

Knowledge Centered Epistemic Utility Theory

Julien Dutant¹ Branden Fitelson²

¹Philosophy @ KCL

julien.dutant@kcl.ac.uk

²Philosophy & Religion @ Northeastern

branden@fitelson.org

Outline

1. Prelude: Two Examples
2. Truth & The Old Lockeanism
3. Knowledge & A New Lockeanism
4. Extras: Two Supplementary Slides

- **Miners** [35, 27]. 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** [15, 31]. 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.*

- From a veritistic perspective, the imperative for belief is:
The Truth Imperative. Believe propositions that are *true!*
- From an evidentialist perspective, the belief imperative is:
The Evidential Imperative. Believe propositions that are *supported by the total evidence (i.e., propositions that are sufficiently probable, given the total evidence)*!
- **Gibbard's Coin** is a case in which these two imperatives *conflict*. Even more stark conflicts between these two imperatives can occur — *e.g.*, Lotteries, Prefaces.
- Veritistic Epistemic Utility Theory (VEUT) can provide a *reconciliation* of these conflicting epistemic imperatives.
- ☞ Even if belief aims at truth, this does not imply that rationality requires us to obey The Truth Imperative.
Veritistic Rational Requirement. Maximize the *expected accuracy* of your beliefs (as gauged by your credences)!
- Next: a formal explication of the VRR.

Prelude ○ Truth & The Old Lockeanism ○●○○ Knowledge & A New Lockeanism ○○○○ Extras ○○ Bibliography

- 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.
- Our agent will form judgments regarding each member of a finite agenda \mathcal{A} of (classical, possible worlds) propositions. They will either believe p , believe $\neg p$, or *suspend* re p .
- We do *not* assume that these qualitative judgments can be *reduced* to $b(\cdot)$. But, we will use $b(\cdot)$ to derive a *veritistic rational requirement* for qualitative belief sets \mathbf{B} (on \mathcal{A}).
- This derivation requires both the agent's credence function $b(\cdot)$ and their *epistemic utility function* [19, 30, 32] $u(\cdot)$.
 - ☞ Following Easwaran [12] & Dorst [10], we assume our agent cares *only* about the *accuracy* of their judgments (*veritism*).
- Specifically, our agent attaches some *positive* utility (r) with having an *accurate* belief, and some *negative* utility ($-\omega$) with having an *inaccurate* belief (where $\omega > r > 0$).

Dutant & Fitelson Knowledge Centered Epistemic Utility Theory 3

Prelude ○ Truth & The Old Lockeanism ○●○○ Knowledge & A New Lockeanism ○○○○ Extras ○○ Bibliography

- Because non-beliefs (*viz.*, suspensions) are neither accurate nor inaccurate (*per se*), our agent will attach *zero* epistemic utility to them, independently of the truth-value of p .
- This yields the following piecewise definition of $u(\cdot, w)$.

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

$$u(\neg Bp, w) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p \text{ is false at } w \\ 0 & \text{if } p \text{ is true at } w \end{cases}$$
- With this *accuracy-centered* epistemic utility function in hand, we can derive our veritistic rational requirement.
- To do so, we'll also need a *decision-theoretic principle*.
- We will adopt an *expected epistemic utility maximization* principle to derive our rational requirement [29, 17, 13, 34].

Dutant & Fitelson Knowledge Centered Epistemic Utility Theory 4

Prelude ○ Truth & The Old Lockeanism ○○○● Knowledge & A New Lockeanism ○○○○ Extras ○○ Bibliography

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

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

- Here, we assume "*act-state independence*": Bp and p are *b-independent* [16, 6, 5, 28]. See Extras for discussion.
- The consequences of (VRR) are rather simple and intuitive. It is straightforward to prove the following result.

Theorem ([12, 10]). An agent with credence function $b(\cdot)$ and belief set \mathbf{B} (over \mathcal{A}) satisfies (VRR) *iff* for all $p \in \mathcal{A}$

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

- ☞ That is, (VRR) *entails* The Evidential Imperative, where the meaning of "sufficiently probable" is determined by the way the agent (relatively) values accuracy vs. inaccuracy.

Dutant & Fitelson Knowledge Centered Epistemic Utility Theory 5

Prelude ○ Truth & The Old Lockeanism ○○○○ Knowledge & A New Lockeanism ●○○○ Extras ○○ Bibliography

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...

Veritism It is good to believe what is *true*.

Gnostic It is good to believe what is *known*.

- If we use the same models, but amend our epistemic utility function to take into account the value of knowing, then we may derive a **Gnostic Rational Requirement (GRR)**. To wit:

world w	$u(Bp, w)$	$u(\neg Bp, w)$
Kp	r	u
$p \ \& \ \neg Kp$	t	v
$\neg p$	$-\omega$	0

Dutant & Fitelson Knowledge Centered Epistemic Utility Theory 6

- We'll assume $u = v = 0$. (If $K \Rightarrow B$, u does not even arise.)
- Obvious constraints: $r \geq t \geq -w$ and $-w < 0$.
- Natural constraint: $w > r$.
- Main choice point:
 - **Pure Gnosticism.** $t = -w$. Believing something you do not know is bad — whether the thing is true or not.
 - **Mixed Gnosticism.** $t > -w$. Believing a truth is better than believing a falsehood — even if you don't know it.
- We'll focus on Pure Gnosticism. Simplest and most radical.
- Given the pure gnostic account of value and a subject's credence function b , we get the following GRR:

Neo-Lockeanism $p \in \mathbf{B}$ iff $b(Kp) \geq \frac{r}{r+w} > 1/2$.

☞ That is, you should believe that p just in case it is sufficiently probable that *you know* that p .

Lottery You own a ticket in a large lottery with 1 winner.

- Is it probable that your ticket is a loser? Yes.
- Is it probable that *you know that* your ticket is a loser? No.

Old (Veritistic) Lockeanism: believe your ticket will lose.

Neo-Lockeanism: do *not* believe your ticket will lose.

Preface You've well researched each claim in the book. You've also well aware that every book like yours ever published contained an error.

- Is it probable that your book contains an error? Yes.
- Is it probable that *you know that* your book contains an error? Yes.

Old/Neo-Lockeanism: believe your book contains an error.

- **Civil cases:** *Preponderance of the Evidence, or balance of probabilities* — *more likely than not* the defendant is guilty.

Clue 55% of the rodeo's audience gatecrashed. John was present. The tears on his pants are consistent with climbing over the fence.

Statistics 55% of the rodeo's audience gatecrashed. John was present.

- Is it probable that John gatecrashed? Yes.
- Is it probable *that we know that* John gatecrashed? Yes in the first, no in the second.
- Compare Blome-Tillmann [1] "the balance of probabilities" means "it is more likely than not *that we know.*"

- Contrast three views:
 - **Lockean:** believe if it is *probably true*.
 - **Neo-lockean:** believe if it is *probably known*.
 - **Gnostic idealist:** believe if it is *known*.

	Lockean	Neo-lockean	Gnostic idealist
Lottery beliefs	✓	✗	✗
Statistical gatecrasher	✓	✗	✗
Preface beliefs	✓	✓	✗
k inconsistent set	✗	✗	✗

- To Sum up:
 - Mixed Gnosticism? Essentially the same implications as Lockeanism.
 - If you think knowledge matters, the epistemic utility framework gives a systematic way of taking it into account.
 - Instead of knowledge, we can plug anything that you may care about ("sensitivity", "safety", "justification", *etc.*).

Prelude ○	Truth & The Old Lockeanism ○○○	Knowledge & A New Lockeanism ○○○○	Extras ●○	Bibliography
--------------	-----------------------------------	--------------------------------------	--------------	--------------

- We mentioned that we assume “act-state independence” (ASI). There are two main reasons we assume (ASI) here.
- If Bp and p are correlated under $b(\cdot)$, then the verdicts delivered by (VRR) 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 [22].
- More importantly, if Bp and p are correlated under $b(\cdot)$, then EUT can yield unintuitive (and/or odd) verdicts (even assuming a “natural” partition of states). See [5, 16, 6, 28].
- For instance, Carr [6] considers cases in which Bp 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 Bp and p have been discussed by various authors (*e.g.*, [16]). The most extreme (and difficult) examples along these lines are the self-referential examples due to Michael Caie [5].

Dutant & Fitelson	Knowledge Centered Epistemic Utility Theory	11
-------------------	---	----

Prelude ○	Truth & The Old Lockeanism ○○○	Knowledge & A New Lockeanism ○○○○	Extras ●○	Bibliography
--------------	-----------------------------------	--------------------------------------	--------------	--------------

- 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 [11, 12] that $(VRR) \Rightarrow (WADA) \Rightarrow (SADA)$.
- Indeed, (WADA) and (SADA) are *very* weak [11]. But, they do constitute non-trivial *necessary requirements* of rationality.

Dutant & Fitelson	Knowledge Centered Epistemic Utility Theory	12
-------------------	---	----

Prelude ○	Truth & The Old Lockeanism ○○○	Knowledge & A New Lockeanism ○○○○	Extras ○○	Bibliography
--------------	-----------------------------------	--------------------------------------	--------------	--------------

- [1] M. Bloome-Tillman, *Sensitivity, Causality, and Statistical Evidence in Courts of Law, Thought*, 2015.
- [2] L. Bonjour, *The Coherence Theory of Empirical Knowledge, Phil. Studies*, 1975.
- [3] R. Briggs, F. Cariani, K. Easwaran, B. Fitelson, *Individual Coherence and Group Coherence, Essays in Collective Epistemology*, J. Lackey (ed.), OUP, 2014.
- [4] J. Broome, *Wide or Narrow Scope, Mind*, 2007.
- [5] M. Caie, *Rational Probabilistic Incoherence, Philosophical Review*, 2013.
- [6] J. Carr, *How to Expect When You're Expecting*, 2014.
- [7] D. Christensen, *Putting Logic in its Place*, OUP, 2007.
- [8] W. Clifford, *The ethics of belief, Contemporary Review*, 1877.
- [9] B. de Finetti, *The Theory of Probability*, Wiley, 1974.
- [10] K. Dorst, *Lockeans Maximize Expected Accuracy*, 2015.
- [11] K. Easwaran and B. Fitelson, *Accuracy, Coherence, and Evidence*, in *Oxford Studies in Epistemology V*, T. Szabo-Gendler and J. Hawthorne (eds.), OUP, 2015.
- [12] K. Easwaran, *Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probability*, *Noûs*, 2015.
- [13] ———, *Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection*, *Philosophy of Science*, 2013.
- [14] R. Foley, *Working Without a Net*, Oxford University Press, 1992.
- [15] A. Gibbard, *Truth and correct belief, Philosophical Issues*, 2005.
- [16] H. Greaves, *Epistemic Decision Theory, Mind*, 2013.
- [17] H. Greaves and D. Wallace, *Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility*, *Mind*, 2006.

Dutant & Fitelson	Knowledge Centered Epistemic Utility Theory	13
-------------------	---	----

Prelude ○	Truth & The Old Lockeanism ○○○	Knowledge & A New Lockeanism ○○○○	Extras ○○	Bibliography
--------------	-----------------------------------	--------------------------------------	--------------	--------------

- [18] A. Hájek, *Arguments for — or Against — Probabilism?*, *BJPS*, 2008.
- [19] C. Hempel, *Deductive-Nomological vs. Statistical Explanation*, 1962.
- [20] W. James, *The will to believe, The New World*, 1896.
- [21] 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.
- [22] ———, *The Foundations of Causal Decision Theory*, CUP, 1999.
- [23] ———, *A Nonpragmatic Vindication of Probabilism, Philosophy of Science*, 1998.
- [24] M. Kaplan, *Coming to Terms with our Human Fallibility*, *PPR*, 2013.
- [25] P. Klein, *The virtues of inconsistency, The Monist*, 1985.
- [26] N. Kolodny, *How Does Coherence Matter?*, *Proc. of the Aristotelian Society*, 2007.
- [27] N. Kolodny and J. MacFarlane, *Ifs and oughts, Journal of Philosophy*, 2010.
- [28] J. Konek and B. Levinstein, *The Foundations of Epistemic Decision Theory*, 2014.
- [29] H. Leitgeb and R. Pettigrew, *An Objective Justification of Bayesianism I & II, Philosophy of Science*, 2010.
- [30] I. Levi, *Gambling with Truth*, 1967.
- [31] J. MacFarlane, *Assessment Sensitivity: Relative Truth and its Applications*, OUP, 2014.
- [32] P. Maher, *Betting on Theories*, 1993.
- [33] D. Nelkin, *The lottery paradox, knowledge, and rationality, Phil. Review*, 2000.
- [34] G. Oddie, *Conditionalization, cogency, and cognitive value, BJPS*, 1997.
- [35] D. Parfit, *What we together do*, unpublished manuscript, 1988.
- [36] ———, *Epistemic Utility Arguments for Probabilism, SEP*, 2011.
- [37] J. Pollock, *Nomic probability and the foundations of induction*, OUP, 1990.
- [38] J. Predd, R. Seringer, E. Loeb, D. Osherson, H.V. Poor and S., Kulkarni, *Probabilistic coherence and proper scoring rules, IEEE Transactions*, 2009.
- [39] F. Ramsey, *Truth and Probability*, 1926.
- [40] S. Ryan, *The epistemic virtues of consistency, Synthese*, 1996.
- [41] M. Titelbaum, *Quitting Certainties*, OUP, 2013.

Dutant & Fitelson	Knowledge Centered Epistemic Utility Theory	14
-------------------	---	----