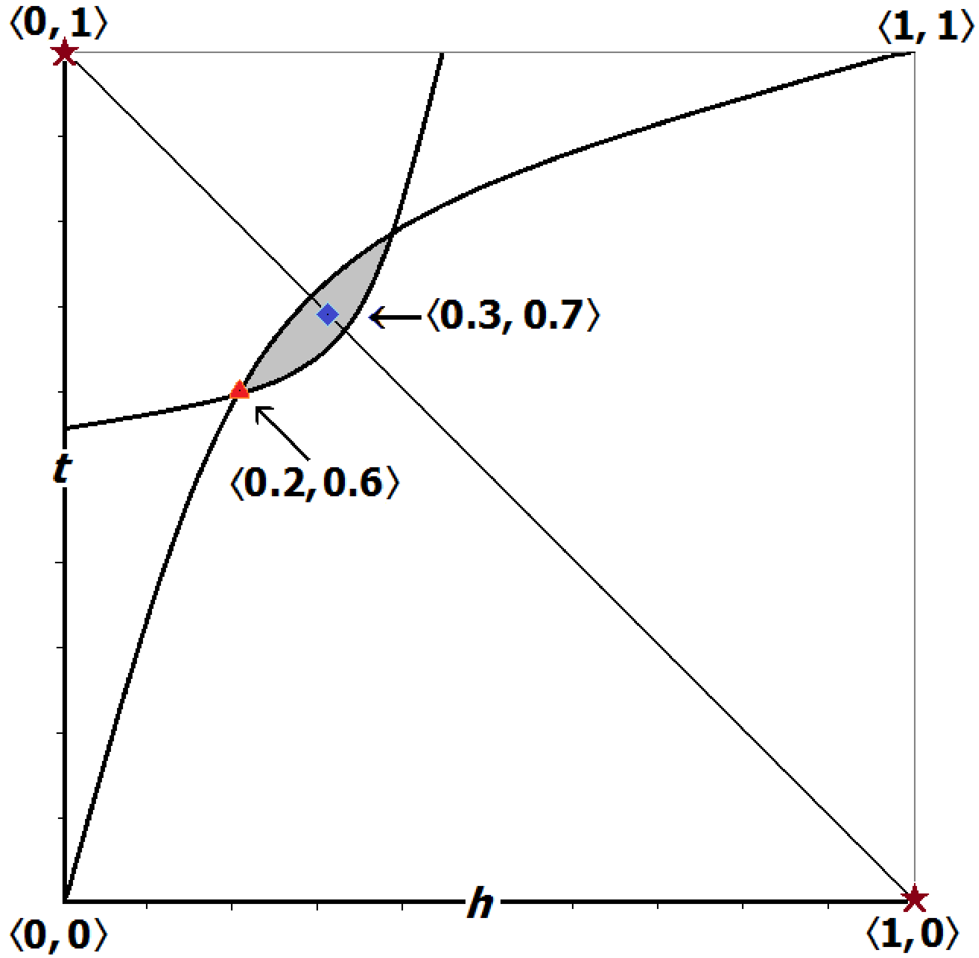

Points ⟨*h*, *t*⟩ are credences for *H* and ~*H*. The points ⟨0, 1⟩ and ⟨1, 0⟩ represent consistent truth-value assignments. The incoherent pair ⟨0.2, 0.6⟩ is accuracy-dominated by the coherent pair ⟨0.3, 0.7⟩.

AN IS NON-NEGOTIABLE

- Just as non-dominance principles are central to the idea that pragmatic value can be represented by utility functions, **AN** is essential to the idea that inaccuracy scores capture coherent ways of valuing 'closeness to truth'.

- **AN** entails that ***c*** is worse than ***b all-epistemic-things-considered*** when ***b*** accuracy-dominates ***c***.
    - *Any* advantage that ***c*** might have over ***b*** in justification is trumped by the fact that ***c*** dominates ***b***.

Problem? What if the totality of the evidence supports a dominated credal state over one that dominates it?

Answer. That would be bad, but (I'll argue) it should never happen if we are doing things right.

AN IS NON-NEGOTIABLE

- Just as non-dominance principles are central to the idea that pragmatic value can be represented by utility functions, **AN** is essential to the idea that inaccuracy scores capture coherent ways of valuing ‘closeness to truth’.

- **AN** entails that ***c*** is worse than ***b all-epistemic-things-considered*** when ***b*** accuracy-dominates ***c***.
  - *Any* advantage that ***c*** might have over ***b*** in justification is trumped by the fact that ***c*** dominates ***b***.

Problem? What if the totality of the evidence supports a dominated credal state over one that dominates it?

Answer. That would be bad, but (I’ll argue) it should never happen if we are doing things right.

AN IS NON-NEGOTIABLE

- Just as non-dominance principles are central to the idea that pragmatic value can be represented by utility functions, **AN** is essential to the idea that inaccuracy scores capture coherent ways of valuing ‘closeness to truth’.

- **AN** entails that ***c*** is worse than ***b all-epistemic-things-considered*** when ***b*** accuracy-dominates ***c***.
  - *Any* advantage that ***c*** might have over ***b*** in justification is trumped by the fact that ***c*** dominates ***b***.

Problem? What if the totality of the evidence supports a dominated credal state over one that dominates it?

Answer. That would be bad, but (I’ll argue) it should never happen if we are doing things right.

AN IS NON-NEGOTIABLE

- Just as non-dominance principles are central to the idea that pragmatic value can be represented by utility functions, **AN** is essential to the idea that inaccuracy scores capture coherent ways of valuing ‘closeness to truth’.

- **AN** entails that ***c*** is worse than ***b all-epistemic-things-considered*** when ***b*** accuracy-dominates ***c***.
    - *Any* advantage that ***c*** might have over ***b*** in justification is trumped by the fact that ***c*** dominates ***b***.

Problem? What if the totality of the evidence supports a dominated credal state over one that dominates it?

Answer. That would be bad, but (I’ll argue) it should never happen if we are doing things right.

Some Evidential Norms for Credences

Truth. If your evidence is that *H* is true, then you should be certain of *H*.

Principal Principle. If your evidence is that the current objective chance of *H* is *h*, and if you have no 'inadmissible' information about *H*, then your evidence requires you to assign a credence of *h* to *H*.

THE CORE CLAIM

A body of evidence provides more justification for one credal state than for another just when it requires a believer to fix a higher estimate for the accuracy of the first state than for that of the second.

- But, when does the evidence requires a believer to "fix a higher estimate for the accuracy of the first state than for that of the second"? Some examples:

    - **Dominance**. If credal state ***b*** has a strictly lower inaccuracy than ***c*** in every world not excluded by the data, then a believer with that data is required to have a higher estimate for ***b***'s accuracy than for ***c***'s.

    - A *probabilistically coherent* believer with credences of $\boldsymbol{b}(H) = x$ and $\boldsymbol{b}(\sim H) = 1 - x$ will estimate the accuracy of any credal state $\langle h, t\rangle$ using its *expected accuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

PP ⇨ If the evidence is $\boldsymbol{chance}(H) = x$ and $\boldsymbol{chance}(\sim H) = 1 - x$, then a rational believer's estimated inaccuracy for $\langle h, t\rangle$ is its *objective expected inaccuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

THE CORE CLAIM

A body of evidence provides more justification for one credal state than for another just when it requires a believer to fix a higher estimate for the accuracy of the first state than for that of the second.

- But, when does the evidence requires a believer to "fix a higher estimate for the accuracy of the first state than for that of the second"? Some examples:
  - **Dominance**. If credal state ***b*** has a strictly lower inaccuracy than ***c*** in every world not excluded by the data, then a believer with that data is required to have a higher estimate for ***b***'s accuracy than for ***c***'s.
  - A *probabilistically coherent* believer with credences of ***b***(H) = x and ***b***($\sim H$) = $1 - x$ will estimate the accuracy of any credal state $\langle h, t \rangle$ using its *expected accuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.
  - **PP** ⇨ If the evidence is ***chance***(H) = x and ***chance***($\sim H$) = $1 - x$, then a rational believer's estimated inaccuracy for $\langle h, t \rangle$ is its *objective expected inaccuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

THE CORE CLAIM

A body of evidence provides more justification for one credal state than for another just when it requires a believer to fix a higher estimate for the accuracy of the first state than for that of the second.

- But, when does the evidence requires a believer to "fix a higher estimate for the accuracy of the first state than for that of the second"? Some examples:

  - **Dominance**. If credal state ***b*** has a strictly lower inaccuracy than ***c*** in every world not excluded by the data, then a believer with that data is required to have a higher estimate for ***b***'s accuracy than for ***c***'s.

  - A *probabilistically coherent* believer with credences of ***b***(H) = x and ***b***($\sim H$) = $1 - x$ will estimate the accuracy of any credal state $\langle h, t\rangle$ using its *expected accuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

  - **PP** ⇨ If the evidence is ***chance***(H) = x and ***chance***($\sim H$) = $1 - x$, then a rational believer's estimated inaccuracy for $\langle h, t\rangle$ is its *objective expected inaccuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

THE CORE CLAIM

A body of evidence provides more justification for one credal state than for another just when it requires a believer to fix a higher estimate for the accuracy of the first state than for that of the second.

- But, when does the evidence requires a believer to "fix a higher estimate for the accuracy of the first state than for that of the second"? Some examples:

  - **Dominance**. If credal state ***b*** has a strictly lower inaccuracy than ***c*** in every world not excluded by the data, then a believer with that data is required to have a higher estimate for ***b***'s accuracy than for ***c***'s.

  - A *probabilistically coherent* believer with credences of ***b***(H) = x and ***b***($\sim H$) = $1 - x$ will estimate the accuracy of any credal state $\langle h, t\rangle$ using its *expected accuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

  - **PP** ⇨ If the evidence is ***chance***(H) = x and ***chance***($\sim H$) = $1 - x$, then a rational believer's estimated inaccuracy for $\langle h, t\rangle$ is its *objective expected inaccuracy*: $x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$.

KEY CONSEQUENCES OF DOMINANCE

★ **Evidence that supports a credal state always provides even more support for any state that accuracy-dominates it.**

★ **Evidence that tells against a state always tells even more strongly against anything that state dominates**.

- These points distill the idea that accuracy-dominated credal states are inferior, *all-epistemic-things-considered*, to states that dominate them.

- They tell us that no conflict between legitimate norms of evidence and the norm of accuracy dominance will ever arise.

Important Consequence: Choosing an inaccuracy score commits us to judgments about which credal states are better justified in light of the available evidence.

KEY CONSEQUENCES OF DOMINANCE

★ **Evidence that supports a credal state always provides even more support for any state that accuracy-dominates it.**

★ **Evidence that tells against a state always tells even more strongly against anything that state dominates**.

- These points distill the idea that accuracy-dominated credal states are inferior, *all-epistemic-things-considered*, to states that dominate them.

- They tell us that no conflict between legitimate norms of evidence and the norm of accuracy dominance will ever arise.

Important Consequence: Choosing an inaccuracy score commits us to judgments about which credal states are better justified in light of the available evidence.

KEY CONSEQUENCES OF DOMINANCE

- ★ **Evidence that supports a credal state always provides even more support for any state that accuracy-dominates it.**

- ★ **Evidence that tells against a state always tells even more strongly against anything that state dominates**.

  - These points distill the idea that accuracy-dominated credal states are inferior, *all-epistemic-things-considered*, to states that dominate them.

  - They tell us that no conflict between legitimate norms of evidence and the norm of accuracy dominance will ever arise.

Important Consequence: Choosing an inaccuracy score commits us to judgments about which credal states are better justified in light of the available evidence.

AN INSTRUCTIVE EXAMPLE

- The evidence: ***chance***(*H*) = 0.2 & ***chance***(~*H*) = 0.8.
- Let's compare the support this data provides for $\langle 0.2, 0.6\rangle$ and for $\langle 0.3, 0.7\rangle$ by asking which is worse (justification-wise):
  - having a credence for *H* that agrees perfectly with *H*'s chance and a credence for ~*H* that falls 0.2 probiles below ~*H*'s chance, or
  - having a credence for *H* that is 0.1 probilie above *H*'s chance and a credence for ~*H* that falls 0.1 probile below ~*H*'s chance.

The answer depends on how we measure inaccuracy!

Brier Score: $\mathbf{\mathit{Exp}}(\mathscr{B}(0.2, 0.6)) = 0.2\cdot\mathscr{B}_1(0.2, 0.6) + 0.8\cdot\mathscr{B}_0(0.2, 0.6) = 0.18$

$\mathbf{\mathit{Exp}}(\mathscr{B}(0.3, 0.7)) = 0.2\cdot\mathscr{B}_1(0.3, 0.7) + 0.8\cdot\mathscr{B}_0(0.3, 0.7) = \mathbf{0.17}$

Root Score: $\mathbf{\mathit{Exp}}(\mathscr{R}(0.2, 0.6)) = 0.2\cdot\mathscr{R}_1(0.2, 0.6) + 0.8\cdot\mathscr{R}_0(0.2, 0.6) = \mathbf{0.5988}$

$\mathbf{\mathit{Exp}}(\mathscr{R}(0.3, 0.7)) = 0.2\cdot\mathscr{R}_1(0.3, 0.7) + 0.8\cdot\mathscr{R}_0(0.3, 0.7) = 0.6055$

An Instructive Example

- The evidence: ***chance***(H) = 0.2 & ***chance***($\sim H$) = 0.8.
- Let's compare the support this data provides for $\langle 0.2, 0.6\rangle$ and for $\langle 0.3, 0.7\rangle$ by asking which is worse (justification-wise):
  - having a credence for H that agrees perfectly with H's chance and a credence for $\sim H$ that falls 0.2 probiles below $\sim H$'s chance, or
  - having a credence for H that is 0.1 probilie above H's chance and a credence for $\sim H$ that falls 0.1 probile below $\sim H$'s chance.

The answer depends on how we measure inaccuracy!

Brier Score: $\boldsymbol{Exp}(\mathscr{B}(0.2, 0.6)) = 0.2\cdot\mathscr{B}_1(0.2, 0.6) + 0.8\cdot\mathscr{B}_0(0.2, 0.6) = 0.18$

$\boldsymbol{Exp}(\mathscr{B}(0.3, 0.7)) = 0.2\cdot\mathscr{B}_1(0.3, 0.7) + 0.8\cdot\mathscr{B}_0(0.3, 0.7) = \mathbf{0.17}$

Root Score: $\boldsymbol{Exp}(\mathscr{R}(0.2, 0.6)) = 0.2\cdot\mathscr{R}_1(0.2, 0.6) + 0.8\cdot\mathscr{R}_0(0.2, 0.6) = \mathbf{0.5988}$

$\boldsymbol{Exp}(\mathscr{R}(0.3, 0.7)) = 0.2\cdot\mathscr{R}_1(0.3, 0.7) + 0.8\cdot\mathscr{R}_0(0.3, 0.7) = 0.6055$

PROBLEM: IT ALL DEPENDS ON THE ACCURACY SCORE!

- The conclusion that ⟨0.3, 0.7⟩ is better justified than ⟨0.2, 0.6⟩ hangs on our use of a score, like ***Brier***, that has ⟨0.3, 0.7⟩ dominating ⟨0.2, 0.6⟩.

- The conclusion that **PP** advocates well-justified credences relies on the use of a ***proper*** scoring rule.

- But, there are ways of measuring 'closeness to the truth' relative to which ⟨0.2, 0.6⟩ dominates ⟨0.3, 0.7⟩, and that contravene **PP**.

- Is there any non-question-begging way to show that the **Central Claim** and **AN** will never conflict with any legitimate evidential norm?

    - We can exclude Square Root because it is improper, but isn't that just an *ad hoc* maneuver designed to get justification relations to work out the way we want? **Answer**: Yes, unabashedly so!

Problem: It All Depends on the Accuracy Score!

- The conclusion that ⟨0.3, 0.7⟩ is better justified than ⟨0.2, 0.6⟩ hangs on our use of a score, like ***Brier***, that has ⟨0.3, 0.7⟩ dominating ⟨0.2, 0.6⟩.
- The conclusion that **PP** advocates well-justified credences relies on the use of a ***proper*** scoring rule.
- But, there are ways of measuring 'closeness to the truth' relative to which ⟨0.2, 0.6⟩ dominates ⟨0.3, 0.7⟩, and that contravene **PP**.
- Is there any non-question-begging way to show that the **Central Claim** and **AN** will never conflict with any legitimate evidential norm?
    - We can exclude Square Root because it is improper, but isn't that just an *ad hoc* maneuver designed to get justification relations to work out the way we want? **Answer**: Yes, unabashedly so!

PROBLEM: IT ALL DEPENDS ON THE ACCURACY SCORE!

- The conclusion that ⟨0.3, 0.7⟩ is better justified than ⟨0.2, 0.6⟩ hangs on our use of a score, like ***Brier***, that has ⟨0.3, 0.7⟩ dominating ⟨0.2, 0.6⟩.
- The conclusion that **PP** advocates well-justified credences relies on the use of a ***proper*** scoring rule.
- But, there are ways of measuring 'closeness to the truth' relative to which ⟨0.2, 0.6⟩ dominates ⟨0.3, 0.7⟩, and that contravene **PP**.
- Is there any non-question-begging way to show that the **Central Claim** and **AN** will never conflict with any legitimate evidential norm?
    - We can exclude Square Root because it is improper, but isn't that just an *ad hoc* maneuver designed to get justification relations to work out the way we want? **Answer**: Yes, unabashedly so!

A Wrong Way to Think About the Problem

- The problem assumes that the goal is to develop a free-standing theory of accuracy, *untainted by evidential considerations*, that rules out scores like Square Root and rules in requirements like Propriety.

- This is the wrong. Evidential considerations are in the picture from the start!
  - The relationship between epistemic norms and accuracy norms is symbiotic, not hierarchical. Epistemic norms are not *derived* from accuracy norms, they *cohere* with them.
  - Inaccuracy scores are ways of measuring 'closeness to truth' that reflect our considered views about how such closeness should valued. Part of our goal in choosing a score is to promote correct epistemic values.
  - From this perspective, it ***is*** entirely legitimate to reject some putative accuracy score because it entails that ***b*** dominates ***c*** when we know ***c*** is better justified than ***b***.

A Wrong Way to Think About the Problem

- The problem assumes that the goal is to develop a free-standing theory of accuracy, *untainted by evidential considerations*, that rules out scores like Square Root and rules in requirements like Propriety.

- This is the wrong. Evidential considerations are in the picture from the start!
    - The relationship between epistemic norms and accuracy norms is symbiotic, not hierarchical. Epistemic norms are not *derived* from accuracy norms, they *cohere* with them.
    - Inaccuracy scores are ways of measuring 'closeness to truth' that reflect our considered views about how such closeness should valued. Part of our goal in choosing a score is to promote correct epistemic values.
    - From this perspective, it ***is*** entirely legitimate to reject some putative accuracy score because it entails that ***b*** dominates ***c*** when we know ***c*** is better justified than ***b***.

A Wrong Way to Think About the Problem

- The problem assumes that the goal is to develop a free-standing theory of accuracy, *untainted by evidential considerations*, that rules out scores like Square Root and rules in requirements like Propriety.

- This is the wrong. Evidential considerations are in the picture from the start!

  - The relationship between epistemic norms and accuracy norms is symbiotic, not hierarchical. Epistemic norms are not *derived* from accuracy norms, they *cohere* with them.

  - Inaccuracy scores are ways of measuring 'closeness to truth' that reflect our considered views about how such closeness should valued. Part of our goal in choosing a score is to promote correct epistemic values.

  - From this perspective, it ***is*** entirely legitimate to reject some putative accuracy score because it entails that ***b*** dominates ***c*** when we know ***c*** is better justified than ***b***.

A Wrong Way to Think About the Problem

- The problem assumes that the goal is to develop a free-standing theory of accuracy, *untainted by evidential considerations*, that rules out scores like Square Root and rules in requirements like Propriety.

- This is the wrong. Evidential considerations are in the picture from the start!

    - The relationship between epistemic norms and accuracy norms is symbiotic, not hierarchical. Epistemic norms are not *derived* from accuracy norms, they *cohere* with them.

    - Inaccuracy scores are ways of measuring 'closeness to truth' that reflect our considered views about how such closeness should valued. Part of our goal in choosing a score is to promote correct epistemic values.

    - From this perspective, it ***is*** entirely legitimate to reject some putative accuracy score because it entails that ***b*** dominates ***c*** when we know ***c*** is better justified than ***b***.

A WRONG WAY TO THINK ABOUT THE PROBLEM

- The problem assumes that the goal is to develop a free-standing theory of accuracy, *untainted by evidential considerations*, that rules out scores like Square Root and rules in requirements like Propriety.

- This is the wrong. Evidential considerations are in the picture from the start!

  - The relationship between epistemic norms and accuracy norms is symbiotic, not hierarchical. Epistemic norms are not *derived* from accuracy norms, they *cohere* with them.

  - Inaccuracy scores are ways of measuring 'closeness to truth' that reflect our considered views about how such closeness should valued. Part of our goal in choosing a score is to promote correct epistemic values.

  - From this perspective, it ***is*** entirely legitimate to reject some putative accuracy score because it entails that ***b*** dominates ***c*** when we know ***c*** is better justified than ***b***.

THE ACCURACY SCORE AS CONSISTENCY TEST

In developing an "accuracy-centered" epistemology for credences, the goal is to show that various evidential norms – **Truth**, **PP**,… – are jointly consistent with the idea that accuracy is the cardinal epistemic good.

- To achieve this goal one must prove that there exists an inaccuracy score $\mathscr{I}$ such that:
  - ★ No norm permits $\mathscr{I}$-dominated credences in any evidential situation.
  - ★ No norm prohibits the credences ***b*** in any evidential situation unless it also prohibits any credences that are $\mathscr{I}$-dominates by ***b***.
  - ★ More generally, no norm permits ***c*** when the available evidence requires the believer to fix a higher estimate for the accuracy of another state ***b***.

Key Point: In proving that the required $\mathscr{I}$ exists one, in effect, shows that there is at least one way of valuing accuracy that reflects the epistemic values that the evidential norms incorporate.

The Accuracy Score as Consistency Test

In developing an "accuracy-centered" epistemology for credences, the goal is to show that various evidential norms – **Truth**, **PP**,… – are jointly consistent with the idea that accuracy is the cardinal epistemic good.

- To achieve this goal one must prove that there exists an inaccuracy score $\mathscr{I}$ such that:
  - ★ No norm permits $\mathscr{I}$-dominated credences in any evidential situation.
  - ★ No norm prohibits the credences ***b*** in any evidential situation unless it also prohibits any credences that are $\mathscr{I}$-dominates by ***b***.
  - ★ More generally, no norm permits ***c*** when the available evidence requires the believer to fix a higher estimate for the accuracy of another state ***b***.

Key Point: In proving that the required $\mathscr{I}$ exists one, in effect, shows that there is at least one way of valuing accuracy that reflects the epistemic values that the evidential norms incorporate.

THE ACCURACY SCORE AS CONSISTENCY TEST

In developing an "accuracy-centered" epistemology for credences, the goal is to show that various evidential norms – **Truth**, **PP**,… – are jointly consistent with the idea that accuracy is the cardinal epistemic good.

- To achieve this goal one must prove that there exists an inaccuracy score $\mathscr{I}$ such that:
  - ★ No norm permits $\mathscr{I}$-dominated credences in any evidential situation.
  - ★ No norm prohibits the credences ***b*** in any evidential situation unless it also prohibits any credences that are $\mathscr{I}$-dominates by ***b***.
  - ★ More generally, no norm permits ***c*** when the available evidence requires the believer to fix a higher estimate for the accuracy of another state ***b***.

Key Point: In proving that the required $\mathscr{I}$ exists one, in effect, shows that there is at least one way of valuing accuracy that reflects the epistemic values that the evidential norms incorporate.

EXAMPLE: 🌩🌩 ABSOLUTE VALUE AND SQUARE ROOT SCORES 🌩🌩

Imagine you have rolled a die 1000 times and have observed these outcomes:

ONE	TWO	THREE	FOUR	FIVE	SIX
100	100	100	200	200	300

Evidential Norm: Based on this evidence, $\langle 0.1, 0.1, 0.1, 0.2, 0.2, 0.3\rangle$ is better justified than $\langle 0, 0, 0, 0.1, 0.1, 0.2\rangle$.

- If we measure inaccuracy using the absolute value score the latter credal state dominates the former! For this score the only undominated credences are those for which some *N* has credence zero. Same for Square Root.
  - To see this, note that when $0 < b_1 \leq b_N$ for all N = ONE, TWO,…, SIX one has $\mathcal{a}_N(\langle b_n\rangle) - \mathcal{a}_N(\langle b_n - b_1\rangle) = 4{\cdot}b_1 > 0$

Here our norms of evidence inform and constrain our accuracy scores.

EXAMPLE: EPISTEMIC EXPERTS

- Imagine a believer, with a probabilistically coherent credence function $\boldsymbol{b}$, who treats some source of (probabilistically coherent) information $\boldsymbol{q}$ as an *expert* about a proposition X.

  Def. $\boldsymbol{q}$ is an **expert** for $\boldsymbol{b}$ about X just in case $\boldsymbol{b}(X \mid \boldsymbol{q}(X) = x) = x$.

  - This means that the believer will defer to $\boldsymbol{q}$'s values when she knows what those values are. (Note: ★ and the Accuracy Argument require $\boldsymbol{q}$ to be a *probability*.)

- An **expert principle** is an evidential norm of the form:

  If you know $\boldsymbol{q}(X) = x$, and if this is all your relevant evidence about X, then x should be your credence for X.

  Examples:

  - $\boldsymbol{q}(X) = @(X) =$ X's actual truth value (always an expert)
  - $\boldsymbol{q}(X) = \boldsymbol{chance}_{now}(X) =$ X's current chance
  - $\boldsymbol{q}(X) = \boldsymbol{chance}_{later}(X) =$ X's chance in an hour

ACCURACY AND THE HIERARCHY OF EPISTEMIC EXPERTS

The rub in any expert principle is the “this is all your evidence about *X*” clause, which can be problematic when experts give conflicting advice.

Def. ***q*** *trumps* ***r*** for ***b*** exactly if $\boldsymbol{b}(X \mid \boldsymbol{q}(X) = x \;\&\; \boldsymbol{r}(X) = y) = x$ for all x, y.

E.g., @ trumps everything, later chances trump earlier chances.

Question: Under what conditions will ***q*** trump ***r*** for ***b***?

- Intuitively, the trumping expert should be the one that ***b*** expects to be most accurate. This turns out to be true!

Fact: Given any accuracy score that satisfies Truth-directedness and Propriety, ***q*** *trumps* ***r*** for ***b*** only if ***b***’s estimate of ***r***’s inaccuracy exceeds ***b***’s estimate of ***q***’s inaccuracy.

- So, this *one* aspect of the theory of epistemic experts can be subsumed into the accuracy-based framework: the expert ‘pecking order’ goes by increasing expected accuracy.

MICHAEL CAIE'S COUNTEREXAMPLE TO ACCURACY DOMINANCE

Caie (2013): "Considerations of accuracy support the claim that an agent may rationally fail to have probabilistically coherent credences."

(#) My current credence in the proposition # expressed by the sentence (#) is not greater than ½.

Notation. ***b*** is my credence function *rigidly designated*. β is my credence function *non-rigidly designated*. # is the proposition $\beta(\#) \ngeq$ ½. # is true just in case either I have no definite credence for # or I have a credence *x* and that number is less than ½.

Caie Claims:

I. Considerations of accuracy dictate that I should assign credence ½ to # and credence 1 to ~#.

II. Since my credence for # *influences* #'s truth-value, accuracy dominance does not apply, and the argument for probabilism is nullified.

MICHAEL CAIE'S COUNTEREXAMPLE TO ACCURACY DOMINANCE

Caie (2013): "Considerations of accuracy support the claim that an agent may rationally fail to have probabilistically coherent credences."

(#) My current credence in the proposition # expressed by the sentence (#) is not greater than ½.

Notation. ***b*** is my credence function *rigidly designated*. β is my credence function *non-rigidly designated*. # is the proposition $\beta(\#) \ngeq ½$. # is true just in case either I have no definite credence for # or I have a credence *x* and that number is less than ½.

Caie Claims:

I. Considerations of accuracy dictate that I should assign credence ½ to # and credence 1 to ~#.

II. Since my credence for # *influences* #'s truth-value, accuracy dominance does not apply, and the argument for probabilism is nullified.

Michael Caie's Counterexample to Accuracy Dominance

Caie (2013): "Considerations of accuracy support the claim that an agent may rationally fail to have probabilistically coherent credences."

(#) My current credence in the proposition # expressed by the sentence (#) is not greater than ½.

Notation. ***b*** is my credence function *rigidly designated*. β is my credence function *non-rigidly designated*. # is the proposition $\beta(\#) \ngeq ½$. # is true just in case either I have no definite credence for # or I have a credence *x* and that number is less than ½.

Caie Claims:

I. Considerations of accuracy dictate that I should assign credence ½ to # and credence 1 to ~#.

II. Since my credence for # *influences* #'s truth-value, accuracy dominance does not apply, and the argument for probabilism is nullified.

DE FACTO ACCURACY VERSUS SELF-ACCURACY

- Let prospective credences be given by pairs ⟨***b***(#), ***b***(~#)⟩ = ⟨*x*, *y*⟩.

***De facto* accuracies,** measured relative to a *fixed* truth-value for #, with arrows pointing toward greater accuracy

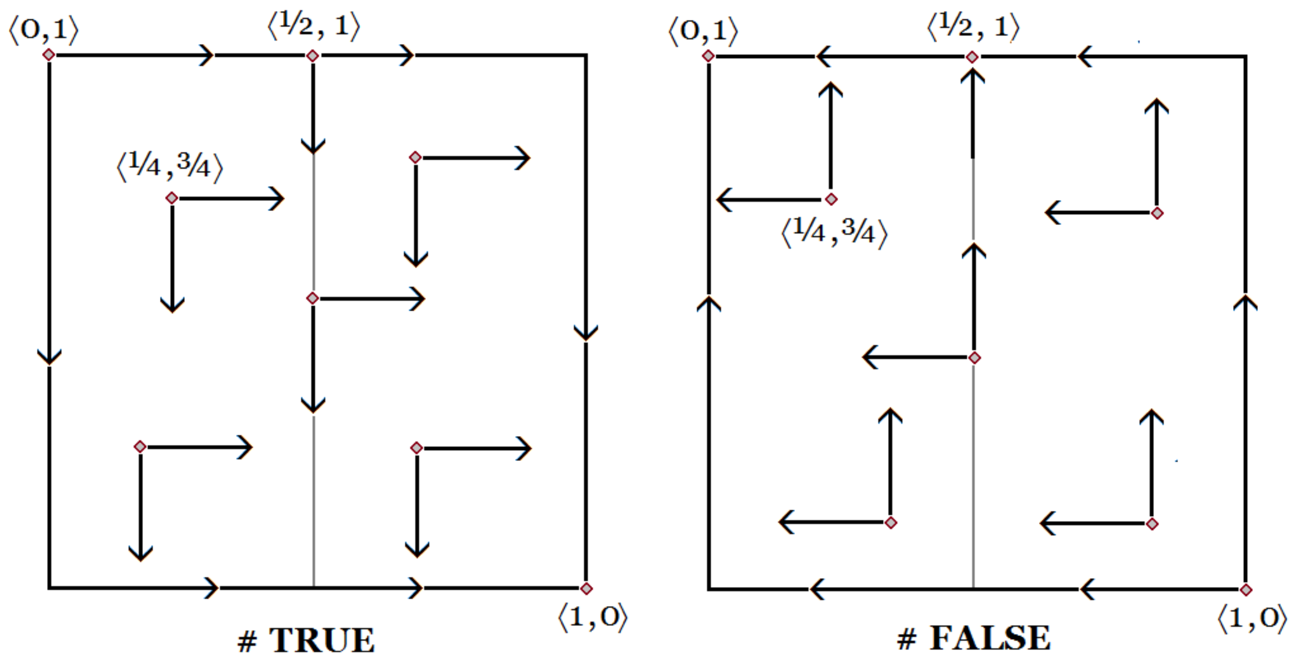

b(#) < ½ or ***b***(#) not defined ***b***(#) ≥ ½

FACTS ABOUT DE FACTO ACCURACY

- If we measure inaccuracy using the Brier score, then ⟨¼, ¾⟩ has a lower *de facto* inaccuracy than ⟨½, 1⟩ whether # is true or false.

This suggests a **dominance argument**:

- ➢ If I have the credences ⟨½, 1⟩, I will see one or both of # and ~# as live epistemic possibilities.
- ➢ If # is true, then ⟨¼, ¾⟩ is more accurate than ⟨½, 1⟩.
- ➢ If # is false, then ⟨¼, ¾⟩ is more accurate than ⟨½, 1⟩.

So, holding ⟨½, 1⟩ commits me to credences that are less accurate than ⟨¼, ¾⟩ in every world I regard as possible, which is supposed to be irrational.

Caie's Claim: This dominance reasoning is invalid!

FACTS ABOUT DE FACTO ACCURACY

- If we measure inaccuracy using the Brier score, then ⟨¼, ¾⟩ has a lower *de facto* inaccuracy than ⟨½, 1⟩ whether # is true or false.

This suggests a **dominance argument**:

- If I have the credences ⟨½, 1⟩, I will see one or both of # and ~# as live epistemic possibilities.
- If # is true, then ⟨¼, ¾⟩ is more accurate than ⟨½, 1⟩.
- If # is false, then ⟨¼, ¾⟩ is more accurate than ⟨½, 1⟩.

So, holding ⟨½, 1⟩ commits me to credences that are less accurate than ⟨¼, ¾⟩ in every world I regard as possible, which is supposed to be irrational.

Caie's Claim: This dominance reasoning is invalid!

FACTS ABOUT DE FACTO ACCURACY

- If we measure inaccuracy using the Brier score, then ⟨¼, ¾⟩ has a lower *de facto* inaccuracy than ⟨½, 1⟩ whether # is true or false.

This suggests a dominance argument:

- ➢ If I have the credences ⟨½, 1⟩, I will see one or both of # and ~# as live epistemic possibilities.
- ➢ If # is true, then ⟨¼, ¾⟩ is more accurate than ⟨½, 1⟩.
- ➢ If # is false, then ⟨¼, ¾⟩ is more accurate than ⟨½, 1⟩.

So, holding ⟨½, 1⟩ commits me to credences that are less accurate than ⟨¼, ¾⟩ in every world I regard as possible, which is supposed to be irrational.

Caie's Claim: This dominance reasoning is invalid!

AN EPISTEMIC DECISION

	#	~#
Adopt ⟨½, 1⟩	6	14
Adopt ⟨¼, ¾⟩	7	15

Values = 16 x (1 – Brier)

- Decision theory tells us that dominance arguments are ***not*** valid when the choice of an act affects the state of the world.

- Caie: Since adopting ⟨½, 1⟩ *makes* # false while adopting ⟨¼, ¾⟩ *makes* it true, the shaded boxes are not genuine possibilities.

	#	~#
Adopt ⟨½, 1⟩		14
Adopt ⟨¼, ¾⟩	7	

AN EPISTEMIC DECISION

	#	~#
Adopt ⟨½, 1⟩	6	14
Adopt ⟨¼, ¾⟩	7	15

Values = 16 x (1 – Brier)

- Decision theory tells us that dominance arguments are ***not*** valid when the choice of an act affects the state of the world.

- Caie: Since adopting ⟨½, 1⟩ *makes* # false while adopting ⟨¼, ¾⟩ *makes* it true, the shaded boxes are not genuine possibilities.

	#	~#
Adopt ⟨½, 1⟩		14
Adopt ⟨¼, ¾⟩	7	

An Epistemic Decision

	#	~#
Adopt ⟨½, 1⟩	6	14
Adopt ⟨¼, ¾⟩	7	15

Values = 16 x (1 – Brier)

- Decision theory tells us that dominance arguments are ***not*** valid when the choice of an act affects the state of the world.

- Caie: Since adopting ⟨½, 1⟩ *makes* # false while adopting ⟨¼, ¾⟩ *makes* it true, the shaded boxes are not genuine possibilities.

	#	~#
Adopt ⟨½, 1⟩		14
Adopt ⟨¼, ¾⟩	7	

SELF-ACCURACY

- Caie asks: How accurate *would* ⟨*x*, *y*⟩ be *were* they my credences?

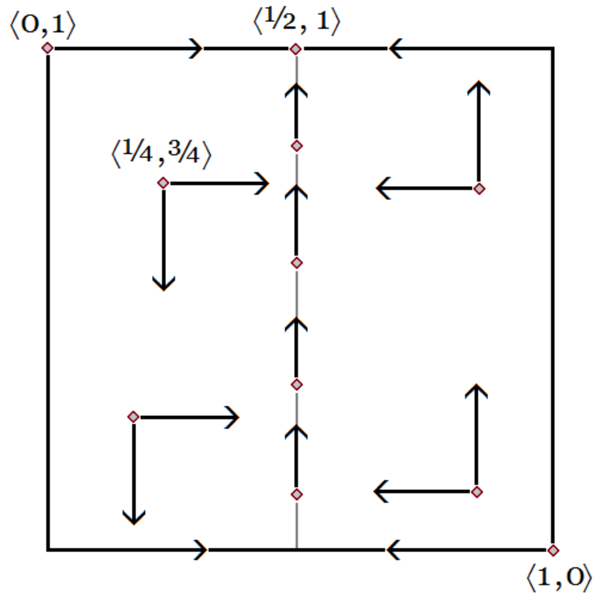

FACT: For any reasonable measure of inaccuracy, minimum self-inaccuracy is (uniquely) attained at ⟨½, 1⟩.

Adopting Credences as Decision Making

Caie's sees the pursuit of doxastic accuracy as a kind of 'choice problem' in which a believer, who assesses things from a bird's eye perspective *in which she has not yet committed to any beliefs,* selects the credences that minimize inaccuracy on the assumption that they are selected.

- Because adopting a credal state determines a truth-value for #, Caie says, that the accuracy-based approach is committed to assessing each potential credal state as an instrument for securing low inaccuracy scores.

- My ability to *choose* the truth-values of # and ~# makes this a *decision under certainty*. I can figure out what to do without recourse to my (introspective) evidence about what I believe about # and ~#.

ADOPTING CREDENCES AS DECISION MAKING

Caie's sees the pursuit of doxastic accuracy as a kind of 'choice problem' in which a believer, who assesses things from a bird's eye perspective *in which she has not yet committed to any beliefs,* selects the credences that minimize inaccuracy on the assumption that they are selected.

- Because adopting a credal state determines a truth-value for #, Caie says, that the accuracy-based approach is committed to assessing each potential credal state as an instrument for securing low inaccuracy scores.

- My ability to *choose* the truth-values of # and ~# makes this a *decision under certainty*. I can figure out what to do without recourse to my (introspective) evidence about what I believe about # and ~#.

ADOPTING CREDENCES AS DECISION MAKING

Caie's sees the pursuit of doxastic accuracy as a kind of 'choice problem' in which a believer, who assesses things from a bird's eye perspective *in which she has not yet committed to any beliefs,* selects the credences that minimize inaccuracy on the assumption that they are selected.

- Because adopting a credal state determines a truth-value for #, Caie says, that the accuracy-based approach is committed to assessing each potential credal state as an instrument for securing low inaccuracy scores.

- My ability to *choose* the truth-values of # and ~# makes this a *decision under certainty*. I can figure out what to do without recourse to my (introspective) evidence about what I believe about # and ~#.

EX ANTE AND EX POST

We must be careful to distinguish the "act" of **adopting** credences from the subsequent state of **having** those credences.

> **($\#_t$)** At time t, my credence in the proposition $\#_t$ expressed by the sentence ($\#_t$) is not greater than ½.

- Imagine that I can adopt credences by taking a pill marked $\langle x, y\rangle$ at time $t = 0$, whose sole effect will be to cause me to have the credences $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1({\sim}\#_1) = y$ at a later time $t = 1$.
  - There is no problem with being certain about $\#_1$'s truth-value at $t = 0$. If I decide to take the $\langle ½, 1\rangle$ pill at $t = 0$ I will become sure that $\#_1$ is false.
    - There is, of course, a big problem with being certain about $\#_1$ at $t = 1$ since knowing at t =1 that you are certain of $\#_1$ means that you know ~ $\#_1$.
    - Following Caie, I am going to assume that nothing prevents a believer from knowing her credences for $\#_1$ at $t = 1$ (though I don't actually believe this).

EX ANTE AND EX POST

We must be careful to distinguish the "act" of **adopting** credences from the subsequent state of **having** those credences.

> **($\#_t$)** At time t, my credence in the proposition $\#_t$ expressed by the sentence ($\#_t$) is not greater than ½.

- Imagine that I can adopt credences by taking a pill marked $\langle x, y\rangle$ at time $t = 0$, whose sole effect will be to cause me to have the credences $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1({\sim}\#_1) = y$ at a later time $t = 1$.
  - There is no problem with being certain about $\#_1$'s truth-value at $t = 0$. If I decide to take the $\langle ½, 1\rangle$ pill at $t = 0$ I will become sure that $\#_1$ is false.
    - There is, of course, a big problem with being certain about $\#_1$ at $t = 1$ since knowing at t =1 that you are certain of $\#_1$ means that you know $\sim \#_1$.
    - Following Caie, I am going to assume that nothing prevents a believer from knowing her credences for $\#_1$ at $t = 1$ (though I don't actually believe this).

EX ANTE AND EX POST

We must be careful to distinguish the "act" of **adopting** credences from the subsequent state of **having** those credences.

> **($\#_t$)** At time t, my credence in the proposition $\#_t$ expressed by the sentence ($\#_t$) is not greater than ½.

- Imagine that I can adopt credences by taking a pill marked $\langle x, y\rangle$ at time $t = 0$, whose sole effect will be to cause me to have the credences $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1(\sim\#_1) = y$ at a later time $t = 1$.
  - There is no problem with being certain about $\#_1$'s truth-value at $t = 0$. If I decide to take the $\langle ½, 1\rangle$ pill at $t = 0$ I will become sure that $\#_1$ is false.
    - There is, of course, a big problem with being certain about $\#_1$ at $t = 1$ since knowing at t =1 that you are certain of $\#_1$ means that you know ~ $\#_1$.
    - Following Caie, I am going to assume that nothing prevents a believer from knowing her credences for $\#_1$ at $t = 1$ (though I don't actually believe this).

EX ANTE AND EX POST

We must be careful to distinguish the “act” of **adopting** credences from the subsequent state of **having** those credences.

($\#_t$) At time t, my credence in the proposition $\#_t$ expressed by the sentence ($\#_t$) is not greater than ½.

- Imagine that I can adopt credences by taking a pill marked $\langle x, y\rangle$ at time $t = 0$, whose sole effect will be to cause me to have the credences $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1({\sim}\#_1) = y$ at a later time $t = 1$.
  - There is no problem with being certain about $\#_1$’s truth-value at $t = 0$. If I decide to take the $\langle ½, 1\rangle$ pill at $t = 0$ I will become sure that $\#_1$ is false.
    - There is, of course, a big problem with being certain about $\#_1$ at $t = 1$ since knowing at t =1 that you are certain of $\#_1$ means that you know ~ $\#_1$.
    - Following Caie, I am going to assume that nothing prevents a believer from knowing her credences for $\#_1$ at $t = 1$ (though I don’t actually believe this).

CAIE'S CLAIM I REVISITED

- Caie maintains this:

  $\mathbf{I}_A$ The accuracy-centered view says that I am rationally obliged to take the ⟨½, 1⟩ pill at $t = 0$ because this maximizes self-accuracy.

- He must also maintain at least one of the following:

  $\mathbf{I}_B$ If it is rational to take the ⟨½, 1⟩ pill at $t = 0$, then it is rational at $t = 1$ to hold the ⟨½, 1⟩ credences.

  $\mathbf{I}_C$ If it is rational to take the ⟨½, 1⟩ pill at $t = 0$, then it is rational at $t = 0$ to endorse* ⟨½, 1⟩ as the optimal credal state at $t = 1$.

$\mathbf{I}_C$ is where the action is!

- Since the $t = 1$ credences undermine themselves (below), the only reason to hold $\mathbf{I}_B$ is the $t = 1$ judgment that any *change* in my credences will be worse, which is basically the same issues that $\mathbf{I}_C$ raises moved along in time.

Caie's Claim I Revisited

- Caie maintains this:

  I$_A$ The accuracy-centered view says that I am rationally obliged to take the ⟨½, 1⟩ pill at t = 0 because this maximizes self-accuracy.

- He must also maintain at least one of the following:

  I$_B$ If it is rational to take the ⟨½, 1⟩ pill at t = 0, then it is rational at t = 1 to hold the ⟨½, 1⟩ credences.

  I$_C$ If it is rational to take the ⟨½, 1⟩ pill at t = 0, then it is rational at t = 0 to endorse* ⟨½, 1⟩ as the optimal credal state at t = 1.

I$_C$ is where the action is!

- Since the t = 1 credences undermine themselves (below), the only reason to hold **I**$_B$ is the t = 1 judgment that any *change* in my credences will be worse, which is basically the same issues that **I**$_C$ raises moved along in time.

CAIE'S CLAIM I REVISITED

- Caie maintains this:

  $\mathbf{I}_A$ The accuracy-centered view says that I am rationally obliged to take the ⟨½, 1⟩ pill at t = 0 because this maximizes self-accuracy.

- He must also maintain at least one of the following:

  $\mathbf{I}_B$ If it is rational to take the ⟨½, 1⟩ pill at t = 0, then it is rational at t = 1 to hold the ⟨½, 1⟩ credences.

  $\mathbf{I}_C$ If it is rational to take the ⟨½, 1⟩ pill at t = 0, then it is rational at t = 0 to endorse* ⟨½, 1⟩ as the optimal credal state at t = 1.

$\mathbf{I}_C$ is where the action is!

- Since the t = 1 credences undermine themselves (below), the only reason to hold $\mathbf{I}_B$ is the t = 1 judgment that any *change* in my credences will be worse, which is basically the same issues that $\mathbf{I}_C$ raises moved along in time.

An Instructive Parallel: De Finetti's Prevision Game

A proposition M is specified. You are given \$2 on the understanding that you must announce two 'previsions' $x, y \in [0, 1]$ and will then have to repay a sum of $\$[(1 - x)^2 + y^2]$ if M is true and $\$[x^2 + (1 - y)^2]$ if M is false.

- Note how the size of your 'penalty' depends on both the actual truth-values, M and $\sim M$, and the 'previsions' you announce.

Usual Claim: A rational expected payoff maximizer's previsions reveal her credences, so that $\boldsymbol{b}(M) = x$ and $\boldsymbol{b}(\sim M) = y$.

- That's not *exactly* right. To derive the desired conclusion we must suppose, as De Finetti implicitly did, that the subject cannot *control* M's truth or falsity.

- When this independence condition fails we can***not*** safely identify previsions with credences.

An Instructive Parallel: De Finetti's Prevision Game

A proposition M is specified. You are given \$2 on the understanding that you must announce two 'previsions' $x, y \in [0, 1]$ and will then have to repay a sum of $\$[(1 - x)^2 + y^2]$ if M is true and $\$[x^2 + (1 - y)^2]$ if M is false.

- Note how the size of your 'penalty' depends on both the actual truth-values, M and $\sim M$, and the 'previsions' you announce.

Usual Claim: A rational expected payoff maximizer's previsions reveal her credences, so that $\boldsymbol{b}(M) = x$ and $\boldsymbol{b}(\sim M) = y$.

- That's not *exactly* right. To derive the desired conclusion we must suppose, as De Finetti implicitly did, that the subject cannot *control* M's truth or falsity.

- When this independence condition fails we can***not*** safely identify previsions with credences.

AN INSTRUCTIVE PARALLEL: DE FINETTI'S PREVISION GAME

A proposition M is specified. You are given \$2 on the understanding that you must announce two 'previsions' x, $y \in [0, 1]$ and will then have to repay a sum of $\$[(1 - x)^2 + y^2]$ if M is true and $\$[x^2 + (1 - y)^2]$ if M is false.

- Note how the size of your 'penalty' depends on both the actual truth-values, M and $\sim M$, and the 'previsions' you announce.

Usual Claim: A rational expected payoff maximizer's previsions reveal her credences, so that $\boldsymbol{b}(M) = x$ and $\boldsymbol{b}(\sim M) = y$.

- That's not *exactly* right. To derive the desired conclusion we must suppose, as De Finetti implicitly did, that the subject cannot *control* M's truth or falsity.

- When this independence condition fails we can***not*** safely identify previsions with credences.

AN INSTRUCTIVE PARALLEL: DE FINETTI'S PREVISION GAME

A proposition M is specified. You are given \$2 on the understanding that you must announce two 'previsions' $x, y \in [0, 1]$ and will then have to repay a sum of $\$[(1 - x)^2 + y^2]$ if M is true and $\$[x^2 + (1 - y)^2]$ if M is false.

- Note how the size of your 'penalty' depends on both the actual truth-values, M and $\sim M$, and the 'previsions' you announce.

Usual Claim: A rational expected payoff maximizer's previsions reveal her credences, so that $\boldsymbol{b}(M) = x$ and $\boldsymbol{b}(\sim M) = y$.

- That's not *exactly* right. To derive the desired conclusion we must suppose, as De Finetti implicitly did, that the subject cannot *control* M's truth or falsity.

- When this independence condition fails we can***not*** safely identify previsions with credences.

THE PREVISION GAME WITHOUT ACT/STATE INDEPENDENCE

Suppose you play the prevision game with

M = The prevision you announce will ***not*** be ½ or greater.

- You walk away with $[2$x – (x^2 + y^2)$] if *M* and $[2$y – (x^2 + y^2)$] if ~*M*.
  - So, the value you announce for *x* determines which payoff schedule is used, but this is *not* influenced by your value for *y*.

- You minimize your penalty by announcing x = ½ and y = 1. (Look familiar?)

- But, ½ is clearly *not* your credence for *M*. When you announce *x* you ***know*** *M* is false! You decided to make it so! (And ½ need not have been your pre-decision credence either.)

What is going on here?

THE PREVISION GAME WITHOUT ACT/STATE INDEPENDENCE

Suppose you play the prevision game with

M = The prevision you announce will ***not*** be ½ or greater.

- You walk away with $[2x – ($x^2 + y^2$)] if M and $[2y – ($x^2 + y^2$)] if ~$M$.
  - So, the value you announce for x determines which payoff schedule is used, but this is *not* influenced by your value for y.

- You minimize your penalty by announcing x = ½ and y = 1. (Look familiar?)

- But, ½ is clearly *not* your credence for M. When you announce x you ***know*** M is false! You decided to make it so! (And ½ need not have been your pre-decision credence either.)

What is going on here?

The Prevision Game Without Act/State Independence

Suppose you play the prevision game with

M = The prevision you announce will ***not*** be ½ or greater.

- You walk away with $[2$x$ – ($x^2 + y^2$)] if M and $[2$y$ – ($x^2 + y^2$)] if ~M.
  - So, the value you announce for x determines which payoff schedule is used, but this is *not* influenced by your value for y.

- You minimize your penalty by announcing x = ½ and y = 1. (Look familiar?)

- But, ½ is clearly *not* your credence for M. When you announce x you ***know*** M is false! You decided to make it so! (And ½ need not have been your pre-decision credence either.)

What is going on here?

THE PREVISION GAME WITHOUT ACT/STATE INDEPENDENCE

Suppose you play the prevision game with

M = The prevision you announce will ***not*** be ½ or greater.

- You walk away with $[2$x$ – ($x^2 + y^2$)] if M and $[2$y$ – ($x^2 + y^2$)] if ~M.
  - So, the value you announce for x determines which payoff schedule is used, but this is *not* influenced by your value for y.

- You minimize your penalty by announcing x = ½ and y = 1. (Look familiar?)

- But, ½ is clearly *not* your credence for M. When you announce x you ***know*** M is false! You decided to make it so! (And ½ need not have been your pre-decision credence either.)

What is going on here?

A DIFFERENCE BETWEEN *x* AND *y*

x does not reflect your credence for *M*, *y does* reflect your credence for ~*M*.

- Because you can manipulate *M*'s truth-value by your choice of *x*, the *x*-value you announce does ***not*** represent your best estimate of *M*'s truth-value. It is also influenced by 'strategic considerations'.
  - Whatever your credence for *M*, the setup makes it impossible for you to announce an *x* that is closer than ½ to *M*'s actual truth-value.
  - So, you must move *x* away from your actual credence for *M* in an effort to manipulate *M*'s truth-value.
- Because *M* is **not** influenced by *y*, the game's payoff structure gives you an incentive to announce your best estimate of ~*M*'s truth-value.
  - Thus, your prevision for ~*M* is not contaminated by the sorts of 'strategic' considerations that figure into your choice of a prevision for *M*.

A DIFFERENCE BETWEEN *x* AND *y*

x does not reflect your credence for *M*, *y does* reflect your credence for ~*M*.

- Because you can manipulate *M*'s truth-value by your choice of *x*, the *x*-value you announce does ***not*** represent your best estimate of *M*'s truth-value. It is also influenced by 'strategic considerations'.
  - Whatever your credence for *M*, the setup makes it impossible for you to announce an *x* that is closer than ½ to *M*'s actual truth-value.
  - So, you must move *x* away from your actual credence for *M* in an effort to manipulate *M*'s truth-value.
- Because *M* is **not** influenced by *y*, the game's payoff structure gives you an incentive to announce your best estimate of ~*M*'s truth-value.
  - Thus, your prevision for ~*M* is not contaminated by the sorts of 'strategic' considerations that figure into your choice of a prevision for *M*.

A DIFFERENCE BETWEEN *x* AND *y*

x does not reflect your credence for *M*, *y does* reflect your credence for ~*M*.

- Because you can manipulate *M*'s truth-value by your choice of *x*, the *x*-value you announce does ***not*** represent your best estimate of *M*'s truth-value. It is also influenced by 'strategic considerations'.
    - Whatever your credence for *M*, the setup makes it impossible for you to announce an *x* that is closer than ½ to *M*'s actual truth-value.
    - So, you must move *x* away from your actual credence for *M* in an effort to manipulate *M*'s truth-value.
- Because *M* is **not** influenced by *y*, the game's payoff structure gives you an incentive to announce your best estimate of ~*M*'s truth-value.
    - Thus, your prevision for ~*M* is not contaminated by the sorts of 'strategic' considerations that figure into your choice of a prevision for *M*.

The Upshot for I_C

Just as announcing x = ½ says zilch about my credence for M when M's truth-value depends on what prevision I announce, so choosing the ⟨½, 1⟩ pill says zilch about the normative status of ½ as a credence for $\#_1$.

- Upon deciding to take pill ⟨½, 1⟩ at $t = 0$, my credence for $\#_1$ goes to ***zero***, and I see the effect of taking the pill as ***de**creasing* my accuracy.
  - Once I decide to take the ⟨½, 1⟩ pill I have the same evidence about $\#_1$'s truth-value that my $t = 1$ self will have.
  - Given that it is impossible for me to get my inaccuracy for $\boldsymbol{b}_1(\#_1)$ below ¼, ⟨½, 1⟩ leaves me *least* inaccurate.
  - But, I still see ⟨½, 1⟩, and every other choice I have the ability to make, as ***de**creasing* my accuracy!

THE UPSHOT FOR $\mathrm{I_C}$

Just as announcing $x = ½$ says zilch about my credence for *M* when *M*'s truth-value depends on what prevision I announce, so choosing the $\langle ½, 1\rangle$ pill says zilch about the normative status of ½ as a credence for $\#_1$.

- Upon deciding to take pill $\langle ½, 1\rangle$ at $t = 0$, my credence for $\#_1$ goes to ***zero***, and I see the effect of taking the pill as ***de****creasing* my accuracy.
  - Once I decide to take the $\langle ½, 1\rangle$ pill I have the same evidence about $\#_1$'s truth-value that my $t = 1$ self will have.
  - Given that it is impossible for me to get my inaccuracy for $\boldsymbol{b}_1(\#_1)$ below ¼, $\langle ½, 1\rangle$ leaves me *least* inaccurate.
  - But, I still see $\langle ½, 1\rangle$, and every other choice I have the ability to make, as ***de****creasing* my accuracy!

THE UPSHOT FOR $\mathbf{I}_C$

Just as announcing x = ½ says zilch about my credence for M when M's truth-value depends on what prevision I announce, so choosing the ⟨½, 1⟩ pill says zilch about the normative status of ½ as a credence for $\#_1$.

- Upon deciding to take pill ⟨½, 1⟩ at $t = 0$, my credence for $\#_1$ goes to ***zero***, and I see the effect of taking the pill as ***de****creasing* my accuracy.
  - Once I decide to take the ⟨½, 1⟩ pill I have the same evidence about $\#_1$'s truth-value that my $t = 1$ self will have.
  - Given that it is impossible for me to get my inaccuracy for $\boldsymbol{b}_1(\#_1)$ below ¼, ⟨½, 1⟩ leaves me *least* inaccurate.
  - But, I still see ⟨½, 1⟩, and every other choice I have the ability to make, as ***de****creasing* my accuracy!

THE GENERAL POINT

- Choosing a future credence for *M* with the aim of minimizing self-accuracy reliably yields a credence for *M* that one *currently* endorses (given the *post-choice* information) ***only when*** the credence chosen reflects one's estimate of *M*'s truth-value.
  - One's choice can only be counted upon to reflect one's best estimate of *M*'s truth-value when the choice does ***not*** causally influence *M*'s truth-value.
  - But, when the "choice" of a future credence can influence *M*'s truth-value the fact that a certain choice minimizes self-accuracy does not entail that those credences are the rational ones to adopt.

Distinguish:

"Rational credences maximize estimated *de facto*-accuracy."

"Anything that maximizes estimated self-accuracy is a rational credence."

Accuracy-based epistemology is committed only to the FIRST claim!

THE GENERAL POINT

- Choosing a future credence for *M* with the aim of minimizing self-accuracy reliably yields a credence for *M* that one *currently* endorses (given the *post-choice* information) ***only when*** the credence chosen reflects one's estimate of *M*'s truth-value.
  - One's choice can only be counted upon to reflect one's best estimate of *M*'s truth-value when the choice does ***not*** causally influence *M*'s truth-value.
  - But, when the "choice" of a future credence can influence *M*'s truth-value the fact that a certain choice minimizes self-accuracy does not entail that those credences are the rational ones to adopt.

Distinguish:

"Rational credences maximize estimated *de facto*-accuracy."

"Anything that maximizes estimated self-accuracy is a rational credence."

Accuracy-based epistemology is committed only to the FIRST claim!

THE GENERAL POINT

- Choosing a future credence for *M* with the aim of minimizing self-accuracy reliably yields a credence for *M* that one *currently* endorses (given the *post-choice* information) ***only when*** the credence chosen reflects one's estimate of *M*'s truth-value.
  - One's choice can only be counted upon to reflect one's best estimate of *M*'s truth-value when the choice does ***not*** causally influence *M*'s truth-value.
  - But, when the "choice" of a future credence can influence *M*'s truth-value the fact that a certain choice minimizes self-accuracy does not entail that those credences are the rational ones to adopt.

Distinguish:

"Rational credences maximize estimated *de facto*-accuracy."

"Anything that maximizes estimated self-accuracy is a rational credence."

Accuracy-based epistemology is committed only to the FIRST claim!

ACCURACY DOMINANCE REVISITED

Caie is right that the accuracy dominance argument does nothing to discredit your $t = 0$ *choice* of future credences for $\#_1$ and $\sim\#_1$, but it does discredit those credences themselves!

- At $t = 1$ and there is no longer anything you can do to change the fact that $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1(\sim\#_1) = y$. (You took the $\langle x, y\rangle$ pill.)

- But, dominance considerations ***do*** apply when the relevant truth-values can no longer be altered.

- Thus, at $t = 1$ we can invoke the dominating position of $\langle 1/4, 3/4\rangle$ to show that $\langle 1/2, 1\rangle$ is epistemically defective (even though it might have been entirely rational to choose the latter credences over the former at a time when one still has the power to decide which of $\#_1$ or $\sim\#_1$ would be true).

RESIST the temptation to reply "if my credences were $\langle 1/4, 3/4\rangle$ I'd be even less accurate than I am now!" True, but the appeal to dominance is ***not*** intended to suggest that you should hold $\langle 1/4, 3/4\rangle$ instead of $\langle 1/4, 3/4\rangle$ (now or ever). It is only meant to show that you should *not* hold $\langle 1/2, 1\rangle$.

ACCURACY DOMINANCE REVISITED

Caie is right that the accuracy dominance argument does nothing to discredit your $t = 0$ *choice* of future credences for $\#_1$ and $\sim\#_1$, but it does discredit those credences themselves!

- At $t = 1$ and there is no longer anything you can do to change the fact that $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1(\sim\#_1) = y$. (You took the $\langle x, y\rangle$ pill.)
- But, dominance considerations ***do*** apply when the relevant truth-values can no longer be altered.
- Thus, at $t = 1$ we can invoke the dominating position of $\langle ¼, ¾\rangle$ to show that $\langle ½, 1\rangle$ is epistemically defective (even though it might have been entirely rational to choose the latter credences over the former at a time when one still has the power to decide which of $\#_1$ or $\sim\#_1$ would be true).
- RESIST the temptation to reply "if my credences were $\langle ¼, ¾\rangle$ I'd be even less accurate than I am now!" True, but the appeal to dominance is ***not*** intended to suggest that you should hold $\langle ¼, ¾\rangle$ instead of $\langle ¼, ¾\rangle$ (now or ever). It is only meant to show that you should *not* hold $\langle ½, 1\rangle$.

ACCURACY DOMINANCE REVISITED

Caie is right that the accuracy dominance argument does nothing to discredit your $t = 0$ *choice* of future credences for $\#_1$ and $\sim\#_1$, but it does discredit those credences themselves!

- At $t = 1$ and there is no longer anything you can do to change the fact that $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1(\sim\#_1) = y$. (You took the $\langle x, y\rangle$ pill.)
- But, dominance considerations ***do*** apply when the relevant truth-values can no longer be altered.
- Thus, at $t = 1$ we can invoke the dominating position of $\langle ¼, ¾\rangle$ to show that $\langle ½, 1\rangle$ is epistemically defective (even though it might have been entirely rational to choose the latter credences over the former at a time when one still has the power to decide which of $\#_1$ or $\sim\#_1$ would be true).

RESIST the temptation to reply "if my credences were $\langle ¼, ¾\rangle$ I'd be even less accurate than I am now!" True, but the appeal to dominance is ***not*** intended to suggest that you should hold $\langle ¼, ¾\rangle$ instead of $\langle ¼, ¾\rangle$ (now or ever). It is only meant to show that you should *not* hold $\langle ½, 1\rangle$.

ACCURACY DOMINANCE REVISITED

Caie is right that the accuracy dominance argument does nothing to discredit your $t = 0$ *choice* of future credences for $\#_1$ and $\sim\#_1$, but it does discredit those credences themselves!

- At $t = 1$ and there is no longer anything you can do to change the fact that $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1(\sim\#_1) = y$. (You took the $\langle x, y\rangle$ pill.)

- But, dominance considerations ***do*** apply when the relevant truth-values can no longer be altered.

- Thus, at $t = 1$ we can invoke the dominating position of $\langle \frac{1}{4}, \frac{3}{4}\rangle$ to show that $\langle \frac{1}{2}, 1\rangle$ is epistemically defective (even though it might have been entirely rational to choose the latter credences over the former at a time when one still has the power to decide which of $\#_1$ or $\sim\#_1$ would be true).

RESIST the temptation to reply "if my credences were $\langle \frac{1}{4}, \frac{3}{4}\rangle$ I'd be even less accurate than I am now!" True, but the appeal to dominance is ***not*** intended to suggest that you should hold $\langle \frac{1}{4}, \frac{3}{4}\rangle$ instead of $\langle \frac{1}{4}, \frac{3}{4}\rangle$ (now or ever). It is only meant to show that you should *not* hold $\langle \frac{1}{2}, 1\rangle$.

ACCURACY DOMINANCE REVISITED

Caie is right that the accuracy dominance argument does nothing to discredit your t = 0 *choice* of future credences for $\#_1$ and $\sim\#_1$, but it does discredit those credences themselves!

- At t = 1 and there is no longer anything you can do to change the fact that $\boldsymbol{b}_1(\#_1) = x$ and $\boldsymbol{b}_1(\sim\#_1) = y$. (You took the $\langle x, y\rangle$ pill.)
- But, dominance considerations ***do*** apply when the relevant truth-values can no longer be altered.
- Thus, at t = 1 we can invoke the dominating position of ⟨¼, ¾⟩ to show that ⟨½, 1⟩ is epistemically defective (even though it might have been entirely rational to choose the latter credences over the former at a time when one still has the power to decide which of $\#_1$ or $\sim\#_1$ would be true).
- RESIST the temptation to reply "if my credences were ⟨¼, ¾⟩ I'd be even less accurate than I am now!" True, but the appeal to dominance is ***not*** intended to suggest that you should hold ⟨¼, ¾⟩ instead of ⟨¼, ¾⟩ (now or ever). It is only meant to show that you should *not* hold ⟨½, 1⟩.

JUST SAY NO TO THE PILLS!

So, should I have probabilistically coherent credences $\boldsymbol{b}_1(\#_1) = 1 - \boldsymbol{b}_1(\sim\#_1)$?

- No! If I am even minimally aware of what my credences are, my evidence will tell me that my credences are far from the most accurate ones to hold.
  - E.g., as Caie shows, if $\boldsymbol{b}_1(\#_1) \ngeq ½$ and $\boldsymbol{b}_1(\beta(\#_1) \ngeq ½) > ½$ then I can deduce that $\#_1$ is false (and thus that my credence is way too high).

- For a minimally self-aware believer, *every* definite credal state for $\#_1$ will be defective. In every such state the believer will possess evidence that justifies her in thinking that some other state has higher *de facto* accuracy.

- The only credal state that does not sin against accuracy in this way, and the one that strikes me as obviously right in the situation, is the one in which $\#_1$ is true because the believer invests no credence whatsoever in its truth (not a sharp one, not an imprecise one, nothing).

Unlike in the Prevision Game, where I offered you $2, there is no epistemic reason at all to play Caie's game.

JUST SAY NO TO THE PILLS!

So, should I have probabilistically coherent credences $\boldsymbol{b}_1(\#_1) = 1 - \boldsymbol{b}_1({\sim}\#_1)$?

- No! If I am even minimally aware of what my credences are, my evidence will tell me that my credences are far from the most accurate ones to hold.
  - E.g., as Caie shows, if $\boldsymbol{b}_1(\#_1) \ngeq ½$ and $\boldsymbol{b}_1(\beta(\#_1) \ngeq ½) > ½$ then I can deduce that $\#_1$ is false (and thus that my credence is way too high).
- For a minimally self-aware believer, *every* definite credal state for $\#_1$ will be defective. In every such state the believer will possess evidence that justifies her in thinking that some other state has higher *de facto* accuracy.
- The only credal state that does not sin against accuracy in this way, and the one that strikes me as obviously right in the situation, is the one in which $\#_1$ is true because the believer invests no credence whatsoever in its truth (not a sharp one, not an imprecise one, nothing).
  - Unlike in the Prevision Game, where I offered you $2, there is no epistemic reason at all to play Caie's game.

JUST SAY NO TO THE PILLS!

So, should I have probabilistically coherent credences $\boldsymbol{b}_1(\#_1) = 1 - \boldsymbol{b}_1({\sim}\#_1)$?

- No! If I am even minimally aware of what my credences are, my evidence will tell me that my credences are far from the most accurate ones to hold.
  - E.g., as Caie shows, if $\boldsymbol{b}_1(\#_1) \ngeq ½$ and $\boldsymbol{b}_1(\beta(\#_1) \ngeq ½) > ½$ then I can deduce that $\#_1$ is false (and thus that my credence is way too high).

- For a minimally self-aware believer, *every* definite credal state for $\#_1$ will be defective. In every such state the believer will possess evidence that justifies her in thinking that some other state has higher *de facto* accuracy.

- The only credal state that does not sin against accuracy in this way, and the one that strikes me as obviously right in the situation, is the one in which $\#_1$ is true because the believer invests no credence whatsoever in its truth (not a sharp one, not an imprecise one, nothing).
  - Unlike in the Prevision Game, where I offered you $2, there is no epistemic reason at all to play Caie's game.

Just Say No To The Pills!

So, should I have probabilistically coherent credences $\boldsymbol{b}_1(\#_1) = 1 - \boldsymbol{b}_1(\sim\#_1)$?

- No! If I am even minimally aware of what my credences are, my evidence will tell me that my credences are far from the most accurate ones to hold.
    - E.g., as Caie shows, if $\boldsymbol{b}_1(\#_1) \not\geq ½$ and $\boldsymbol{b}_1(\beta(\#_1) \not\geq ½) > ½$ then I can deduce that $\#_1$ is false (and thus that my credence is way too high).

- For a minimally self-aware believer, *every* definite credal state for $\#_1$ will be defective. In every such state the believer will possess evidence that justifies her in thinking that some other state has higher *de facto* accuracy.

- The only credal state that does not sin against accuracy in this way, and the one that strikes me as obviously right in the situation, is the one in which $\#_1$ is true because the believer invests no credence whatsoever in its truth (not a sharp one, not an imprecise one, nothing).
    - Unlike in the Prevision Game, where I offered you $2, there is no epistemic reason at all to play Caie's game.

JUST SAY NO TO THE PILLS!

So, should I have probabilistically coherent credences $\boldsymbol{b}_1(\#_1) = 1 - \boldsymbol{b}_1(\sim\#_1)$?

- No! If I am even minimally aware of what my credences are, my evidence will tell me that my credences are far from the most accurate ones to hold.
  - E.g., as Caie shows, if $\boldsymbol{b}_1(\#_1) \ngeq ½$ and $\boldsymbol{b}_1(\beta(\#_1) \ngeq ½) > ½$ then I can deduce that $\#_1$ is false (and thus that my credence is way too high).

- For a minimally self-aware believer, *every* definite credal state for $\#_1$ will be defective. In every such state the believer will possess evidence that justifies her in thinking that some other state has higher *de facto* accuracy.

- The only credal state that does not sin against accuracy in this way, and the one that strikes me as obviously right in the situation, is the one in which $\#_1$ is true because the believer invests no credence whatsoever in its truth (not a sharp one, not an imprecise one, nothing).
  - Unlike in the Prevision Game, where I offered you \$2, there is no epistemic reason at all to play Caie's game.

USEFUL REFERENCES

Michael Caie, “Rational Probabilistic Incoherence” https://dl.dropboxusercontent.com/u/8245565/RPI_PR_Rev_09_08_12.pdf

K. Easwaran and B. Fitelson (2012) “An ‘Evidentialist’ Worry about Joyce’s Argument for Probabilism,” *Dialectica* **66**: 425-433.

Gibbard, A. (2008) “Rational Credence and the Value of Truth,” in T. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol. 2. Oxford: Clarendon Press.

Goldman, Alvin (2010) “Epistemic Relativism and Reasonable Disagreement,” in Richard Feldman & Ted Warfield (eds.), *Disagreement*. London and New York: Oxford University Press.

Lindley, D. (1982) “Scoring Rules and the Inevitability of Probability,” *International Statistical Review* 50: 1–26.

Joyce, J. (1998) “A Non-Pragmatic Vindication of Probabilism,” *Philosophy of Science* 65: 575–603.

Joyce, J. (2009) “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief,” in Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*. Synthese (2009)

Lewis, D. (1980) “A Subjectivist’s Guide to Objective Chance,” in *Studies in Inductive Logic and Probability*, edited by R. Jeffrey, vol. 2. Berkeley: University of California Press: 263–94.

J. Predd, R. Seiringer, E. H. Lieb, D. Osherson, V. Poor, and S. Kulkarni (2009). “Probabilistic Coherence and Proper Scoring Rules.” *IEEE Transactions on Information Theory*, **55**(10): 4786–4792, 2009.

Pettigrew, Richard (forthcoming) “A New Epistemic Utility Argument for the Principal Principle,” *Episteme*