- 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 (like deductive consistency) supervene on *logical* 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* [32] requirements.

  (iii) *Epistemic* coherence involves *distinctively* epistemic values (specifically: *accuracy* [19] and *evidential support* [7]).
    - This is to be distinguished from *pragmatic* coherence (*e.g.*, immunity from dutch books [30], and the like [18]).

  (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 [32].

- Here is a — perhaps *the* — "paradigm" CR [29, 31, 28, 23].
  - **The Consistency Requirement for Belief** (CB). Agents should have *sets* of beliefs that are *logically consistent*.
- (CB) follows from the following (narrow-scope) *norm*:
  - **The Truth Norm for Belief** (TB). Agents should have beliefs that are *true* (*i.e.*, each *individual* belief should be true).
- Alethic norms [(CB)/(TB)] can conflict with evidential norms.
  - **The Evidential Norm for Belief** (EB). Agents should have beliefs that are *supported by the evidence*.
- In some cases (*e.g.*, preface cases), agents satisfy (EB) while violating (CB) — this generates an alethic/evidential *conflict*.
- Such alethic/evidential conflicts needn't give rise to states that receive an (overall) evaluation as *ir*rational (nor must they inevitably give rise to rational *dilemmas*) [6, 25, 15, 24].
- We'll refer to the claim that there exist *some* such cases as *the datum*. Foley's [15] 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 not argue for *the datum* here. We think Foley [15], Christensen [6], Kolodny [25], and others have made a compelling case for it. Today, it is our *point of departure*. But, we do have our own favorite (first-order) Preface case.

First-Order Preface Paradox. John is an excellent empirical scientist. He has devoted his entire (long and esteemed) scientific career to gathering and assessing the evidence that is relevant to the following first-order, empirical hypothesis: (H) all scientific/empirical books of sufficient complexity contain at least one false claim. By the end of his career, John is ready to publish his masterpiece, which is an exhaustive, encyclopedic, 15-volume (scientific/empirical) book which aims to summarize (all) the evidence that contemporary empirical science takes to be relevant to H. John sits down to write the Preface to his masterpiece. Rather than reflecting on his own fallibility, John simply reflects on the contents of (the main text of) his book, which constitutes *very strong inductive evidence in favor of* H. On this basis, John (inductively) infers H. But, John also believes each of the individual claims asserted in the main text of the book. Thus, because John believes (indeed, knows) that his masterpiece instantiates the antecedent of H, the (total) set of John's (rational/justified) beliefs is inconsistent.

☞ John's **B** is *alethically*, but *not* ***evidentially***, inconsistent.

Evidential Consistency (EC). A judgment set is *evidentially consistent* just in case there exists *some* (*possible*) body of total evidence E which supports each of its members.

- Some philosophers construe *the datum* as reason to believe that (⋆) *there are no coherence requirements for full belief*.
- Christensen [6] thinks (a) *credences* do have coherence requirements (*probabilism*); (⋆) full beliefs do *not*; (b) what *seem* to be CRs for full belief can be explained *via* (a).
- Kolodny [25] agrees with (⋆), but he disagrees with (a) and (b). He thinks (c) full belief is *explanatorily indispensable*; (d) there are *no* coherence requirements for *any* judgments; (e) what *seem* to be CRs for full belief can be explained *via* (EB).
- Christensen & Kolodny *agree — trivially*, *via* (⋆) — that:

  (†) *If* there are *any* coherence requirements for full belief, *then* (CB) is a coherence requirement for full belief.
- We [2, 12] agree with Christensen on (a) and Kolodny on (c), but we disagree with them on (⋆), (d), (e), and (†). We'll explain how to ground "conflict-proof" CRs for full belief, by analogy with Joyce's [22, 20] argument(s) for probabilism.

- We begin with some background assumptions/notation.
  - $B(p) \overset{\text{def}}{=} S$ believes that p. $D(p) \overset{\text{def}}{=} S$ disbelieves that p.
- S makes judgments regarding propositions in a (finite) *agenda* ($\mathcal{A}$) of (classical, possible-worlds) propositions. We'll use "**B**" to denote the *set* of S's judgments on $\mathcal{A}$.

☞ We're only evaluating *explicit judgments* (*on* $\mathcal{A}$) — we assume nothing about off-agenda commitments.

- We'll make two key assumptions about B/D on $\mathcal{A}$. The first assumption is integral to the framework. The second assumption is made for simplicity (and can be relaxed).
  - **Accuracy conditions**. $B(p)$ $[D(p)]$ is accurate iff p is T [F].

  - **Opinionation**. $B(p) \vee D(p)$.
- See Extras (27) for Kenny's [11] relaxation of Opinionation.
- We will assume *belief/world independence*. [Extras (15) contains a problematic example in which this assumption fails.]

- Now, we can explain how our new CRs were discovered, by analogy with Joyce's [22, 20] argument(s) for probabilism.
- Both arguments can be seen as involving three key steps.
- **Step 1**: Define $\mathring{\mathbf{B}}_w$ — the *vindicated (viz., alethically ideal* or *perfectly accurate)* judgment set (on $\mathcal{A}$), at world w.
  - $\mathring{\mathbf{B}}_w$ contains $B(p)$ $[D(p)]$ iff p is true (false) at w.
    - Heuristically, we can think of $\mathring{\mathbf{B}}_w$ as the set of judgments that an omniscient agent would have (on $\mathcal{A}$, at w).
- **Step 2**: Define $d(\mathbf{B}, \mathring{\mathbf{B}}_w)$ — a measure of distance between **B** and $\mathring{\mathbf{B}}_w$. That is, a measure of **B**'s *distance from vindication*.
  - $d(\mathbf{B}, \mathring{\mathbf{B}}_w) \overset{\text{def}}{=}$ the number of inaccurate judgments in **B** at w.
    - *Hamming distance* [9] between the binary vectors **B**, $\mathring{\mathbf{B}}_w$.
- **Step 3**: Adopt a *fundamental epistemic principle*, which uses $d(\mathbf{B}, \mathring{\mathbf{B}}_w)$ to ground a coherence requirement for **B**.
- This last step is the philosophically crucial one...

- Given our choices at Steps 1 and 2, there is ***a*** choice we can make at Step 3 that will yield (CB) as a requirement for **B**.

  Possible Vindication (PV). There exists some possible world w at which all of the judgments in **B** are accurate. Or, to put this more formally, in terms of d: $(\exists w)[d(\mathbf{B}, \mathring{\mathbf{B}}_w) = 0]$.
- Possible vindication is *one* way we could go here. But, our framework is much more general than the classical one. It allows for many other choices of fundamental principle.
- Like Joyce [22, 20] — who makes the analogous move with credences, to ground probabilism — we retreat from (PV) to the weaker: *avoidance of (weak) dominance in* $d(\mathbf{B}, \mathring{\mathbf{B}}_w)$.

  Weak Accuracy-Dominance Avoidance (WADA).

  > There does *not* exist an alternative belief set $\mathbf{B}'$ such that:
  > (i) $(\forall w)[d(\mathbf{B}', \mathring{\mathbf{B}}_w) \le d(\mathbf{B}, \mathring{\mathbf{B}}_w)]$, and
  > (ii) $(\exists w)[d(\mathbf{B}', \mathring{\mathbf{B}}_w) < d(\mathbf{B}, \mathring{\mathbf{B}}_w)]$.
- Completing Step 3 in this way reveals new CRs for **B**...

- Ideally, we want a coherence requirement that [like (CB)] can be motivated by considerations of *accuracy* (*viz.*, a CR that is *entailed by* alethic requirements such as TB/CB/PV).
- ☞ But, in light of (*e.g.*) preface cases, we also want a CR that is *weaker* than (CB). More precisely, we want a CR that is weaker than (CB) *in such a way that it is also entailed by* (EB).
- We can show that our new CRs [*e.g.*, (WADA)] fit the bill, *if* we assume the following "probabilistic-evidentialist" necessary condition for the satisfaction of (EB).

  Necessary Condition for Satisfying (EB). **B** satisfies (EB), *i.e.*, all judgments in **B** are *supported by the evidence*, ***only** if*:

  ($\mathcal{R}$) There exists *some* Pr-function that probabilifies (*i.e.*, assigns Pr greater than 1/2 to) each belief in **B** and dis-probabilifies (*i.e.*, assigns Pr less than 1/2 to) each disbelief in **B**.
- "Probabilistic-evidentialists" will disagree about *which* $\Pr(\cdot)$ undergirds (EB) [5, 33, 16, 21]; but, they agree on (EB) ⇒ ($\mathcal{R}$). Indeed, advocates of (PrE) will hold that (EB) ⇒ (EC) ⇒ ($\mathcal{R}$).

- Here are the logical relationships between key norms:

Truth Norm for Belief:	(TB)
	⇓ ⇑̸
Consistency Norm for Belief (*viz.*, PV):	(CB)/(PV)
	⇓ ⇑̸
☞ Weak Accuracy-Dominance Avoidance:	(WADA)
	⇑ ⇓̸
Evidential Norm for Belief:	(EB)

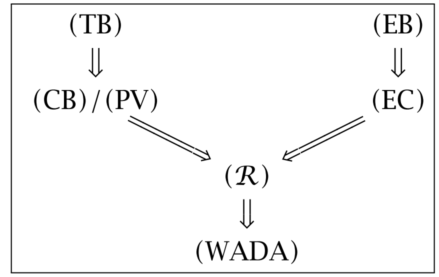

- See slide #18 for a bigger map w/11 requirements/norms.

- There are many advantages to adopting ($\mathcal{R}$), rather than (WADA), as our (ultimate) CR for full belief. Here are a few:
  - First, (WADA) is (intuitively) *too weak* to serve as our (ultimate) CR — $\{B(p), B(\neg p)\}$ may be *non-dominated*, as the following table reveals (*ditto* for $\{D(p), D(\neg p)\}$).

	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**	+	−	+	+	−	−	−	+

  - ($\mathcal{R}$) ⇒ (NCP) $D(p) \equiv B(\neg p)$, which *rules-out* $\{B(p), B(\neg p)\}$.
  - ($\mathcal{R}$) is strictly stronger than (WADA) + (NCP). Indeed, we conjecture that ($\mathcal{R}$) is the strongest CR (uncontroversially) entailed by both alethic and evidential considerations.
  - ($\mathcal{R}$) entails (WADA_d), for *any additive distance measure d*. In this sense, ($\mathcal{R}$) is *robust* across choices of d.
  - (WADA) only makes sense for *finite* agendas, whereas ($\mathcal{R}$) is potentially applicable to *infinite* agendas (if there be such).

- It is useful to draw an analogy between the norms and requirements we've been discussing, and principles in rational choice theory. The Decision-Theoretic Analogy.

Epistemic Principle	Analogous Decision-Theoretic Principle
(TB)	(AMU) Do ϕ only if ϕ maximizes utility in the *actual* world.
(CB)	(PMU) Do ϕ only if ϕ maximizes u in *some possible* world.
($\mathcal{R}$)	(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 (TB), (AMU) is *not* a *requirement of rationality*; and, like (CB), (PMU) isn't a rational requirement either. Moreover, also like (CB), seeing this requires "paradoxical" cases [26].
- As Foley (*op. cit.*) explains, (CB) is *too demanding*. But, ($\mathcal{R}$) and (WADA) are *not* — they do *not* "pressure us to believe only those propositions that are (close to) certain for us".

- Sharon Ryan [31] 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.

- Michael Caie [4] writes about an example involving self-reference, which causes problems for Joyce-style (accuracy-dominance) arguments for *probabilism*.
- 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 ($\mathcal{R}$).

	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)?

$\mathcal{B}$	$\mathbf{B}_1$	$\mathbf{B}_2$
$\neg X \,\&\, \neg Y$	D	D
$X \,\&\, \neg Y$	D	D
$X \,\&\, Y$	D	D
$\neg X \,\&\, Y$	D	D
$\neg Y$	D	D
$X \equiv Y$	D	D
$\neg X$	B	B
X	B	D
$\neg(X \equiv Y)$	D	D
Y	D	D
$X \vee \neg Y$	B	B
$\neg X \vee \neg Y$	B	B
$\neg X \vee Y$	B	B
$X \vee Y$	D	B
$X \vee \neg X$	B	B
$X \,\&\, \neg X$	D	D

We have the following four facts regarding $\mathbf{B}_1$ and $\mathbf{B}_2$:

(1) $\mathbf{B}_1$ is *weakly* dominated in distance from vindication by $\mathbf{B}_2$ (this is easily verified by simple counting). Thus, $\mathbf{B}_1$ *violates* (WADA).

(2) $\mathbf{B}_1$ is *not strictly* dominated in distance from vindication by *any* belief set over $\mathcal{B}$ (this can be verified *via exhaustive search* on the set of all belief sets over $\mathcal{B}$). Thus, $\mathbf{B}_1$ *satisfies* (SADA).

(3) $\mathbf{B}_2$ is *not weakly* dominated (in distance from vindication) by *any* belief set over $\mathcal{B}$ (this can be verified *via exhaustive search* on the set of all belief sets over $\mathcal{B}$). Thus, $\mathbf{B}_2$ *satisfies* (WADA).

(4) $\mathbf{B}_2$ is *not* represented (in the sense of Definition 2) by *any* probability function on $\mathcal{B}$, since the set $\mathbf{B}_2$ contains two contradictory judgment pairs: $\{D(Y), D(\neg Y)\}$ and $\{D(X \equiv Y), D(\neg(X \equiv Y))\}$. Therefore, $\mathbf{B}_2$ *violates* ($\mathcal{R}$).

(TB) S ought believe p iff p is true.

(PV) $(\exists w)[d(\mathbf{B}, \mathring{\mathbf{B}}_w) = 0]$. That is, $\mathbf{B}$ is *deductively consistent*.

(SADA) $\nexists \mathbf{B}'$ such that: $(\forall w)[d(\mathbf{B}', \mathring{\mathbf{B}}_w) < d(\mathbf{B}, \mathring{\mathbf{B}}_w)]$.

(NW$_2$S) $\nexists \beta \subseteq \mathbf{B}$ s.t.: $(\forall w)$ [> 1/2 *of the members of* β *are inaccurate at* w].

($\mathcal{R}_r$) $\exists$ a probability function $\Pr(\cdot)$ such that, $\forall p \in \mathcal{A}$:

$$B(p) \textit{ iff } \Pr(p) > r, \text{ and } D(p) \textit{ iff } \Pr(p) < 1 - r.$$

(EB) S ought believe p iff p is supported by the (*actual*) evidence.

(EC) There is *some* (*possible*) E which supports each p in S's $\mathbf{B}$.

(NWS) $\nexists \beta \subseteq \mathbf{B}$ s.t.: $(\forall w)$ [≥ 1/2 *of the members of* β *are inaccurate at* w] & $(\exists w)$ [> 1/2 *of the members of* β *are inaccurate at* w].

(WADA) $\nexists \mathbf{B}'$ s.t.: $(\forall w)[d(\mathbf{B}', \mathring{\mathbf{B}}_w) \leq d(\mathbf{B}, \mathring{\mathbf{B}}_w)] \,\&\, (\exists w)[d(\mathbf{B}', \mathring{\mathbf{B}}_w) < d(\mathbf{B}, \mathring{\mathbf{B}}_w)]$.

(NW$_1$S) $\nexists \beta \subseteq \mathbf{B}$ s.t.: $(\forall w)$ [≥ 1/2 *of the members of* β *are inaccurate at* w].

(NCP) S disbelieves p iff S believes $\neg p$ [*i.e.*, $D(p) \equiv B(\neg p)$].

- Here is what the logical relations look like, among all of the 11 requirements & norms for (opinionated) full belief.

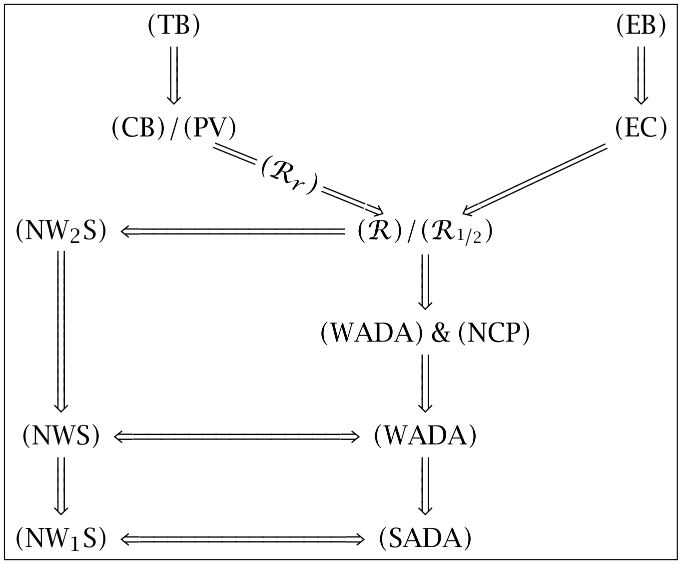

- Proof of the central result that $(\mathcal{R}) \Rightarrow$ (WADA).

Let Pr be a probability function that represents **B** in sense of $(\mathcal{R})$. Consider the expected distance from vindication of a belief set — the sum of $\Pr(w) \cdot d(\mathbf{B}, \mathring{\mathbf{B}}_w)$. Since $d(\mathbf{B}, \mathring{\mathbf{B}}_w)$ is a sum of components for each proposition (1 if **B** disagrees with w on the proposition and 0 if they agree), and since expectations are linear, the expected distance from vindication is the sum of the expectation of these components. The expectation of the component for disbelieving p is $\Pr(p)$ while the expectation of the component for believing p is $1 - \Pr(p)$. Thus, if $\Pr(p) > 1/2$ then believing p is the attitude that uniquely minimizes the expectation, while if $\Pr(p) < 1/2$ then disbelieving p is the attitude that uniquely minimizes the expectation. Thus, since Pr represents **B**, this means that **B** has strictly lower expected distance from vindication than any other belief set with respect to Pr. Suppose, for *reductio*, that some **B**′ (weakly) dominates **B**. Then, **B**′ must be no farther from vindication than **B** in any world, and thus **B**′ must have expected distance from vindication no greater than that of **B**. But **B** has strictly lower expected distance from vindication than any other belief set. Contradiction. ∴ **B** must be non-dominated.

- The key to our central theorem that $(\mathcal{R}) \Rightarrow$ (WADA) is that our inaccuracy measure $d(\mathbf{B}, \mathring{\mathbf{B}}_w)$ is *evidentially proper*.

Definition (Evidential Propriety)

Suppose a judgment set **J** of type $\mathfrak{J}$ is supported by the evidence. That is, suppose there exists some evidential probability function $\Pr(\cdot)$ which represents **J** (in the appropriate sense of "represents" for sets of type $\mathfrak{J}$). If this is sufficient to ensure that **J** minimizes expected inaccuracy (relative to Pr), according to the measure of inaccuracy $\mathfrak{I}(\mathbf{J}, \mathring{\mathbf{J}}_w)$, then we will say that the measure $\mathfrak{I}$ is **evidentially proper**.

- If an inaccuracy measure is evidentially *im*proper, then some probabilistically representable judgment sets will be *ruled out* as *irrational via* accuracy-dominance (WADA).
- This would engender a *conflict* between alethic and evidential requirements for judgment, which is exactly what coherence requirements are *not* supposed to do.
- In our book [13], evidential propriety plays a central role.

- Proof of the claim that (NWS) $\Leftrightarrow$ (WADA).
  - ($\Leftarrow$) We'll prove the contrapositive. Suppose that some $\mathbf{S} \subseteq \mathbf{B}$ is a witnessing set. Let **B**′ agree with **B** on all judgments outside **S** and disagree with **B** on all judgments in **S**. By the definition of a witnessing set, **B**′ weakly dominates **B** in distance from vindication $[d(\mathbf{B}, \mathring{\mathbf{B}}_w)]$.
  - ($\Rightarrow$) [Contrapositive again.] Suppose **B** is dominated, *i.e.*, that there is some **B**′ that weakly dominates **B** in distance from vindication $[d(\mathbf{B}, \mathring{\mathbf{B}}_w)]$. Let $\mathbf{S} \subseteq \mathbf{B}$ be the set of judgments on which **B** and **B**′ disagree. Then, **S** is a witnessing set.
- A similar proof can be given for: (NW$_1$S) $\Leftrightarrow$ (SADA).
- We also know that $(\mathcal{R}) \Rightarrow$ (NW$_2$S). See next slide for a proof.
  - The converse (NW$_2$S) $\overset{?}{\Rightarrow}$ $(\mathcal{R})$ remained open for several years, but was recently settled (negatively) — see slide #24.
- One final (positive) result: $(\mathcal{R})$ is *strictly stronger* than the conjunction (WADA) & (NCP). See slide #23 for a proof.

Theorem

$(\mathcal{R}) \Rightarrow (\text{NW}_2\text{S})$.

Proof.

In our proof (slide #19) of the claim that $(\mathcal{R}) \Rightarrow$ (WADA), we established that if Pr represents **B**, then **B** has strictly lower expected distance from vindication than any other belief set with respect to Pr. Assume, for *reductio*, that $\mathbf{S} \subseteq \mathbf{B}$ is a witnessing$_2$ set for **B**. Let $\mathbf{B}'$ agree with **B** on all judgments outside **S** and disagree with **B** on all judgments in **S**. Then by the definition of a witnessing$_2$ set, $\mathbf{B}'$ must be no farther from vindication than **B** in any world. But this contradicts the fact that **B** has strictly lower expected distance from vindication than $\mathbf{B}'$ with respect to Pr. So the witnessing$_2$ set must not exist. □

Theorem

(NDB & NCP) $\not\Rightarrow$ $(\mathcal{R})$. [In other words, (WADA & NCP) $\not\Rightarrow$ $(\mathcal{R})$.]

Proof.

Let there be six possible worlds, $w_1, w_2, w_3, w_4, w_5, w_6$. And, let $\mathcal{A} \stackrel{\text{def}}{=} \{p_1, p_2, p_3, p_4\}$, where the p_i are defined as follows.

$$p_1 \stackrel{\text{def}}{=} \{w_1, w_2, w_3\} \qquad p_2 \stackrel{\text{def}}{=} \{w_1, w_4, w_5\}$$

$$p_3 \stackrel{\text{def}}{=} \{w_2, w_4, w_6\} \qquad p_4 \stackrel{\text{def}}{=} \{w_3, w_5, w_6\}$$

Let $\mathbf{B} \stackrel{\text{def}}{=} \{B(p_1), B(p_2), B(p_3), B(p_4)\}$. **B** is a witnessing$_2$ set, since, in every w_i, *exactly half* of the beliefs in **B** are accurate. So, by $(\mathcal{R}) \Rightarrow (\text{NW}_2\text{S})$, **B** *violates* $(\mathcal{R})$. But, **B** *satisfies* (NDB), since every belief set on $\mathcal{A}$ has an expected distance from vindication of 2, relative to the uniform Pr-distribution, which implies that no belief set on $\mathcal{A}$ dominates any other belief set on $\mathcal{A}$. Finally, **B** satisfies (NCP), since every pair of beliefs in **B** is consistent. □

Theorem

$(\text{NW}_2\text{S}) \not\Rightarrow (\mathcal{R})$.

Proof.

Let there be twelve possible worlds, $w_1, \ldots, w_{12}$. And, let $\mathcal{A} \stackrel{\text{def}}{=} \{p_1, \ldots, p_6\}$, where the p_i are defined as follows.[a]

$$p_1 \stackrel{\text{def}}{=} \{w_1, w_2, w_3, w_4, w_8\} \qquad p_2 \stackrel{\text{def}}{=} \{w_1, w_2, w_5, w_6, w_9\}$$

$$p_3 \stackrel{\text{def}}{=} \{w_1, w_3, w_5, w_7, w_{10}\} \qquad p_4 \stackrel{\text{def}}{=} \{w_1, w_4, w_6, w_7, w_{11}\}$$

$$p_5 \stackrel{\text{def}}{=} \{w_2, w_3, w_4, w_5, w_6, w_7, w_{12}\}$$

$$p_6 \stackrel{\text{def}}{=} \{w_8, w_9, w_{10}, w_{11}, w_{12}\}$$

Let $\mathbf{B} \stackrel{\text{def}}{=} \{B(p_1), B(p_2), B(p_3), B(p_4), B(p_5), B(p_6)\}$. It can be shown that (a) **B** contains no witnessing$_2$ set (**B** *satisfies* NW_2S), but (b) **B** has no probabilistic representation (**B** *violates* $\mathcal{R}$). □

[a] This counterexample was discovered by Johannes Marti (ILLC).

Parametric Family of Requirements Between $(\mathcal{R})$ and (CB)

$(\mathcal{R}_r)$ There is a probability function Pr such that, for all $p \in \mathcal{A}$:

(i) **B** contains $B(p)$ iff $\Pr(p) > r$, and

(ii) **B** contains $D(p)$ iff $\Pr(p) < 1 - r$,

where $r \in [1/2, 1)$.

- Let $\mathbb{B}_n$ denote the class of minimal inconsistent belief sets of size n — each member of $\mathbb{B}_n$ is an inconsistent judgment set of size n containing no inconsistent proper subset.
- Let $\mathbf{B}_n$ be a member of $\mathbb{B}_n$, *i.e.*, $\mathbf{B}_n$ consists of n propositions, there is no world in which all of these n propositions are true, but for each proper subset $\mathbf{B} \subset \mathbf{B}_n$ there is a world in which all members of **B** are true.

Theorem

For all $n \geq 2$, if $r \geq \frac{n-1}{n}$ then $(\mathcal{R}_r)$ rules out each member of $\mathbb{B}_n$, while if $r < \frac{n-1}{n}$, then $(\mathcal{R}_r)$ rules out no member of $\mathbb{B}_n$.

Proof.

Let $\mathbf{B}_n \stackrel{\text{def}}{=} \phi_1, \ldots, \phi_n$. Let each w_i be a world in which ϕ_i is false, but all other members of $\mathbf{B}_n$ are true. Let Pr be the probability function that assigns $1/n$ to each world w_i and 0 to all other worlds. If $r < n{-}1/n$, then Pr shows $\mathbf{B}_n$ satisfies $(\mathcal{R}_r)$. This establishes the second half of the Theorem.

For the first half of the Theorem, we proceed *via reductio*. Suppose (for *reductio*) $\mathbf{B}_n$ is a member of $\mathbb{B}_n$ that is *not* ruled out by $(\mathcal{R}_{n-1/n})$. Then there must be some Pr such that for each i, $\Pr(\phi_i) > n{-}1/n$. Therefore, for each i, $\Pr(\neg\phi_i) < 1/n$. Now, since the disjunction of finitely many propositions is at most as probable as the sum of their individual probabilities, we must have $\Pr(\neg\phi_1 \vee \ldots \vee \neg\phi_n) < 1$. But, since $\mathbf{B}_n$ is inconsistent, $\neg\phi_1 \vee \ldots \vee \neg\phi_n$ is a tautology, and must have probability 1. Contradiction. So $\mathbf{B}_n$ must be ruled out by $(\mathcal{R}_{n-1/n})$. □

- Kenny has written a paper [11] that explains how to relax the assumption of Opinionation in our framework.
- Our approach is equivalent to assigning (in)accurate judgments a *score* of (−1) +1, and calculating the *total score* of **B** (at w) as the *sum* of the scores of all $p \in \mathcal{A}$.
- Kenny's Generalizations: (a) allow scores of $-\mathfrak{w}$ and $+\mathfrak{r}$, where $\mathfrak{w} \geq \mathfrak{r} > 0$, and (b) allow S to *suspend on* p $[S(p)]$, where all suspensions are given a *neutral* score of *zero*.
- This generalization of our framework leads to an elegant analogue of our central Theorem that $(\mathcal{R})$ entails (WADA).

Theorem. An agent S will avoid (strict) dominance in *total score* ***if*** their belief set **B** can be represented as follows:

$(\mathfrak{R})$ There exists a probability function $\Pr(\cdot)$ such that, $\forall p \in \mathcal{A}$:

$$B(p) \textit{ iff } \Pr(p) > \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$

$$D(p) \textit{ iff } \Pr(p) < 1 - \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}},$$

$$S(p) \textit{ iff } \Pr(p) \in \left[1 - \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}, \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right].$$

- We (along with Rachael Briggs and Fabrizio Cariani) [2] are investigating various applications of this new approach.
- One interesting application is to *judgment aggregation*. *E.g.*,
  - Majority rule aggregations of the judgments of a bunch of agents — each of whom satisfy (PV) — *need not* satisfy (PV).
- **Q**: does majority rule preserve *our* notion of coherence, *viz.*, 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* [*i.e.*, satisfy (PV)], then their majority profile *must* satisfy (WADA).

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

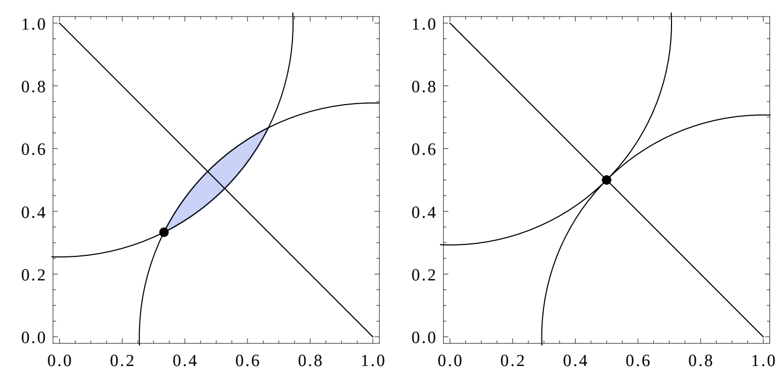

- **Simplest case of dF's Theorem [8].** The diagonal lines are the *probabilistic* b's (on $\langle P, \neg P\rangle$). The point $\langle 1, 0\rangle$ ($\langle 0, 1\rangle$) corresponds to the world in which P is true (false).

Theorem (de Finetti [8]). b is *non*-probabilistic $\Leftrightarrow \exists\, b'(\cdot)$ which is (Euclidean) *closer to* $v_w(\cdot)$ *in every possible world.*

- The plot on the left (right) explains the ⇒ (⇐) direction.

- In the first part of the book, I will present an argument for the key claim that "probabilities reflect evidence". That is:

(PrE) In each epistemic context (determined by a body of total evidence E), there is a (sharp, numerical) function $s(p, E)$ which measures *the degree to which E supports p* (for each p in $\mathcal{A}$), where $s(\cdot, E)$ is a *probability function* $\Pr(\cdot)$.

- The argument in Part I of the book is a variant of Joyce's [20] argument that *credences* ought to be *probabilistic*.
- This argument trades (only) on three assumptions regarding *measures* $\mathcal{I}(b, w)$ *of the gradational inaccuracy* of a credence function b at a possible world w.
- In general, measures of credal inaccuracy are measures of "distance" between a credence function b and the *indicator function* v_w at w (which determines the *alethic ideal* at w).
- These measures are assumed to be *continuous, truth-directed, and* ***probabilistically admissible.***

Probabilistic Admissibility (PA). A credal inaccuracy measure $\mathcal{I}$ is *probabilistically admissible* iff it never implies that any *probabilistic* credence function is (weakly) *dominated* in gradational inaccuracy.

- Here's how the argument of Part I goes (at a high level).
  1. The gradational inaccuracy of b (at w) is measured by a *continuous, truth-directed, probabilistically admissible* $\mathcal{I}$.
  2. If (1) is true, then $b(\cdot)$ is non-(weakly)-dominated in $\mathcal{I}$-accuracy *iff* b is *probabilistic*.
  3. If b is (weakly) dominated in $\mathcal{I}$-accuracy, then — *no matter what the total evidence E is* — $b(\cdot) \neq s(\cdot, E)$.
  4. $\therefore$ If (1) is true, then $b(\cdot) = s(\cdot, E)$ (for *some* body of total evidence E) *only if* $b(\cdot)$ — and $\therefore$ $s(\cdot, E)$ — is probabilistic.
- We can simplify steps 3 & 4 *via* the credal analogue of (EC).
  3. If $b(\cdot)$ is $\mathcal{I}$-dominated, then $b(\cdot)$ is *evidentially inconsistent*.
  4. $\therefore$ If (1) is true, then $b(\cdot)$ is evidentially (and alethically) consistent *only if* $b(\cdot)$ — and $\therefore$ $s(\cdot, E)$ — is *probabilistic*.

[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*, to appear in *Essays in Collective Epistemology*, J. Lackey (*ed.*), OUP.
[3] J. Broome, *Wide or Narrow Scope?*, *Mind*, 2007.
[4] M. Caie, *Rational Probabilistic Incoherence*, *The Philosophical Review*, 2013.
[5] R. Carnap, *Logical Foundations of Probability*, U. of Chicago, 2nd ed., 1962.
[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] M. Deza and E. Deza, *Encyclopedia of Distances*, Springer, 2009.
[10] C. Duddy and A. Piggins, *A measure of distance between judgment sets*, *Social Choice and Welfare*, 2012.
[11] K. Easwaran, *Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probability*, manuscript.
[12] K. Easwaran and B. Fitelson, *Accuracy, Coherence, and Evidence*, to appear in *Oxford Studies in Epistemology V*, T. Szabo-Gendler and J. Hawthorne (*eds.*), OUP.
[13] B. Fitelson, *Coherence*, book manuscript, 2014.
[14] B. Fitelson and D. McCarthy, *Toward an Epistemic Foundation for Comparative Confidence*, manuscript, 2014.

[15] R. Foley, *Working Without a Net*, Oxford University Press, 1992.
[16] R. Fumerton, *Metaepistemology and Skepticism*, Rowman & Littlefield, 1995.
[17] A. Hájek, *A Puzzle About Degree of Belief*, manuscript.
[18] ______, *Arguments for — or Against — Probabilism?*, BJPS, 2008.
[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] ______, *How Probabilities Reflect Evidence*, *Philosophical Perspectives 19*, 2005.
[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] D. Miller, *Out of error: Further essays on critical rationalism*, Ashgate, 2006.
[28] D. Nelkin, *The lottery paradox, knowledge, and rationality*, *Phil. Review*, 2000.
[29] J. Pollock, *Nomic probability and the foundations of induction*, OUP, 1990.
[30] F. Ramsey, *Truth and Probability*, 1928.
[31] S. Ryan, *The epistemic virtues of consistency*, *Synthese*, 1996.
[32] M. Titelbaum, *Quitting Certainties*, OUP, 2013.
[33] T. Williamson, *Knowledge and its Limits*, Oxford University Press, 2000.