Knowledge Centered Epistemic Utility Theory

Julien Dutant[1] Branden Fitelson[2]

[1]Philosophy @ KCL

julien.dutant@kcl.ac.uk

[2]Philosophy & Religion @ Northeastern

branden@fitelson.org

Outline

1. Prelude: Two Examples and an Overview of the Talk
2. Truth & The Old Lockeanism
3. Knowledge & A New Lockeanism
4. Extras: Some Supplementary Slides (hopefully useful in discussion)

- **Miners** [34, 26]. You are standing in front of two mine shafts (*A* and *B*). Flood waters are approaching. You know that ten miners are in one of the shafts, but you don't know which (*e.g.*, their location was determined by the result of a fair coin toss). You have enough sand bags to block one of the shafts. If the miners are in *A*, then blocking *A* saves all 10 miners (and, hence, minimizes disutility, *i.e.*, # of dead miners). If the miners are in *B*, then blocking *B* minimizes disutility. If you block neither *A* nor *B*, the water will be divided, and only the lowest miner in the shaft will die.
  Claim. *It is rationally permissible to block neither A nor B*.
- **Gibbard's Coin** [14, 30]. A fair coin has been tossed (and you have no information about how it landed). If it landed Heads (*H*), then believing *H* is the attitude which minimizes (epistemic) disutility (*viz.*, *inaccuracy*). If it landed Tails (*T*), then believing *T* is the attitude which minimizes inaccuracy.
  Claim. *It is rationally permissible to believe neither H nor T*.

☞ *It can be rationally permissible to (knowingly) occupy a state, which does **not** minimize disutility — in **any** possible world.*

- **Miners** [34, 26]. You are standing in front of two mine shafts (*A* and *B*). Flood waters are approaching. You know that ten miners are in one of the shafts, but you don't know which (*e.g.*, their location was determined by the result of a fair coin toss). You have enough sand bags to block one of the shafts. If the miners are in *A*, then blocking *A* saves all 10 miners (and, hence, minimizes disutility, *i.e.*, # of dead miners). If the miners are in *B*, then blocking *B* minimizes disutility. If you block neither *A* nor *B*, the water will be divided, and only the lowest miner in the shaft will die. **Claim**. *It is rationally permissible to block neither A nor B*.
- **Gibbard's Coin** [14, 30]. A fair coin has been tossed (and you have no information about how it landed). If it landed Heads (*H*), then believing *H* is the attitude which minimizes (epistemic) disutility (*viz.*, *inaccuracy*). If it landed Tails (*T*), then believing *T* is the attitude which minimizes inaccuracy. **Claim**. *It is rationally permissible to believe neither H nor T*.

☞ *It can be rationally permissible to (knowingly) occupy a state, which does **not** minimize disutility — in **any** possible world.*

- **Miners** [34, 26]. You are standing in front of two mine shafts (*A* and *B*). Flood waters are approaching. You know that ten miners are in one of the shafts, but you don't know which (*e.g.*, their location was determined by the result of a fair coin toss). You have enough sand bags to block one of the shafts. If the miners are in *A*, then blocking *A* saves all 10 miners (and, hence, minimizes disutility, *i.e.*, # of dead miners). If the miners are in *B*, then blocking *B* minimizes disutility. If you block neither *A* nor *B*, the water will be divided, and only the lowest miner in the shaft will die. **Claim**. *It is rationally permissible to block neither A nor B*.
- **Gibbard's Coin** [14, 30]. A fair coin has been tossed (and you have no information about how it landed). If it landed Heads (*H*), then believing *H* is the attitude which minimizes (epistemic) disutility (*viz.*, *inaccuracy*). If it landed Tails (*T*), then believing *T* is the attitude which minimizes inaccuracy. **Claim**. *It is rationally permissible to believe neither H nor T*.

☞ *It can be rationally permissible to (knowingly) occupy a state, which does **not** minimize disutility — in **any** possible world.*

- **Miners** [34, 26]. You are standing in front of two mine shafts (*A* and *B*). Flood waters are approaching. You know that ten miners are in one of the shafts, but you don't know which (*e.g.*, their location was determined by the result of a fair coin toss). You have enough sand bags to block one of the shafts. If the miners are in *A*, then blocking *A* saves all 10 miners (and, hence, minimizes disutility, *i.e.*, # of dead miners). If the miners are in *B*, then blocking *B* minimizes disutility. If you block neither *A* nor *B*, the water will be divided, and only the lowest miner in the shaft will die. **Claim**. *It is rationally permissible to block neither A nor B*.
- **Gibbard's Coin** [14, 30]. A fair coin has been tossed (and you have no information about how it landed). If it landed Heads (H), then believing H is the attitude which minimizes (epistemic) disutility (*viz.*, *inaccuracy*). If it landed Tails (T), then believing T is the attitude which minimizes inaccuracy. **Claim**. *It is rationally permissible to believe neither* H *nor* T.

☞ *It can be rationally permissible to (knowingly) occupy a state, which does **not** minimize disutility — in **any** possible world.*

- **Miners** [34, 26]. You are standing in front of two mine shafts (*A* and *B*). Flood waters are approaching. You know that ten miners are in one of the shafts, but you don't know which (*e.g.*, their location was determined by the result of a fair coin toss). You have enough sand bags to block one of the shafts. If the miners are in *A*, then blocking *A* saves all 10 miners (and, hence, minimizes disutility, *i.e.*, # of dead miners). If the miners are in *B*, then blocking *B* minimizes disutility. If you block neither *A* nor *B*, the water will be divided, and only the lowest miner in the shaft will die. **Claim**. *It is rationally permissible to block neither A nor B*.
- **Gibbard's Coin** [14, 30]. A fair coin has been tossed (and you have no information about how it landed). If it landed Heads (*H*), then believing *H* is the attitude which minimizes (epistemic) disutility (*viz.*, *inaccuracy*). If it landed Tails (*T*), then believing *T* is the attitude which minimizes inaccuracy. **Claim**. *It is rationally permissible to believe neither H nor T*.

☞ *It can be rationally permissible to (knowingly) occupy a state, which does **not** minimize disutility — in **any** possible world.*

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).
  - (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).
    - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.
  - (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.
    - So, we'll *not* be discussing any *diachronic* [40] requirements.
  - (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).
  - (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.
  - (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).
  - (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).
    - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.
  - (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent *S at a single time t*.
    - So, we'll *not* be discussing any *diachronic* [40] requirements.
  - (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).
  - (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.
  - (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).
  - (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).
    - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.
  - (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.
    - So, we'll *not* be discussing any *diachronic* [40] requirements.
  - (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).
  - (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.
  - (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).
  (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).
    - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.

  (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.
    - So, we'll *not* be discussing any *diachronic* [40] requirements.

  (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).

  (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.

  (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).
  - (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).
    - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.
  - (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.
    - So, we'll *not* be discussing any *diachronic* [40] requirements.
  - (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).
  - (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.
  - (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).
  (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).
    - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.

  (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.
    - So, we'll *not* be discussing any *diachronic* [40] requirements.

  (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).

  (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.

  (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).

  (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).

  - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.

  (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.

  - So, we'll *not* be discussing any *diachronic* [40] requirements.

  (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).

  - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).

  (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.

  (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).

  (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).

  - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.

  (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.

  - So, we'll *not* be discussing any *diachronic* [40] requirements.

  (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).

  - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).

  (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.

  (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Today's talk is about (i) formal, (ii) synchronic, (iii) epistemic (iv) coherence (v) requirements (of ideal rationality).

  (i) *Formal* coherence is to be distinguished from other sorts of coherence discussed in contemporary epistemology (*e.g.*, in some empirical, truth/knowledge-conducive sense [1]).

  - Our notions of coherence will supervene on *logical* (and *formal probabilistic*) properties of judgment sets.

  (ii) *Synchronic* coherence has to do with the coherence of a set of judgments held by an agent S *at a single time* t.

  - So, we'll *not* be discussing any *diachronic* [40] requirements.

  (iii) *Epistemic* coherence involves *distinctively* epistemic values (*e.g.*, *accuracy* [19], *evidential support* [7], *knowledge* [Meno]).

  - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [38], and the like [17]).

  (iv) *Coherence* has to do with how a set of judgments "hangs together". CRs are *wide-scope* [3], global requirements.

  (v) *Requirements* are *evaluative*; they give *necessary* conditions for (ideal) epistemic rationality of a doxastic state [40].

- Here is a — perhaps *the* — "paradigm" CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — "paradigm" CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — “paradigm” CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if *S* is Alethically Ideal, then *S*’s beliefs are consistent).
  - **Alethic Ideal** (for belief). *S* should (*alethically, ideally*) believe (disbelieve) that *p* just in case *p* is true (false).
- We’ve already seen (**Gibbard’s Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). *S* should (*evidentially, ideally*) believe (disbelieve) *p* if *S*’s total evidence supports (counter-supports) *p*. Otherwise, *S* should *suspend* on *p*.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We’ll refer to the claim that there exist *some* such cases as *the datum*. Foley’s [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — “paradigm” CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — "paradigm" CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — "paradigm" CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — "paradigm" CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

- Here is a — perhaps *the* — "paradigm" CR [36, 39, 32, 23].
  - **The Consistency Requirement for Belief**. Agents should have *sets* of beliefs that are *logically consistent*.
- The Consistency Requirement is implied by The Alethic Ideal (*i.e.*, if S is Alethically Ideal, then S's beliefs are consistent).
  - **Alethic Ideal** (for belief). S should (*alethically, ideally*) believe (disbelieve) that p just in case p is true (false).
- We've already seen (**Gibbard's Coin**) that The Alethic Ideal can come into *conflict* with The Evidential Ideal.
  - **The Evidential Ideal** (for belief). S should (*evidentially, ideally*) believe (disbelieve) p if S's total evidence supports (counter-supports) p. Otherwise, S should *suspend* on p.
- More subtle cases reveal that The Consistency Requirement can also conflict with The Evidential Ideal [6, 25, 13, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [13] explanation of *the datum* is helpful.

"...if the avoidance of recognizable inconsistency were an absolute prerequisite of rational belief, we could not rationally believe each member of a set of propositions and also rationally believe of this set that at least one of its members is false. But this in turn pressures us to be unduly cautious. It pressures us to believe only those propositions that are certain or at least close to certain for us, since otherwise we are likely to have reasons to believe that at least one of these propositions is false. At first glance, the requirement that we avoid recognizable inconsistency seems little enough to ask in the name of rationality. It asks only that we avoid certain error. It turns out, however, that this is far too much to ask."

- We will offer an explication of Foley's (Old) Lockeanism. The idea: *Epistemic rationality requires minimization of expected inaccuracy*. Later, we will examine a New (Knowledge Centered) Lockeanism, based on a refinement of this idea.

"...if the avoidance of recognizable inconsistency were an absolute prerequisite of rational belief, we could not rationally believe each member of a set of propositions and also rationally believe of this set that at least one of its members is false. But this in turn pressures us to be unduly cautious. It pressures us to believe only those propositions that are certain or at least close to certain for us, since otherwise we are likely to have reasons to believe that at least one of these propositions is false. At first glance, the requirement that we avoid recognizable inconsistency seems little enough to ask in the name of rationality. It asks only that we avoid certain error. It turns out, however, that this is far too much to ask."

- We will offer an explication of Foley's (Old) Lockeanism. The idea: *Epistemic rationality requires minimization of expected inaccuracy*. Later, we will examine a New (Knowledge Centered) Lockeanism, based on a refinement of this idea.

- We assume that our agent has a credence function $b(\cdot)$, which is *probabilistic*. This allows us to use $b(\cdot)$ to define notions of (subjective) *expected* (epistemic) utility.
- We assume that our agent takes exactly one of three qualitative attitudes (B, D, S) toward each member of a finite agenda $\mathcal{A}$ of (classical, possible worlds) propositions.
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *rational coherence constraint* for qualitative judgment sets $\mathbf{B}$ (on $\mathcal{A}$).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [18, 29, 31] $u(\cdot)$.
  - ☞ Following Easwaran [11] & Dorst [9], we assume our agent cares *only* about whether their judgments are *accurate*.
- Specifically, our agent attaches some *positive* utility (r) with making an *accurate* judgment, and some *negative* utility ($-w$) with making an *inaccurate* judgment (where $w > r > 0$).

- We assume that our agent has a credence function $b(\cdot)$, which is *probabilistic*. This allows us to use $b(\cdot)$ to define notions of (subjective) *expected* (epistemic) utility.
- We assume that our agent takes exactly one of three qualitative attitudes (B, D, S) toward each member of a finite agenda $\mathcal{A}$ of (classical, possible worlds) propositions.
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *rational coherence constraint* for qualitative judgment sets $\mathbf{B}$ (on $\mathcal{A}$).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [18, 29, 31] $u(\cdot)$.
  - ☞ Following Easwaran [11] & Dorst [9], we assume our agent cares *only* about whether their judgments are *accurate*.
- Specifically, our agent attaches some *positive* utility (r) with making an *accurate* judgment, and some *negative* utility ($-w$) with making an *inaccurate* judgment (where $w > r > 0$).

- We assume that our agent has a credence function $b(\cdot)$, which is *probabilistic*. This allows us to use $b(\cdot)$ to define notions of (subjective) *expected* (epistemic) utility.
- We assume that our agent takes exactly one of three qualitative attitudes (B, D, S) toward each member of a finite agenda $\mathcal{A}$ of (classical, possible worlds) propositions.
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *rational coherence constraint* for qualitative judgment sets **B** (on $\mathcal{A}$).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [18, 29, 31] $u(\cdot)$.
  - ☞ Following Easwaran [11] & Dorst [9], we assume our agent cares *only* about whether their judgments are *accurate*.
- Specifically, our agent attaches some *positive* utility ($\mathfrak{r}$) with making an *accurate* judgment, and some *negative* utility ($-\mathfrak{w}$) with making an *inaccurate* judgment (where $\mathfrak{w} > \mathfrak{r} > 0$).

- We assume that our agent has a credence function $b(\cdot)$, which is *probabilistic*. This allows us to use $b(\cdot)$ to define notions of (subjective) *expected* (epistemic) utility.
- We assume that our agent takes exactly one of three qualitative attitudes (B, D, S) toward each member of a finite agenda $\mathcal{A}$ of (classical, possible worlds) propositions.
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *rational coherence constraint* for qualitative judgment sets $\mathbf{B}$ (on $\mathcal{A}$).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [18, 29, 31] $u(\cdot)$.
  - ☞ Following Easwaran [11] & Dorst [9], we assume our agent cares *only* about whether their judgments are *accurate*.
- Specifically, our agent attaches some *positive* utility ($\mathfrak{r}$) with making an *accurate* judgment, and some *negative* utility ($-\mathfrak{w}$) with making an *inaccurate* judgment (where $\mathfrak{w} > \mathfrak{r} > 0$).

- We assume that our agent has a credence function $b(\cdot)$, which is *probabilistic*. This allows us to use $b(\cdot)$ to define notions of (subjective) *expected* (epistemic) utility.
- We assume that our agent takes exactly one of three qualitative attitudes (B, D, S) toward each member of a finite agenda $\mathcal{A}$ of (classical, possible worlds) propositions.
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *rational coherence constraint* for qualitative judgment sets $\mathbf{B}$ (on $\mathcal{A}$).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [18, 29, 31] $u(\cdot)$.
  - ☞ Following Easwaran [11] & Dorst [9], we assume our agent cares *only* about whether their judgments are *accurate*.
- Specifically, our agent attaches some *positive* utility ($\mathfrak{r}$) with making an *accurate* judgment, and some *negative* utility ($-\mathfrak{w}$) with making an *inaccurate* judgment (where $\mathfrak{w} > \mathfrak{r} > 0$).

- We assume that our agent has a credence function $b(\cdot)$, which is *probabilistic*. This allows us to use $b(\cdot)$ to define notions of (subjective) *expected* (epistemic) utility.
- We assume that our agent takes exactly one of three qualitative attitudes (B, D, S) toward each member of a finite agenda $\mathcal{A}$ of (classical, possible worlds) propositions.
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *rational coherence constraint* for qualitative judgment sets $\mathbf{B}$ (on $\mathcal{A}$).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [18, 29, 31] $u(\cdot)$.
  - ☞ Following Easwaran [11] & Dorst [9], we assume our agent cares *only* about whether their judgments are *accurate*.
- Specifically, our agent attaches some *positive* utility ($\mathfrak{r}$) with making an *accurate* judgment, and some *negative* utility ($-\mathfrak{w}$) with making an *inaccurate* judgment (where $\mathfrak{w} > \mathfrak{r} > 0$).

- Because suspensions are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to suspensions $S(p)$, independently of the truth-value of p.
- Thus, we have the following piecewise definition of $u(\cdot, w)$.

$$u(B(p), w) \stackrel{\text{def}}{=} \begin{cases} -\mathfrak{w} & \text{if } p \text{ is false at } w \\ \mathfrak{r} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(D(p), w) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{r} & \text{if } p \text{ is false at } w \\ -\mathfrak{w} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(S(p), w) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$

- With this *accuracy-centered* epistemic utility function in hand, we can derive a naïve EUT coherence requirement.

- Because suspensions are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to suspensions $S(p)$, independently of the truth-value of p.
- Thus, we have the following piecewise definition of $u(\cdot, w)$.

$$u(B(p), w) \stackrel{\text{def}}{=} \begin{cases} -\mathfrak{w} & \text{if } p \text{ is false at } w \\ \mathfrak{r} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(D(p), w) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{r} & \text{if } p \text{ is false at } w \\ -\mathfrak{w} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(S(p), w) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$

- With this *accuracy-centered* epistemic utility function in hand, we can derive a naïve EUT coherence requirement.

- Because suspensions are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to suspensions $S(p)$, independently of the truth-value of p.
- Thus, we have the following piecewise definition of $u(\cdot, w)$.

$$u(B(p), w) \stackrel{\text{def}}{=} \begin{cases} -\mathfrak{w} & \text{if } p \text{ is false at } w \\ \mathfrak{r} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(D(p), w) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{r} & \text{if } p \text{ is false at } w \\ -\mathfrak{w} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(S(p), w) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$

- With this *accuracy-centered* epistemic utility function in hand, we can derive a naïve EUT coherence requirement.

- Because suspensions are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to suspensions $S(p)$, independently of the truth-value of p.
- Thus, we have the following piecewise definition of $u(\cdot, w)$.

$$u(B(p), w) \stackrel{\text{def}}{=} \begin{cases} -\mathfrak{w} & \text{if } p \text{ is false at } w \\ \mathfrak{r} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(D(p), w) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{r} & \text{if } p \text{ is false at } w \\ -\mathfrak{w} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(S(p), w) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$

- With this *accuracy-centered* epistemic utility function in hand, we can derive a naïve EUT coherence requirement.

- Because suspensions are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to suspensions $S(p)$, independently of the truth-value of p.
- Thus, we have the following piecewise definition of $u(\cdot, w)$.

$$u(B(p), w) \overset{\text{def}}{=} \begin{cases} -\mathfrak{w} & \text{if } p \text{ is false at } w \\ \mathfrak{r} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(D(p), w) \overset{\text{def}}{=} \begin{cases} \mathfrak{r} & \text{if } p \text{ is false at } w \\ -\mathfrak{w} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(S(p), w) \overset{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$

- With this *accuracy-centered* epistemic utility function in hand, we can derive a naïve EUT coherence requirement.

- Because suspensions are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to suspensions $S(p)$, independently of the truth-value of p.
- Thus, we have the following piecewise definition of $u(\cdot, w)$.

$$u(B(p), w) \stackrel{\text{def}}{=} \begin{cases} -\mathfrak{w} & \text{if } p \text{ is false at } w \\ \mathfrak{r} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(D(p), w) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{r} & \text{if } p \text{ is false at } w \\ -\mathfrak{w} & \text{if } p \text{ is true at } w \end{cases}$$

$$u(S(p), w) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$

- With this *accuracy-centered* epistemic utility function in hand, we can derive a naïve EUT coherence requirement.

- To do so, we'll also need a *decision-theoretic principle*.
- As we saw, applications of EUT to grounding probabilism as a (synchronic) requirement for $b(\cdot)$ typically appeal to a *non-dominance* (in epistemic utility) principle [20, 37, 35].
- But, some authors apply an *expected epistemic utility maximization* (or *expected inaccuracy minimization*) principle to derive rational requirements [28, 16, 12, 33].

> **Coherence**. An agent's belief set **B** over an agenda $\mathcal{A}$ should, from the point of view of their own credence function $b(\cdot)$, *maximize expected epistemic utility* (or *minimize expected inaccuracy*). That is, **B** should maximize
>
> $$EEU(\mathbf{B}, b) \stackrel{\text{def}}{=} \sum_{p \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w)$$
>
> where $\mathbf{B}(p)$ is the agent's attitude toward p, and $W \stackrel{\text{def}}{=} \bigcup \mathcal{A}$.

- We also assume "*act-state independence*": $\mathbf{B}(p)$ and p are *b-independent* [15, 5, 4, 27]. See Extras for discussion.

- To do so, we'll also need a *decision-theoretic principle*.
- As we saw, applications of EUT to grounding probabilism as a (synchronic) requirement for $b(\cdot)$ typically appeal to a *non-dominance* (in epistemic utility) principle [20, 37, 35].
- But, some authors apply an *expected epistemic utility maximization* (or *expected inaccuracy minimization*) principle to derive rational requirements [28, 16, 12, 33].

> **Coherence**. An agent's belief set $\mathbf{B}$ over an agenda $\mathcal{A}$ should, from the point of view of their own credence function $b(\cdot)$, *maximize expected epistemic utility* (or *minimize expected inaccuracy*). That is, $\mathbf{B}$ should maximize
>
> $$EEU(\mathbf{B}, b) \stackrel{\text{def}}{=} \sum_{p \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w)$$
>
> where $\mathbf{B}(p)$ is the agent's attitude toward p, and $W \stackrel{\text{def}}{=} \bigcup \mathcal{A}$.

- We also assume "*act-state independence*": $\mathbf{B}(p)$ and p are *b-independent* [15, 5, 4, 27]. See Extras for discussion.

- To do so, we'll also need a *decision-theoretic principle*.
- As we saw, applications of EUT to grounding probabilism as a (synchronic) requirement for $b(\cdot)$ typically appeal to a *non-dominance* (in epistemic utility) principle [20, 37, 35].
- But, some authors apply an *expected epistemic utility maximization* (or *expected inaccuracy minimization*) principle to derive rational requirements [28, 16, 12, 33].

> **Coherence**. An agent's belief set $\mathbf{B}$ over an agenda $\mathcal{A}$ should, from the point of view of their own credence function $b(\cdot)$, *maximize expected epistemic utility* (or *minimize expected inaccuracy*). That is, $\mathbf{B}$ should maximize
>
> $$EEU(\mathbf{B}, b) \stackrel{\text{def}}{=} \sum_{p \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w)$$
>
> where $\mathbf{B}(p)$ is the agent's attitude toward p, and $W \stackrel{\text{def}}{=} \bigcup \mathcal{A}$.

- We also assume "*act-state independence*": $\mathbf{B}(p)$ and p are *b-independent* [15, 5, 4, 27]. See Extras for discussion.

- To do so, we'll also need a *decision-theoretic principle*.
- As we saw, applications of EUT to grounding probabilism as a (synchronic) requirement for $b(\cdot)$ typically appeal to a *non-dominance* (in epistemic utility) principle [20, 37, 35].
- But, some authors apply an *expected epistemic utility maximization* (or *expected inaccuracy minimization*) principle to derive rational requirements [28, 16, 12, 33].

> **Coherence**. An agent's belief set $\mathbf{B}$ over an agenda $\mathcal{A}$ should, from the point of view of their own credence function $b(\cdot)$, *maximize expected epistemic utility* (or *minimize expected inaccuracy*). That is, $\mathbf{B}$ should maximize
>
> $$EEU(\mathbf{B}, b) \stackrel{\text{def}}{=} \sum_{p \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w)$$
>
> where $\mathbf{B}(p)$ is the agent's attitude toward p, and $W \stackrel{\text{def}}{=} \bigcup \mathcal{A}$.

- We also assume "*act-state independence*": $\mathbf{B}(p)$ and p are *b-independent* [15, 5, 4, 27]. See Extras for discussion.

- To do so, we'll also need a *decision-theoretic principle*.
- As we saw, applications of EUT to grounding probabilism as a (synchronic) requirement for $b(\cdot)$ typically appeal to a *non-dominance* (in epistemic utility) principle [20, 37, 35].
- But, some authors apply an *expected epistemic utility maximization* (or *expected inaccuracy minimization*) principle to derive rational requirements [28, 16, 12, 33].

Coherence. An agent's belief set $\mathbf{B}$ over an agenda $\mathcal{A}$ should, from the point of view of their own credence function $b(\cdot)$, *maximize expected epistemic utility* (or *minimize expected inaccuracy*). That is, $\mathbf{B}$ should maximize

$$EEU(\mathbf{B}, b) \stackrel{\text{def}}{=} \sum_{p \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w)$$

where $\mathbf{B}(p)$ is the agent's attitude toward p, and $W \stackrel{\text{def}}{=} \bigcup \mathcal{A}$.

- We also assume "*act-state independence*": $\mathbf{B}(p)$ and p are *b-independent* [15, 5, 4, 27]. See Extras for discussion.

- To do so, we'll also need a *decision-theoretic principle*.
- As we saw, applications of EUT to grounding probabilism as a (synchronic) requirement for $b(\cdot)$ typically appeal to a *non-dominance* (in epistemic utility) principle [20, 37, 35].
- But, some authors apply an *expected epistemic utility maximization* (or *expected inaccuracy minimization*) principle to derive rational requirements [28, 16, 12, 33].

> **Coherence**. An agent's belief set $\mathbf{B}$ over an agenda $\mathcal{A}$ should, from the point of view of their own credence function $b(\cdot)$, *maximize expected epistemic utility* (or *minimize expected inaccuracy*). That is, $\mathbf{B}$ should maximize
>
> $$EEU(\mathbf{B}, b) \overset{\text{def}}{=} \sum_{p \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(p), w)$$
>
> where $\mathbf{B}(p)$ is the agent's attitude toward p, and $W \overset{\text{def}}{=} \bigcup \mathcal{A}$.

- We also assume "*act-state independence*": $\mathbf{B}(p)$ and p are *b-independent* [15, 5, 4, 27]. See Extras for discussion.

- The consequences of **Coherence** are rather simple and intuitive. It is straightforward to prove the following result.

Theorem ([11, 9]). An agent with credence function $b(\cdot)$ and qualitative judgment set $\mathbf{B}$ over agenda $\mathcal{A}$ satisfies **Coherence** *if and only if* for all $p \in \mathcal{A}$

$$B(p) \in \mathbf{B} \text{ iff } b(p) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$

$$D(p) \in \mathbf{B} \text{ iff } b(p) < \frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}},$$

$$S(p) \in \mathbf{B} \text{ iff } b(p) \in \left[\frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- In other words, **Coherence** *entails Lockean representability*, where the Lockean thresholds are determined by the way the agent (relatively) values accuracy *vs.* inaccuracy.
- This provides an elegant, EUT-based explanation of why Lockean representability is a rational requirement for agents with *both* credences *and* qualitative attitudes.

☞ As Dorst [9] puts it: *Lockeans maximize expected accuracy*.

- The consequences of **Coherence** are rather simple and intuitive. It is straightforward to prove the following result.

> **Theorem** ([11, 9]). An agent with credence function $b(\cdot)$ and qualitative judgment set $\mathbf{B}$ over agenda $\mathcal{A}$ satisfies **Coherence** *if and only if* for all $p \in \mathcal{A}$

$$B(p) \in \mathbf{B} \text{ iff } b(p) > \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$

$$D(p) \in \mathbf{B} \text{ iff } b(p) < \tfrac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}},$$

$$S(p) \in \mathbf{B} \text{ iff } b(p) \in \left[\tfrac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- In other words, **Coherence** *entails Lockean representability*, where the Lockean thresholds are determined by the way the agent (relatively) values accuracy *vs.* inaccuracy.
- This provides an elegant, EUT-based explanation of why Lockean representability is a rational requirement for agents with *both* credences *and* qualitative attitudes.

☞ As Dorst [9] puts it: *Lockeans maximize expected accuracy*.

- The consequences of **Coherence** are rather simple and intuitive. It is straightforward to prove the following result.

> **Theorem** ([11, 9]). An agent with credence function $b(\cdot)$ and qualitative judgment set **B** over agenda $\mathcal{A}$ satisfies **Coherence** *if and only if* for all $p \in \mathcal{A}$
>
> $$B(p) \in \mathbf{B} \text{ iff } b(p) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$
> $$D(p) \in \mathbf{B} \text{ iff } b(p) < \frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}},$$
> $$S(p) \in \mathbf{B} \text{ iff } b(p) \in \left[\frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- In other words, **Coherence** *entails Lockean representability*, where the Lockean thresholds are determined by the way the agent (relatively) values accuracy *vs.* inaccuracy.
- This provides an elegant, EUT-based explanation of why Lockean representability is a rational requirement for agents with *both* credences *and* qualitative attitudes.

☞ As Dorst [9] puts it: *Lockeans maximize expected accuracy*.

- The consequences of **Coherence** are rather simple and intuitive. It is straightforward to prove the following result.

> **Theorem** ([11, 9]). An agent with credence function $b(\cdot)$ and qualitative judgment set $\mathbf{B}$ over agenda $\mathcal{A}$ satisfies **Coherence** *if and only if* for all $p \in \mathcal{A}$
>
> $$B(p) \in \mathbf{B} \text{ iff } b(p) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$
>
> $$D(p) \in \mathbf{B} \text{ iff } b(p) < \frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}},$$
>
> $$S(p) \in \mathbf{B} \text{ iff } b(p) \in \left[\frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- In other words, **Coherence** *entails Lockean representability*, where the Lockean thresholds are determined by the way the agent (relatively) values accuracy *vs.* inaccuracy.
- This provides an elegant, EUT-based explanation of why Lockean representability is a rational requirement for agents with *both* credences *and* qualitative attitudes.

☞ As Dorst [9] puts it: *Lockeans maximize expected accuracy*.

- The consequences of **Coherence** are rather simple and intuitive. It is straightforward to prove the following result.

> **Theorem** ([11, 9]). An agent with credence function $b(\cdot)$ and qualitative judgment set $\mathbf{B}$ over agenda $\mathcal{A}$ satisfies **Coherence** *if and only if* for all $p \in \mathcal{A}$
>
> $$B(p) \in \mathbf{B} \text{ iff } b(p) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$
>
> $$D(p) \in \mathbf{B} \text{ iff } b(p) < \frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}},$$
>
> $$S(p) \in \mathbf{B} \text{ iff } b(p) \in \left[\frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- In other words, **Coherence** *entails Lockean representability*, where the Lockean thresholds are determined by the way the agent (relatively) values accuracy *vs.* inaccuracy.
- This provides an elegant, EUT-based explanation of why Lockean representability is a rational requirement for agents with *both* credences *and* qualitative attitudes.

☞ As Dorst [9] puts it: *Lockeans maximize expected accuracy*.

- The consequences of **Coherence** are rather simple and intuitive. It is straightforward to prove the following result.

> **Theorem** ([11, 9]). An agent with credence function $b(\cdot)$ and qualitative judgment set **B** over agenda $\mathcal{A}$ satisfies **Coherence** *if and only if* for all $p \in \mathcal{A}$
>
> $$B(p) \in \mathbf{B} \text{ iff } b(p) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$
> $$D(p) \in \mathbf{B} \text{ iff } b(p) < \frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}},$$
> $$S(p) \in \mathbf{B} \text{ iff } b(p) \in \left[\frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- In other words, **Coherence** *entails Lockean representability*, where the Lockean thresholds are determined by the way the agent (relatively) values accuracy *vs.* inaccuracy.
- This provides an elegant, EUT-based explanation of why Lockean representability is a rational requirement for agents with *both* credences *and* qualitative attitudes.

☞ As Dorst [9] puts it: *Lockeans maximize expected accuracy*.

- In the *Meno* (97e–98a), Socrates says:

> *For true opinions, as long as they remain, are a fine thing and all they do is good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down ... That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down...*

- Our epistemic utility function (for belief) only assigned positive value to *correctness*. What about *knowledge*?
- Nothing in our (teleological) framework for epistemic utility theory rules out attaching (additional) value to *knowledge*, *over and above* the value we place on correctness/accuracy.
- There are various ways one might refine/alter our naïve (accuracy centered) epistemic utility function, so as to take account of this Meno-style *value of knowledge*.
- I will now describe some models (help needed here!).

- In the *Meno* (97e–98a), Socrates says:

> *For true opinions, as long as they remain, are a fine thing and all they do is good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down ... That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down...*

- Our epistemic utility function (for belief) only assigned positive value to *correctness*. What about *knowledge*?
- Nothing in our (teleological) framework for epistemic utility theory rules out attaching (additional) value to *knowledge*, *over and above* the value we place on correctness/accuracy.
- There are various ways one might refine/alter our naïve (accuracy centered) epistemic utility function, so as to take account of this Meno-style *value of knowledge*.
- I will now describe some models (help needed here!).

- In the *Meno* (97e–98a), Socrates says:

> *For true opinions, as long as they remain, are a fine thing and all they do is good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down ... That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down...*

- Our epistemic utility function (for belief) only assigned positive value to *correctness*. What about *knowledge*?
- Nothing in our (teleological) framework for epistemic utility theory rules out attaching (additional) value to *knowledge*, *over and above* the value we place on correctness/accuracy.
- There are various ways one might refine/alter our naïve (accuracy centered) epistemic utility function, so as to take account of this Meno-style *value of knowledge*.
- I will now describe some models (help needed here!).

- In the *Meno* (97e–98a), Socrates says:

> *For true opinions, as long as they remain, are a fine thing and all they do is good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down . . . That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down. . .*

- Our epistemic utility function (for belief) only assigned positive value to *correctness*. What about *knowledge*?
- Nothing in our (teleological) framework for epistemic utility theory rules out attaching (additional) value to *knowledge*, *over and above* the value we place on correctness/accuracy.
- There are various ways one might refine/alter our naïve (accuracy centered) epistemic utility function, so as to take account of this Meno-style *value of knowledge*.
- I will now describe some models (help needed here!).

- In the *Meno* (97e–98a), Socrates says:

> *For true opinions, as long as they remain, are a fine thing and all they do is good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down ... That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down...*

- Our epistemic utility function (for belief) only assigned positive value to *correctness*. What about *knowledge*?
- Nothing in our (teleological) framework for epistemic utility theory rules out attaching (additional) value to *knowledge*, *over and above* the value we place on correctness/accuracy.
- There are various ways one might refine/alter our naïve (accuracy centered) epistemic utility function, so as to take account of this Meno-style *value of knowledge*.
- I will now describe some models (help needed here!).

- In the *Meno* (97e–98a), Socrates says:

> *For true opinions, as long as they remain, are a fine thing and all they do is good, but they are not willing to remain long, and they escape from a man's mind, so that they are not worth much until one ties them down . . . That is why knowledge is prized higher than correct opinion, and knowledge differs from correct opinion in being tied down. . .*

- Our epistemic utility function (for belief) only assigned positive value to *correctness*. What about *knowledge*?
- Nothing in our (teleological) framework for epistemic utility theory rules out attaching (additional) value to *knowledge*, *over and above* the value we place on correctness/accuracy.
- There are various ways one might refine/alter our naïve (accuracy centered) epistemic utility function, so as to take account of this Meno-style *value of knowledge*.
- I will now describe some models (help needed here!).

- Here's the most general class of models we've come up with:

world (w)	$b(w)$	$u(B(p), w)$
$K(p)$	a	$\mathfrak{x}$
$p \,\&\, \neg K(p)$	b	$\mathfrak{y}$
$\neg p \,\&\, \neg K(\neg p)$	c	$\mathfrak{z}$
$K(\neg p)$	$1 - (\mathsf{a} + \mathsf{b} + \mathsf{c})$	$\mathfrak{u}$

- When we represent things at this level of generality, we realize there are (at least) two key *choice points* here.

  (1) Are knowledge and truth *both* positively valuable (or is knowledge *the only* state that has positive value)? That is: should we have *both* $\mathfrak{x} > 0$ *and* $\mathfrak{y} > 0$, or *just* $\mathfrak{x} > 0$?

  (2) Should truth be more valuable than falsehood, even *within the state of ignorance*? That is, should we have $\mathfrak{y} > \mathfrak{z}$?

- These choices — especially (1) — will impact the kinds of "Lockean Theses" that fall out of the models (*via* MEEU).

- Here's the most general class of models we've come up with:

world (w)	$b(w)$	$u(B(p), w)$
$K(p)$	a	$\mathfrak{x}$
$p \,\&\, \neg K(p)$	b	$\mathfrak{y}$
$\neg p \,\&\, \neg K(\neg p)$	c	$\mathfrak{z}$
$K(\neg p)$	1 − (a + b + c)	$\mathfrak{u}$

- When we represent things at this level of generality, we realize there are (at least) two key *choice points* here.

  (1) Are knowledge and truth *both* positively valuable (or is knowledge *the only* state that has positive value)? That is: should we have *both* $\mathfrak{x} > 0$ *and* $\mathfrak{y} > 0$, or *just* $\mathfrak{x} > 0$?

  (2) Should truth be more valuable than falsehood, even *within the state of ignorance*? That is, should we have $\mathfrak{y} > \mathfrak{z}$?

- These choices — especially (1) — will impact the kinds of "Lockean Theses" that fall out of the models (*via* MEEU).

- Here's the most general class of models we've come up with:

world (w)	$b(w)$	$u(B(p), w)$
$K(p)$	a	$\mathfrak{x}$
$p \,\&\, \neg K(p)$	b	$\mathfrak{y}$
$\neg p \,\&\, \neg K(\neg p)$	c	$\mathfrak{z}$
$K(\neg p)$	$1 - (\mathsf{a} + \mathsf{b} + \mathsf{c})$	$\mathfrak{u}$

- When we represent things at this level of generality, we realize there are (at least) two key *choice points* here.

  (1) Are knowledge and truth *both* positively valuable (or is knowledge *the only* state that has positive value)? That is: should we have *both* $\mathfrak{x} > 0$ *and* $\mathfrak{y} > 0$, or *just* $\mathfrak{x} > 0$?

  (2) Should truth be more valuable than falsehood, even *within the state of ignorance*? That is, should we have $\mathfrak{y} > \mathfrak{z}$?

- These choices — especially (1) — will impact the kinds of "Lockean Theses" that fall out of the models (*via* MEEU).

- Here's the most general class of models we've come up with:

world (w)	$b(w)$	$u(B(p), w)$
$K(p)$	a	$\mathfrak{x}$
$p \,\&\, \neg K(p)$	b	$\mathfrak{y}$
$\neg p \,\&\, \neg K(\neg p)$	c	$\mathfrak{z}$
$K(\neg p)$	$1 - (\mathsf{a} + \mathsf{b} + \mathsf{c})$	$\mathfrak{u}$

- When we represent things at this level of generality, we realize there are (at least) two key *choice points* here.

  (1) Are knowledge and truth *both* positively valuable (or is knowledge *the only* state that has positive value)? That is: should we have *both* $\mathfrak{x} > 0$ *and* $\mathfrak{y} > 0$, or *just* $\mathfrak{x} > 0$?

  (2) Should truth be more valuable than falsehood, even *within the state of ignorance*? That is, should we have $\mathfrak{y} > \mathfrak{z}$?

- These choices — especially (1) — will impact the kinds of "Lockean Theses" that fall out of the models (*via* MEEU).

- Here's the most general class of models we've come up with:

world (w)	$b(w)$	$u(B(p), w)$
$K(p)$	a	$\mathfrak{x}$
$p \,\&\, \neg K(p)$	b	$\mathfrak{y}$
$\neg p \,\&\, \neg K(\neg p)$	c	$\mathfrak{z}$
$K(\neg p)$	$1 - (\mathsf{a} + \mathsf{b} + \mathsf{c})$	$\mathfrak{u}$

- When we represent things at this level of generality, we realize there are (at least) two key *choice points* here.

  (1) Are knowledge and truth *both* positively valuable (or is knowledge *the only* state that has positive value)? That is: should we have *both* $\mathfrak{x} > 0$ *and* $\mathfrak{y} > 0$, or *just* $\mathfrak{x} > 0$?

  (2) Should truth be more valuable than falsehood, even *within the state of ignorance*? That is, should we have $\mathfrak{y} > \mathfrak{z}$?

- These choices — especially (1) — will impact the kinds of "Lockean Theses" that fall out of the models (*via* MEEU).

- Here's the most general class of models we've come up with:

world (w)	$b(w)$	$u(B(p), w)$
$K(p)$	a	$\mathfrak{x}$
$p \mathbin{\&} \neg K(p)$	b	$\mathfrak{y}$
$\neg p \mathbin{\&} \neg K(\neg p)$	c	$\mathfrak{z}$
$K(\neg p)$	$1 - (\mathsf{a} + \mathsf{b} + \mathsf{c})$	$\mathfrak{u}$

- When we represent things at this level of generality, we realize there are (at least) two key *choice points* here.
  (1) Are knowledge and truth *both* positively valuable (or is knowledge *the only* state that has positive value)? That is: should we have *both* $\mathfrak{x} > 0$ *and* $\mathfrak{y} > 0$, or *just* $\mathfrak{x} > 0$?
  (2) Should truth be more valuable than falsehood, even *within the state of ignorance*? That is, should we have $\mathfrak{y} > \mathfrak{z}$?
- These choices — especially (1) — will impact the kinds of "Lockean Theses" that fall out of the models (*via* MEEU).

- The simplest models would be ones in which *only* knowledge has positive value (*i.e.*, greater than suspension).
- Such models — which answer (1) in the negative — will all yield constraints of the following (general) form:

  K-Coherence. $B(p)$ is rationally permissible just in case $b(K(p))$ is "sufficiently high," where "sufficiently high" may be rather complicated (and it may depend on the other credences the agent assigns), but it will always have to be *greater than* $1/2$, provided only that the penalties for non-knowledge are greater than the reward for knowledge.

- Models which answer (1) in the *affirmative*, are far more complex, and can be compatible with $b(K(p))$ being *arbitrarily low*. We don't have a full characterization of those models, but we have some special cases worked out.
- Let's focus on the simplest models (and Lotteries). First, a review of *accuracy*-centered models (and Lotteries).

- The simplest models would be ones in which *only* knowledge has positive value (*i.e.*, greater than suspension).
- Such models — which answer (1) in the negative — will all yield constraints of the following (general) form:

> **K-Coherence.** $B(p)$ is rationally permissible just in case $b(K(p))$ is "sufficiently high," where "sufficiently high" may be rather complicated (and it may depend on the other credences the agent assigns), but it will always have to be *greater than* $1/2$, provided only that the penalties for non-knowledge are greater than the reward for knowledge.

- Models which answer (1) in the *affirmative*, are far more complex, and can be compatible with $b(K(p))$ being *arbitrarily low*. We don't have a full characterization of those models, but we have some special cases worked out.
- Let's focus on the simplest models (and Lotteries). First, a review of *accuracy*-centered models (and Lotteries).

- The simplest models would be ones in which *only* knowledge has positive value (*i.e.*, greater than suspension).
- Such models — which answer (1) in the negative — will all yield constraints of the following (general) form:

  K-Coherence. $B(p)$ is rationally permissible just in case $b(K(p))$ is "sufficiently high," where "sufficiently high" may be rather complicated (and it may depend on the other credences the agent assigns), but it will always have to be *greater than* 1/2, provided only that the penalties for non-knowledge are greater than the reward for knowledge.

- Models which answer (1) in the *affirmative*, are far more complex, and can be compatible with $b(K(p))$ being *arbitrarily low*. We don't have a full characterization of those models, but we have some special cases worked out.
- Let's focus on the simplest models (and Lotteries). First, a review of *accuracy*-centered models (and Lotteries).

- The simplest models would be ones in which *only* knowledge has positive value (*i.e.*, greater than suspension).
- Such models — which answer (1) in the negative — will all yield constraints of the following (general) form:

> **K-Coherence**. $B(p)$ is rationally permissible just in case $b(K(p))$ is "sufficiently high," where "sufficiently high" may be rather complicated (and it may depend on the other credences the agent assigns), but it will always have to be *greater than* $1/2$, provided only that the penalties for non-knowledge are greater than the reward for knowledge.

- Models which answer (1) in the *affirmative*, are far more complex, and can be compatible with $b(K(p))$ being *arbitrarily low*. We don't have a full characterization of those models, but we have some special cases worked out.
- Let's focus on the simplest models (and Lotteries). First, a review of *accuracy*-centered models (and Lotteries).

- The simplest models would be ones in which *only* knowledge has positive value (*i.e.*, greater than suspension).
- Such models — which answer (1) in the negative — will all yield constraints of the following (general) form:

> **K-Coherence**. $B(p)$ is rationally permissible just in case $b(K(p))$ is "sufficiently high," where "sufficiently high" may be rather complicated (and it may depend on the other credences the agent assigns), but it will always have to be *greater than* $1/2$, provided only that the penalties for non-knowledge are greater than the reward for knowledge.

- Models which answer (1) in the *affirmative*, are far more complex, and can be compatible with $b(K(p))$ being *arbitrarily low*. We don't have a full characterization of those models, but we have some special cases worked out.
- Let's focus on the simplest models (and Lotteries). First, a review of *accuracy*-centered models (and Lotteries).

- Suppose our (naïve) agent has a belief set $\mathbf{B}_n$ on a *minimal inconsistent* agenda of size n (*e.g.*, $(n-1)$-ticket lottery).

 Theorem ([10]). For all $n \geq 2$ and any probability function $\Pr(\cdot)$, the $\Pr(\cdot)$-Lockean-representability of $\mathbf{B}_n$ (with threshold t) *entails* deductive consistency of $\mathbf{B}_n$ *iff* $t \geq \frac{n-1}{n}$.

- If we combine this with the Easwaran/Dorst **Coherence** theorem, we get the following conditions under which the **Coherence** of $\mathbf{B}_n$ *entails* the *consistency* of $\mathbf{B}_n$.

 Theorem. For all $n \geq 2$, an agent with an accuracy-centered utility function u, a credence function $b(\cdot)$, and a belief set $\mathbf{B}_n$, the **Coherence** of $\mathbf{B}_n$ entails the consistency of $\mathbf{B}_n$ *iff*

 (†) $$\mathrm{w} \geq (n-1) \cdot \mathrm{r}.$$

☞ Insisting that **Coherence** *implies consistency* (wrt $\mathbf{B}_n$) *requires (naïve) agents to disvalue inaccuracy at least* $(n-1)$ *times as much as they value accuracy.*

- Suppose our (naïve) agent has a belief set $\mathbf{B}_n$ on a *minimal inconsistent* agenda of size n (*e.g.*, $(n-1)$–ticket lottery).

  Theorem ([10]). For all $n \geq 2$ and any probability function $\Pr(\cdot)$, the $\Pr(\cdot)$-Lockean-representability of $\mathbf{B}_n$ (with threshold t) *entails* deductive consistency of $\mathbf{B}_n$ *iff* $t \geq \frac{n-1}{n}$.

- If we combine this with the Easwaran/Dorst **Coherence** theorem, we get the following conditions under which the **Coherence** of $\mathbf{B}_n$ *entails* the *consistency* of $\mathbf{B}_n$.

  Theorem. For all $n \geq 2$, an agent with an accuracy-centered utility function u, a credence function $b(\cdot)$, and a belief set $\mathbf{B}_n$, the **Coherence** of $\mathbf{B}_n$ entails the consistency of $\mathbf{B}_n$ *iff*

  (†) $$\mathfrak{w} \geq (n-1) \cdot \mathfrak{r}.$$

☞ Insisting that **Coherence** *implies consistency* (wrt $\mathbf{B}_n$) *requires (naïve) agents to disvalue inaccuracy at least* $(n-1)$ *times as much as they value accuracy.*

- Suppose our (naïve) agent has a belief set $\mathbf{B}_n$ on a *minimal inconsistent* agenda of size n (*e.g.*, $(n-1)$–ticket lottery).

  Theorem ([10]). For all $n \geq 2$ and any probability function $\Pr(\cdot)$, the $\Pr(\cdot)$-Lockean-representability of $\mathbf{B}_n$ (with threshold t) *entails* deductive consistency of $\mathbf{B}_n$ *iff* $t \geq \frac{n-1}{n}$.

- If we combine this with the Easwaran/Dorst **Coherence** theorem, we get the following conditions under which the **Coherence** of $\mathbf{B}_n$ *entails* the *consistency* of $\mathbf{B}_n$.

  Theorem. For all $n \geq 2$, an agent with an accuracy-centered utility function u, a credence function $b(\cdot)$, and a belief set $\mathbf{B}_n$, the **Coherence** of $\mathbf{B}_n$ entails the consistency of $\mathbf{B}_n$ *iff*

  $$(\dagger) \qquad \mathfrak{w} \geq (n-1) \cdot \mathfrak{r}.$$

☞ Insisting that **Coherence** *implies consistency* (wrt $\mathbf{B}_n$) *requires (naïve) agents to disvalue inaccuracy at least* $(n-1)$ *times as much as they value accuracy.*

- Suppose our (naïve) agent has a belief set $\mathbf{B}_n$ on a *minimal inconsistent* agenda of size n (*e.g.*, $(n-1)$–ticket lottery).

  Theorem ([10]). For all $n \geq 2$ and any probability function $\Pr(\cdot)$, the $\Pr(\cdot)$-Lockean-representability of $\mathbf{B}_n$ (with threshold t) *entails* deductive consistency of $\mathbf{B}_n$ *iff* $t \geq \frac{n-1}{n}$.

- If we combine this with the Easwaran/Dorst **Coherence** theorem, we get the following conditions under which the **Coherence** of $\mathbf{B}_n$ *entails* the *consistency* of $\mathbf{B}_n$.

  Theorem. For all $n \geq 2$, an agent with an accuracy-centered utility function u, a credence function $b(\cdot)$, and a belief set $\mathbf{B}_n$, the **Coherence** of $\mathbf{B}_n$ entails the consistency of $\mathbf{B}_n$ *iff*

  $$(\dagger) \qquad \mathfrak{w} \geq (n-1) \cdot \mathfrak{r}.$$

☞ Insisting that **Coherence** *implies consistency* (wrt $\mathbf{B}_n$) *requires (naïve) agents to disvalue inaccuracy at least* $(n-1)$ *times as much as they value accuracy.*

- Suppose our (naïve) agent has a belief set $\mathbf{B}_n$ on a *minimal inconsistent* agenda of size n (*e.g.*, $(n-1)$–ticket lottery).

  Theorem ([10]). For all $n \geq 2$ and any probability function $\Pr(\cdot)$, the $\Pr(\cdot)$-Lockean-representability of $\mathbf{B}_n$ (with threshold t) *entails* deductive consistency of $\mathbf{B}_n$ *iff* $t \geq \frac{n-1}{n}$.

- If we combine this with the Easwaran/Dorst **Coherence** theorem, we get the following conditions under which the **Coherence** of $\mathbf{B}_n$ *entails* the *consistency* of $\mathbf{B}_n$.

  Theorem. For all $n \geq 2$, an agent with an accuracy-centered utility function u, a credence function $b(\cdot)$, and a belief set $\mathbf{B}_n$, the **Coherence** of $\mathbf{B}_n$ entails the consistency of $\mathbf{B}_n$ *iff*

  (†) $$\mathfrak{w} \geq (n-1) \cdot \mathfrak{r}.$$

☞ Insisting that **Coherence** *implies consistency* (wrt $\mathbf{B}_n$) *requires (naïve) agents to disvalue inaccuracy at least* $(n-1)$ *times as much as they value accuracy.*

- Suppose our (naïve) agent has a belief set $\mathbf{B}_n$ on a *minimal inconsistent* agenda of size n (*e.g.*, $(n-1)$–ticket lottery).

  Theorem ([10]). For all $n \geq 2$ and any probability function $\Pr(\cdot)$, the $\Pr(\cdot)$-Lockean-representability of $\mathbf{B}_n$ (with threshold t) *entails* deductive consistency of $\mathbf{B}_n$ *iff* $t \geq \frac{n-1}{n}$.

- If we combine this with the Easwaran/Dorst **Coherence** theorem, we get the following conditions under which the **Coherence** of $\mathbf{B}_n$ *entails* the *consistency* of $\mathbf{B}_n$.

  Theorem. For all $n \geq 2$, an agent with an accuracy-centered utility function u, a credence function $b(\cdot)$, and a belief set $\mathbf{B}_n$, the **Coherence** of $\mathbf{B}_n$ entails the consistency of $\mathbf{B}_n$ *iff*

  (†) $$\mathfrak{w} \geq (n-1) \cdot \mathfrak{r}.$$

☞ Insisting that **Coherence** *implies consistency* (wrt $\mathbf{B}_n$) *requires (naïve) agents to disvalue inaccuracy at least* $(n-1)$ *times as much as they value accuracy*.

- The simplest knowledge centered models will all require (**K-Coherence**) that an agent believe p *only if* $b(K(p)) > 1/2$.
- Thus, according to such models, the standard lottery beliefs can be irrational, and not because agents are (or ought to be) certain/near certain (or "stable") in their beliefs.
- Thus, Lottery beliefs can be irrational, because (a) *only knowledge* has positive epistemic utility, and (b) maximizing *expected* EU will force such agents to believe *only claims which they are (sufficiently) confident that they **know***.
- I, for one, am *not* confident that I know (any) lottery propositions. So, as applied *to me*, they *prohibit* me from believing that (*e.g.*) my lottery ticket will lose.
- Similar applications can be formulated for Moorean beliefs, beliefs based "solely on statistical evidence", *etc.*

- The simplest knowledge centered models will all require (**K-Coherence**) that an agent believe p *only if* $b(K(p)) > 1/2$.
- Thus, according to such models, the standard lottery beliefs can be irrational, and not because agents are (or ought to be) certain/near certain (or "stable") in their beliefs.
- Thus, Lottery beliefs can be irrational, because (a) *only knowledge* has positive epistemic utility, and (b) maximizing *expected* EU will force such agents to believe *only claims which they are (sufficiently) confident that they* ***know***.
- I, for one, am *not* confident that I know (any) lottery propositions. So, as applied *to me*, they *prohibit* me from believing that (*e.g.*) my lottery ticket will lose.
- Similar applications can be formulated for Moorean beliefs, beliefs based "solely on statistical evidence", *etc.*

- The simplest knowledge centered models will all require (**K-Coherence**) that an agent believe p *only if* $b(K(p)) > 1/2$.
- Thus, according to such models, the standard lottery beliefs can be irrational, and not because agents are (or ought to be) certain/near certain (or "stable") in their beliefs.
- Thus, Lottery beliefs can be irrational, because (a) *only knowledge* has positive epistemic utility, and (b) maximizing *expected* EU will force such agents to believe *only claims which they are (sufficiently) confident that they* ***know***.
- I, for one, am *not* confident that I know (any) lottery propositions. So, as applied *to me*, they *prohibit* me from believing that (*e.g.*) my lottery ticket will lose.
- Similar applications can be formulated for Moorean beliefs, beliefs based "solely on statistical evidence", *etc.*

- The simplest knowledge centered models will all require (**K-Coherence**) that an agent believe p *only if* $b(K(p)) > 1/2$.
- Thus, according to such models, the standard lottery beliefs can be irrational, and not because agents are (or ought to be) certain/near certain (or "stable") in their beliefs.
- Thus, Lottery beliefs can be irrational, because (a) *only knowledge* has positive epistemic utility, and (b) maximizing *expected* EU will force such agents to believe *only claims which they are (sufficiently) confident that they* ***know***.
- I, for one, am *not* confident that I know (any) lottery propositions. So, as applied *to me*, they *prohibit* me from believing that (*e.g.*) my lottery ticket will lose.
- Similar applications can be formulated for Moorean beliefs, beliefs based "solely on statistical evidence", *etc.*

- The simplest knowledge centered models will all require (**K-Coherence**) that an agent believe p *only if* $b(K(p)) > 1/2$.
- Thus, according to such models, the standard lottery beliefs can be irrational, and not because agents are (or ought to be) certain/near certain (or "stable") in their beliefs.
- Thus, Lottery beliefs can be irrational, because (a) *only knowledge* has positive epistemic utility, and (b) maximizing *expected* EU will force such agents to believe *only claims which they are (sufficiently) confident that they **know***.
- I, for one, am *not* confident that I know (any) lottery propositions. So, as applied *to me*, they *prohibit* me from believing that (*e.g.*) my lottery ticket will lose.
- Similar applications can be formulated for Moorean beliefs, beliefs based "solely on statistical evidence", *etc.*

- By way of summary, it is useful to think about the analogy between the norms we've been discussing, and principles of rational choice theory: **The Decision-Theoretic Analogy.**

Epistemic Principle	Analogous Decision-Theoretic Principle
Alethic Ideal	(AMU) Do ϕ only if ϕ maximizes utility in the *actual* world.
Consistency	(PMU) Do ϕ only if ϕ maximizes u in *some possible* world.
Coherence	(MEU) Do ϕ only if ϕ maximizes EU (relative to *some* Pr).
(WADA)	(WDOM) Do ϕ only if ϕ is *not weakly dominated* in utility.
(SADA)	(SDOM) Do ϕ only if ϕ is *not strictly dominated* in utility.

- Like the **Alethic Ideal**, (AMU) is *not* a *requirement of rationality*; and, like **Consistency**, (PMU) isn't a rational requirement either (this was the lesson of **Miners** [34, 26]).
- As Foley (*op. cit.*) explains, **Consistency** is *too demanding*. But, **Coherence** is *not* — it does *not* "pressure us to believe only those propositions that are (close to) certain for us".

- By way of summary, it is useful to think about the analogy between the norms we've been discussing, and principles of rational choice theory: **The Decision-Theoretic Analogy.**

Epistemic Principle	Analogous Decision-Theoretic Principle
Alethic Ideal	(AMU) Do ϕ only if ϕ maximizes utility in the *actual* world.
Consistency	(PMU) Do ϕ only if ϕ maximizes u in *some possible* world.
Coherence	(MEU) Do ϕ only if ϕ maximizes EU (relative to *some* Pr).
(WADA)	(WDOM) Do ϕ only if ϕ is *not weakly dominated* in utility.
(SADA)	(SDOM) Do ϕ only if ϕ is *not strictly dominated* in utility.

- Like the **Alethic Ideal**, (AMU) is *not* a *requirement of rationality*; and, like **Consistency**, (PMU) isn't a rational requirement either (this was the lesson of **Miners** [34, 26]).
- As Foley (*op. cit.*) explains, **Consistency** is *too demanding*. But, **Coherence** is *not* — it does *not* "pressure us to believe only those propositions that are (close to) certain for us".

- By way of summary, it is useful to think about the analogy between the norms we've been discussing, and principles of rational choice theory: **The Decision-Theoretic Analogy**.

Epistemic Principle	Analogous Decision-Theoretic Principle
Alethic Ideal	(AMU) Do ϕ only if ϕ maximizes utility in the *actual* world.
Consistency	(PMU) Do ϕ only if ϕ maximizes u in *some possible* world.
Coherence	(MEU) Do ϕ only if ϕ maximizes EU (relative to *some* Pr).
(WADA)	(WDOM) Do ϕ only if ϕ is *not weakly dominated* in utility.
(SADA)	(SDOM) Do ϕ only if ϕ is *not strictly dominated* in utility.

- Like the **Alethic Ideal**, (AMU) is *not* a *requirement of rationality*; and, like **Consistency**, (PMU) isn't a rational requirement either (this was the lesson of **Miners** [34, 26]).
- As Foley (*op. cit.*) explains, **Consistency** is *too demanding*. But, **Coherence** is *not* — it does *not* "pressure us to believe only those propositions that are (close to) certain for us".

- By way of summary, it is useful to think about the analogy between the norms we've been discussing, and principles of rational choice theory: **The Decision-Theoretic Analogy.**

Epistemic Principle	Analogous Decision-Theoretic Principle
Alethic Ideal	(AMU) Do ϕ only if ϕ maximizes utility in the *actual* world.
Consistency	(PMU) Do ϕ only if ϕ maximizes u in *some possible* world.
Coherence	(MEU) Do ϕ only if ϕ maximizes EU (relative to *some* Pr).
(WADA)	(WDOM) Do ϕ only if ϕ is *not weakly dominated* in utility.
(SADA)	(SDOM) Do ϕ only if ϕ is *not strictly dominated* in utility.

- Like the **Alethic Ideal**, (AMU) is *not* a *requirement of rationality*; and, like **Consistency**, (PMU) isn't a rational requirement either (this was the lesson of **Miners** [34, 26]).
- As Foley (*op. cit.*) explains, **Consistency** is *too demanding*. But, **Coherence** is *not* — it does *not* "pressure us to believe only those propositions that are (close to) certain for us".

- I've been presenting epistemic requirements as if they applied to "doxastic acts" of believing, disbelieving or suspending judgment (or assigning some credence).
- Strictly speaking, I should present both epistemic and prudential requirements as constraints on *preferences*.
- For instance, the key evaluative claim about **Miners** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{C} \succ \mathsf{A} \sim \mathsf{B}$$

is not irrational — because it is aligned with the agent's expected utility ranking (where $\mathsf{C} \stackrel{\text{def}}{=}$ blocking neither shaft).

- Similarly, the key evaluative claim about **Gibbard's Coin** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{S} \succ \mathsf{H} \sim \mathsf{T}$$

is not irrational — since it is aligned with expected *epistemic* utility (where $\mathsf{S} \stackrel{\text{def}}{=}$ believing neither H nor T).

- I've been presenting epistemic requirements as if they applied to "doxastic acts" of believing, disbelieving or suspending judgment (or assigning some credence).
- Strictly speaking, I should present both epistemic and prudential requirements as constraints on *preferences*.
- For instance, the key evaluative claim about **Miners** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{C} \succ \mathsf{A} \sim \mathsf{B}$$

is not irrational — because it is aligned with the agent's expected utility ranking (where $\mathsf{C} \stackrel{\text{def}}{=}$ blocking neither shaft).

- Similarly, the key evaluative claim about **Gibbard's Coin** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{S} \succ \mathsf{H} \sim \mathsf{T}$$

is not irrational — since it is aligned with expected *epistemic* utility (where $\mathsf{S} \stackrel{\text{def}}{=}$ believing neither H nor T).

- I've been presenting epistemic requirements as if they applied to "doxastic acts" of believing, disbelieving or suspending judgment (or assigning some credence).
- Strictly speaking, I should present both epistemic and prudential requirements as constraints on *preferences*.
- For instance, the key evaluative claim about **Miners** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{C} \succ \mathsf{A} \sim \mathsf{B}$$

is not irrational — because it is aligned with the agent's expected utility ranking (where $\mathsf{C} \stackrel{\text{def}}{=}$ blocking neither shaft).

- Similarly, the key evaluative claim about **Gibbard's Coin** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{S} \succ \mathsf{H} \sim \mathsf{T}$$

is not irrational — since it is aligned with expected *epistemic* utility (where $\mathsf{S} \stackrel{\text{def}}{=}$ believing neither H nor T).

- I've been presenting epistemic requirements as if they applied to "doxastic acts" of believing, disbelieving or suspending judgment (or assigning some credence).
- Strictly speaking, I should present both epistemic and prudential requirements as constraints on *preferences*.
- For instance, the key evaluative claim about **Miners** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{C} \succ \mathsf{A} \sim \mathsf{B}$$

  is not irrational — because it is aligned with the agent's expected utility ranking (where $\mathsf{C} \overset{\text{def}}{=}$ blocking neither shaft).
- Similarly, the key evaluative claim about **Gibbard's Coin** is (strictly speaking) that the (partial) *preference ranking*

$$\mathsf{S} \succ \mathsf{H} \sim \mathsf{T}$$

  is not irrational — since it is aligned with expected *epistemic* utility (where $\mathsf{S} \overset{\text{def}}{=}$ believing neither H nor T).

- If an agent does not have (precise) credences, expected inaccuracy minimization will not be an apt coherence requirement. But, we can still say *something* here.
- We can appeal to *non-dominance* requirements, such as:

Weak Accuracy-Dominance Avoidance (WADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(i) $(\forall w)[u(\mathbf{B}', w) \leq u(\mathbf{B}, w)]$, and

(ii) $(\exists w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

Strict Accuracy-Dominance Avoidance (SADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(iii) $(\forall w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

- It turns out [10, 11] that **Coherence** ⇒ (WADA) ⇒ (SADA).
- Indeed, (WADA) and (SADA) are *very* weak [10]. But, they do constitute non-trivial *necessary requirements* of rationality.

- If an agent does not have (precise) credences, expected inaccuracy minimization will not be an apt coherence requirement. But, we can still say *something* here.
- We can appeal to *non-dominance* requirements, such as:

Weak Accuracy-Dominance Avoidance (WADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(i) $(\forall w)[u(\mathbf{B}', w) \leq u(\mathbf{B}, w)]$, and

(ii) $(\exists w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

Strict Accuracy-Dominance Avoidance (SADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(iii) $(\forall w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

- It turns out [10, 11] that **Coherence** ⇒ (WADA) ⇒ (SADA).
- Indeed, (WADA) and (SADA) are *very* weak [10]. But, they do constitute non-trivial *necessary requirements* of rationality.

- If an agent does not have (precise) credences, expected inaccuracy minimization will not be an apt coherence requirement. But, we can still say *something* here.
- We can appeal to *non-dominance* requirements, such as:

Weak Accuracy-Dominance Avoidance (WADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(i) $(\forall w)[u(\mathbf{B}', w) \leq u(\mathbf{B}, w)]$, and

(ii) $(\exists w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

Strict Accuracy-Dominance Avoidance (SADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(iii) $(\forall w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

- It turns out [10, 11] that **Coherence** $\Rightarrow$ (WADA) $\Rightarrow$ (SADA).
- Indeed, (WADA) and (SADA) are *very* weak [10]. But, they do constitute non-trivial *necessary requirements* of rationality.

- If an agent does not have (precise) credences, expected inaccuracy minimization will not be an apt coherence requirement. But, we can still say *something* here.
- We can appeal to *non-dominance* requirements, such as:

Weak Accuracy-Dominance Avoidance (WADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(i) $(\forall w)[u(\mathbf{B}', w) \leq u(\mathbf{B}, w)]$, and

(ii) $(\exists w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

Strict Accuracy-Dominance Avoidance (SADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(iii) $(\forall w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

- It turns out [10, 11] that **Coherence** ⇒ (WADA) ⇒ (SADA).
- Indeed, (WADA) and (SADA) are *very* weak [10]. But, they do constitute non-trivial *necessary requirements* of rationality.

- If an agent does not have (precise) credences, expected inaccuracy minimization will not be an apt coherence requirement. But, we can still say *something* here.
- We can appeal to *non-dominance* requirements, such as:

Weak Accuracy-Dominance Avoidance (WADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(i) $(\forall w)[u(\mathbf{B}', w) \leq u(\mathbf{B}, w)]$, and

(ii) $(\exists w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

Strict Accuracy-Dominance Avoidance (SADA).

There does *not* exist an alternative belief set $\mathbf{B}'$ such that:

(iii) $(\forall w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

- It turns out [10, 11] that **Coherence** $\Rightarrow$ (WADA) $\Rightarrow$ (SADA).
- Indeed, (WADA) and (SADA) are *very* weak [10]. But, they do constitute non-trivial *necessary requirements* of rationality.

- If an agent does not have (precise) credences, expected inaccuracy minimization will not be an apt coherence requirement. But, we can still say *something* here.
- We can appeal to *non-dominance* requirements, such as:

Weak Accuracy-Dominance Avoidance (WADA).

> There does *not* exist an alternative belief set $\mathbf{B}'$ such that:
>
> (i) $(\forall w)[u(\mathbf{B}', w) \leq u(\mathbf{B}, w)]$, and
>
> (ii) $(\exists w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

Strict Accuracy-Dominance Avoidance (SADA).

> There does *not* exist an alternative belief set $\mathbf{B}'$ such that:
>
> (iii) $(\forall w)[u(\mathbf{B}', w) < u(\mathbf{B}, w)]$.

- It turns out [10, 11] that **Coherence** $\Rightarrow$ (WADA) $\Rightarrow$ (SADA).
- Indeed, (WADA) and (SADA) are *very* weak [10]. But, they do constitute non-trivial *necessary requirements* of rationality.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

    The Closure of Rational Belief Principle (CRBP).
    If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

    The No Known Contradictions Principle (NKCP).
    If S knows (at t) that $\perp$ is a logical contradiction, then it would *not* be rational for S to believe $\perp$ (at t).

    The Conjunction Principle (CP).
    If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\perp)\} \subseteq \mathbf{B}$ is incoherent.

$\neg$(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

  The Closure of Rational Belief Principle (CRBP).
  If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

  The No Known Contradictions Principle (NKCP).
  If S knows (at t) that $\perp$ is a logical contradiction, then it would *not* be rational for S to believe $\perp$ (at t).

  The Conjunction Principle (CP).
  If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \models q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\perp)\} \subseteq \mathbf{B}$ is incoherent.

¬(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

  The Closure of Rational Belief Principle (CRBP).
  If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

  The No Known Contradictions Principle (NKCP).
  If S knows (at t) that $\bot$ is a logical contradiction, then it would *not* be rational for S to believe $\bot$ (at t).

  The Conjunction Principle (CP).
  If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any **B** s.t. $\{B(p), D(q)\} \subseteq$ **B** is incoherent.

(NCB) Any **B** such that $\{B(\bot)\} \subseteq$ **B** is incoherent.

$\neg$(CP) *Not* every **B** s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq$ **B** is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

  The Closure of Rational Belief Principle (CRBP).
  If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

  The No Known Contradictions Principle (NKCP).
  If S knows (at t) that $\bot$ is a logical contradiction, then it would *not* be rational for S to believe $\bot$ (at t).

  The Conjunction Principle (CP).
  If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\bot)\} \subseteq \mathbf{B}$ is incoherent.

¬(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

  The Closure of Rational Belief Principle (CRBP).
  If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

  The No Known Contradictions Principle (NKCP).
  If S knows (at t) that $\perp$ is a logical contradiction, then it would *not* be rational for S to believe $\perp$ (at t).

  The Conjunction Principle (CP).
  If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\perp)\} \subseteq \mathbf{B}$ is incoherent.

$\neg$(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

  The Closure of Rational Belief Principle (CRBP).
  If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

  The No Known Contradictions Principle (NKCP).
  If S knows (at t) that $\bot$ is a logical contradiction, then it would *not* be rational for S to believe $\bot$ (at t).

  The Conjunction Principle (CP).
  If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\bot)\} \subseteq \mathbf{B}$ is incoherent.

$\neg$(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

> **The Closure of Rational Belief Principle** (CRBP).
> If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.
>
> **The No Known Contradictions Principle** (NKCP).
> If S knows (at t) that $\bot$ is a logical contradiction, then it would *not* be rational for S to believe $\bot$ (at t).
>
> **The Conjunction Principle** (CP).
> If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\bot)\} \subseteq \mathbf{B}$ is incoherent.

$\neg$(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- Sharon Ryan [39] gives an argument *for* (CB) as a rational requirement, which makes use of these three premises.

  The Closure of Rational Belief Principle (CRBP).
  If S rationally believes p at t and S knows (at t) that p entails q, then it would be rational for S to believe q at t.

  The No Known Contradictions Principle (NKCP).
  If S knows (at t) that $\perp$ is a logical contradiction, then it would *not* be rational for S to believe $\perp$ (at t).

  The Conjunction Principle (CP).
  If S rationally believes p at t and S rationally believes q at t, then it would be rational for S to believe $\ulcorner p \,\&\, q \urcorner$ at t.

- Ryan's (CRBP) & (NKCP) have analogues in our framework (which *are* coherence requirements). But, (CP) does *not*.

(SPC) If $p \vDash q$, then any $\mathbf{B}$ s.t. $\{B(p), D(q)\} \subseteq \mathbf{B}$ is incoherent.

(NCB) Any $\mathbf{B}$ such that $\{B(\perp)\} \subseteq \mathbf{B}$ is incoherent.

¬(CP) *Not* every $\mathbf{B}$ s.t. $\{B(p), B(q), D(p \,\&\, q)\} \subseteq \mathbf{B}$ is incoherent.

- I mentioned above that we assume "act-state independence" (ASI). There are two main reasons we assume (ASI) here.
- If $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then the verdicts delivered by **Coherence** can be *partition-sensitive*, *i.e.*, they can depend on the way in which the underlying set of doxastic possibilities is partitioned or carved up [21].
- More importantly, if $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then EUT can yield unintuitive (and/or odd) verdicts (even assuming a "natural" partition of states). See [4, 15, 5, 27].
- For instance, Carr [5] considers cases in which $B(p)$ and p are *positively* correlated (*e.g.*, believing you will do a handstand makes it much more likely that you will).
- Examples involving *negative* correlation between $B(p)$ and p have been discussed by various authors (*e.g.*, [15]). The most extreme (and difficult) examples along these lines are the self-referential examples due to Michael Caie [4].

- I mentioned above that we assume "act-state independence" (ASI). There are two main reasons we assume (ASI) here.
- If $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then the verdicts delivered by **Coherence** can be *partition-sensitive*, *i.e.*, they can depend on the way in which the underlying set of doxastic possibilities is partitioned or carved up [21].
- More importantly, if $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then EUT can yield unintuitive (and/or odd) verdicts (even assuming a "natural" partition of states). See [4, 15, 5, 27].
- For instance, Carr [5] considers cases in which $B(p)$ and p are *positively* correlated (*e.g.*, believing you will do a handstand makes it much more likely that you will).
- Examples involving *negative* correlation between $B(p)$ and p have been discussed by various authors (*e.g.*, [15]). The most extreme (and difficult) examples along these lines are the self-referential examples due to Michael Caie [4].

- I mentioned above that we assume "act-state independence" (ASI). There are two main reasons we assume (ASI) here.
- If $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then the verdicts delivered by **Coherence** can be *partition-sensitive*, *i.e.*, they can depend on the way in which the underlying set of doxastic possibilities is partitioned or carved up [21].
- More importantly, if $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then EUT can yield unintuitive (and/or odd) verdicts (even assuming a "natural" partition of states). See [4, 15, 5, 27].
- For instance, Carr [5] considers cases in which $B(p)$ and p are *positively* correlated (*e.g.*, believing you will do a handstand makes it much more likely that you will).
- Examples involving *negative* correlation between $B(p)$ and p have been discussed by various authors (*e.g.*, [15]). The most extreme (and difficult) examples along these lines are the self-referential examples due to Michael Caie [4].

- I mentioned above that we assume "act-state independence" (ASI). There are two main reasons we assume (ASI) here.
- If $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then the verdicts delivered by **Coherence** can be *partition-sensitive*, *i.e.*, they can depend on the way in which the underlying set of doxastic possibilities is partitioned or carved up [21].
- More importantly, if $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then EUT can yield unintuitive (and/or odd) verdicts (even assuming a "natural" partition of states). See [4, 15, 5, 27].
- For instance, Carr [5] considers cases in which $B(p)$ and p are *positively* correlated (*e.g.*, believing you will do a handstand makes it much more likely that you will).
- Examples involving *negative* correlation between $B(p)$ and p have been discussed by various authors (*e.g.*, [15]). The most extreme (and difficult) examples along these lines are the self-referential examples due to Michael Caie [4].

- I mentioned above that we assume “act-state independence” (ASI). There are two main reasons we assume (ASI) here.
- If $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then the verdicts delivered by **Coherence** can be *partition-sensitive*, *i.e.*, they can depend on the way in which the underlying set of doxastic possibilities is partitioned or carved up [21].
- More importantly, if $\mathbf{B}(p)$ and p are correlated under $b(\cdot)$, then EUT can yield unintuitive (and/or odd) verdicts (even assuming a “natural” partition of states). See [4, 15, 5, 27].
- For instance, Carr [5] considers cases in which $B(p)$ and p are *positively* correlated (*e.g.*, believing you will do a handstand makes it much more likely that you will).
- Examples involving *negative* correlation between $B(p)$ and p have been discussed by various authors (*e.g.*, [15]). The most extreme (and difficult) examples along these lines are the self-referential examples due to Michael Caie [4].

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence**.

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	−	+	−	−	×	×	×	×
w_2	**T**	**F**	×	×	×	×	−	−	−	+

- The "×"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

(P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence**.

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	−	+	−	−	×	×	×	×
w_2	**T**	**F**	×	×	×	×	−	−	−	+

- The "×"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence**.

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	$-$	$+$	$-$	$-$	$\times$	$\times$	$\times$	$\times$
w_2	**T**	**F**	$\times$	$\times$	$\times$	$\times$	$-$	$-$	$-$	$+$

- The "$\times$"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence**.

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	$-$	$+$	$-$	$-$	$\times$	$\times$	$\times$	$\times$
w_2	**T**	**F**	$\times$	$\times$	$\times$	$\times$	$-$	$-$	$-$	$+$

- The "$\times$"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence.**

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	−	+	−	−	×	×	×	×
w_2	**T**	**F**	×	×	×	×	−	−	−	+

- The "×"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence.**

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	$-$	$+$	$-$	$-$	$\times$	$\times$	$\times$	$\times$
w_2	**T**	**F**	$\times$	$\times$	$\times$	$\times$	$-$	$-$	$-$	$+$

- The "$\times$"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence.**

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	$-$	$+$	$-$	$-$	$\times$	$\times$	$\times$	$\times$
w_2	**T**	**F**	$\times$	$\times$	$\times$	$\times$	$-$	$-$	$-$	$+$

- The "$\times$"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- Caie's original example involved (only) *credences* [4]. It was designed to undermine Joycean (accuracy-dominance) arguments for *probabilism* as a requirement for $b(\cdot)$.
- There are analogous examples for full belief. Consider:

  (P) S does not believe that P. [$\neg B(\ulcorner P \urcorner)$.]

- One can argue (Caie-style) that the only non-dominated (opinionated) belief sets on $\{P, \neg P\}$ are $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$, which are both *ruled-out* by **Coherence.**

	P	$\neg P$	$B(P)$	$B(\neg P)$	$B(P)$	$D(\neg P)$	$D(P)$	$B(\neg P)$	$D(P)$	$D(\neg P)$
w_1	**F**	**T**	$-$	$+$	$-$	$-$	$\times$	$\times$	$\times$	$\times$
w_2	**T**	**F**	$\times$	$\times$	$\times$	$\times$	$-$	$-$	$-$	$+$

- The "$\times$"s indicate that these worlds are *ruled-out* (*a priori*) by the definition of P. As such, the only non-dominated belief sets seem to be $\{B(P), B(\neg P)\}$ and $\{D(P), D(\neg P)\}$.
- If this Caie-style reasoning is correct, then it shows that *some of our assumptions must go*. But, which one(s)?

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe.** Wherever B-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe**. Wherever B-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe**. Wherever B-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe.** Wherever B-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe.** Wherever B-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe.** Wherever B-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of the approach in [10].
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a group of consistent agents *need not* be consistent.
- **Q**: does majority rule preserve *our* notion(s) of coherence, *e,g,.*, is (WADA) preserved by MR? **A**: yes (on simple, atomic + truth-functional agendas), but *not on all possible agendas*.
  - There are (not merely atomic + truth-functional) agendas A and sets of judges J ($|A| \geq 5$, $|J| \geq 5$) that (severally) satisfy (WADA), while their majority profile *violates* (WADA).
- *But*, *if* a set of judges is (severally) *consistent* (or merely **Coherent**), *then* their majority profile *must* be **Coherent**.

☞ **Recipe.** Wherever **B**-*consistency* runs into paradox, substitute *coherence* (in *our* sense), and see what happens.

[1] L. Bonjour, *The Coherence Theory of Empirical Knowledge*, *Phil. Studies*, 1975.

[2] R. Briggs, F. Cariani, K. Easwaran, B. Fitelson, *Individual Coherence and Group Coherence*, *Essays in Collective Epistemology*, J. Lackey (*ed.*), OUP, 2014.

[3] J. Broome, *Wide or Narrow Scope*, *Mind*, 2007.

[4] M. Caie, *Rational Probabilistic Incoherence*, *Philosophical Review*, 2013.

[5] J. Carr, *How to Expect When You're Expecting*, 2014.

[6] D. Christensen, *Putting Logic in its Place*, OUP, 2007.

[7] W. Clifford, *The ethics of belief*, *Contemporary Review*, 1877.

[8] B. de Finetti, *The Theory of Probability*, Wiley, 1974.

[9] K. Dorst, *Lockeans Maximize Expected Accuracy*, 2015.

[10] K. Easwaran and B. Fitelson, *Accuracy, Coherence, and Evidence*, in *Oxford Studies in Epistemology V*, T. Szabo-Gendler and J. Hawthorne (*eds.*), OUP, 2015.

[11] K. Easwaran, *Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probability*, *Noûs*, 2015.

[12] ______, *Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection*, *Philosophy of Science*, 2013.

[13] R. Foley, *Working Without a Net*, Oxford University Press, 1992.

[14] A. Gibbard, *Truth and correct belief*, *Philosophical Issues*, 2005.

[15] H. Greaves, *Epistemic Decision Theory*, *Mind*, 2013.

[16] H. Greaves and D. Wallace, *Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility*, *Mind*, 2006.

[17] A. Hájek, *Arguments for — or Against — Probabilism?*, *BJPS*, 2008.

[18] C. Hempel, *Deductive-Nomological vs. Statistical Explanation*, 1962.

[19] W. James, *The will to believe*, *The New World*, 1896.

[20] J. Joyce, *Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief*, in F. Huber and C. Schmidt-Petri (*eds.*), *Degrees of Belief*, 2009.

[21] ______, *The Foundations of Causal Decision Theory*, CUP, 1999.

[22] ______, *A Nonpragmatic Vindication of Probabilism*, *Philosophy of Science*, 1998.

[23] M. Kaplan, *Coming to Terms with our Human Fallibility*, *PPR*, 2013.

[24] P. Klein, *The virtues of inconsistency*, *The Monist*, 1985.

[25] N. Kolodny, *How Does Coherence Matter?*, *Proc. of the Aristotelian Society*, 2007.

[26] N. Kolodny and J. MacFarlane, *Ifs and oughts*, *Journal of Philosophy*, 2010.

[27] J. Konek and B. Levinstein, *The Foundations of Epistemic Decision Theory*, 2014.

[28] H. Leitgeb and R. Pettigrew, *An Objective Justification of Bayesianism I & II*, *Philosophy of Science*, 2010.

[29] I. Levi, *Gambling with Truth*, 1967.

[30] J. MacFarlane, *Assessment Sensitivity: Relative Truth and its Applications*, OUP, 2014.

[31] P. Maher, *Betting on Theories*, 1993.

[32] D. Nelkin, *The lottery paradox, knowledge, and rationality*, *Phil. Review*, 2000.

[33] G. Oddie, *Conditionalization, cogency, and cognitive value*, *BJPS*, 1997.

[34] D. Parfit, *What we together do*, unpublished manuscript, 1988.

[35] ______, *Epistemic Utility Arguments for Probabilism*, *SEP*, 2011.

[36] J. Pollock, *Nomic probability and the foundations of induction*, OUP, 1990.

[37] J. Predd, R. Seringer, E. Loeb, D. Osherson, H.V. Poor and S., Kulkarni, *Probabilistic coherence and proper scoring rules*, *IEEE Transactions*, 2009.

[38] F. Ramsey, *Truth and Probability*, 1926.

[39] S. Ryan, *The epistemic virtues of consistency*, *Synthese*, 1996.

[40] M. Titelbaum, *Quitting Certainties*, OUP, 2013.