Learning & Value Change

J. Dmitri Gallow

Modality & Method Workshop · Center for Formal Epistemology, CMU · June 10, 2017

1 Introduction

1. Daniel's beliefs are irrational but true; Melissa's beliefs are rational but false. So rational belief is not true belief. But isn't there still *some* connection between rationality and truth?

2. A tempting thought: Daniel's beliefs were true, but they were *likely* to be false.

3. The accuracy-firster attempts to vindicate this tempting thought.

   (a) If the accuracy-firster is right, then Daniel should have expected his beliefs to be less accurate than other beliefs he could have adopted instead.

4. The accuracy-first project is to *derive* all evidential norms from:

   (a) The axiological claim that (properly measured) accuracy is the sole epistemic good; together with

   (b) Consequentialist deontic norms like 'it is rational to maximize expected epistemic value' and 'opinions which are epistemic-value-dominated are irrational'.

5. The point of today's talk:

   (a) Existing accuracy-first approaches to rational learning presuppose evidential norms which are not explained in terms of the single-minded pursuit of accuracy.

   (b) Alternatives are needed.

   (c) I have one to offer.

2 Bayesianism

6. At a time t, your opinions are representable with a *credal state* $< \mathscr{W}, \mathscr{A}, c_t >$, where

   (a) $\mathscr{W} = \{w_1, w_2, \dots w_N\}$ is a finite set of doxastically possible worlds;

   (b) $\mathscr{A} \subseteq \wp(\mathscr{W})$ is a set of propositions; and

   (c) $c_t : \mathscr{A} \to [0, 1]$ is your time t credence function which represents the stength of your belief in all propositions in $\mathscr{A}$.

7. The Bayesian account of rational learning: you should be a probabilistic conditionalizer.

   Probabilism
   At all times t, c_t should be a probability function.

   Conditionalization
   There should be some ur-prior credence function c such that, for all times t and all $A, E \in \mathscr{A}$ such that E could be your total evidence at t,

   $$c_{t,E}(A) = c(A \mid E)$$

   (a) '$c_{t,E}$' is the credence function you are disposed to adopt, at t, upon receiving the total evidence E.

3 Epistemic Value

8. '$\mathcal{V}(c, w)$' is the epistemic value of holding the credence function c at world w.

   (a) Accuracy-first epistemology claims that $\mathcal{V}(c, w)$ is entirely a function of the *accuracy* of c at w.

9. '$\mathcal{V}_c(c^*)$' is how epistemically valuable the credence function c^* is, from the standpoint of the credence function c.

   (a) Leitgeb & Pettigrew (2010): if your credence is a probability, p, then, for all c,

$$\mathcal{V}_p(c) \stackrel{!}{=} \sum_{w \in \mathscr{W}} \mathcal{V}(c, w) \cdot p(w)$$

   (b) This is the consequentialist deontic norm which says (epistemic) acts are choiceworthy to the extent that they maximize expected (epistemic) value.

> Propriety
> An epistemic value function $\mathcal{V}$ is *proper* iff, for every probability function p and every credence function $c \neq p$,
>
> $$\mathcal{V}_p(c) < \mathcal{V}_p(p)$$

10. Why Propriety? Two arguments.[1]

   (a) The first appeals to epistemic conservativism:

   P1. For any probability p, there is some evidence you could have that would make it permissible to have p as your credence function.

   P2. If another credence function c is at least as valuable as your own, then it is permissible to adopt c as your credence function, even without receiving any evidence

   P3. It is impermissible to change your credences without receiving evidence.

   C1. So, epistemic value must be proper

   (b) The second appeals to *immodesty* as a rational requirement:

   P1. For any probability p, there is some evidence you could have that would make it permissible to have p as your credence function.

   P4. Rationality requires you to think that your own credences are epistemically better than any other credences you could have held instead.

   C1. So, epistemic value must be proper

11. Predd et al. (2009) show that, if $\mathcal{V}$ is a proper measure of accuracy, then every non-probabilistic credence is *accuracy dominated* by some probabilistic credence, and no probabilistic credence function is accuracy dominated.

   (a) Thus, the above arguments, if successful, vindicate the rational norm Probabilism in terms of accuracy and accuracy alone.

4 Conditionalization & Accuracy

4.1 Take 1

12. Leitgeb & Pettigrew (2010): if p is your (probabilistic) credence, then you should be disposed, upon learning E, to adopt a new credence which

[1] See, *e.g.*, Oddie (1997), Joyce (2009), and Pettigrew (2011).

maximizes expected epistemic value in all possibilities consistent with E.

$$p_E \stackrel{!}{=} \arg\max_c \left\{ \sum_{w \in E} p(w) \cdot \mathcal{V}(c, w) \right\}$$

13. This norm, together with the following theorem,

Theorem 1 (Generalized from Leitgeb & Pettigrew (2010)). *If $\mathcal{V}$ is proper, then, for any probabilistic p and any E,*

$$\arg\max_c \left\{ \sum_{w \in E} p(w) \cdot \mathcal{V}(c, w) \right\} = p(- \mid E)$$

entails that, if $\mathcal{V}$ is a proper accuracy measure, then $p_E \stackrel{!}{=} p(- \mid E)$ (Conditionalization)

14. This affords the following argument for Conditionalization:

P5. Upon learning that E, you should be disposed to adopt a new credence which maximizes expected epistemic value in all possibilities consistent with E.

P6. Epistemic value is (properly measured) accuracy.

P7. Theorem 1.

C2. Upon learning that E, you should be disposed to conditionalize on E.

15. Note: P5 does not follow from, and in fact conflicts with, the norm to maximize expected epistemic value. Why should we accept this norm?

    (a) Leitgeb & Pettigrew appear to presuppose a 2-stage theory of rational learning.

        i. Stage 1: upon learning E, you eliminate all $\neg E$ worlds from $\mathscr{W}$.
        ii. Stage 2: you use your prior (no longer probabilistic) credences over the remaining worlds to pick a posterior which maximizes expected epistemic value.

    (b) The elimination of worlds at stage 1 either:

        i. relies upon an evidential norm like "do not treat a world as epistemically possible if it is incompatible with your evidence"; or
        ii. is treated as an brute and not rationally evaluable fact.

    (c) In the first case, we've failed to reduce all evidential norms to the pursuit of accuracy; in the second, we must deny that becoming certain that climate change is a hoax perpetrated by the Chinese after a snowfall is irrational.

4.2 Take 2

16. Leitgeb & Pettigrew (2010): if p is your (probabilistic) credence, then you should be disposed, upon learning E, to adopt a new credence which maximizes expected epistemic value *amongst those credence functions consistent with your evidence*.

$$p_E \stackrel{!}{=} \arg\max_{\substack{c : c(E)=1 \\ c(\neg E)=0}} \left\{ \sum_{w \in E} p(w) \cdot \mathcal{V}(c, w) \right\} \qquad (\star)$$

17. The solution to this maximization problem depends upon which proper accuracy measure you use. In the case of the quadratic, Leitgeb & Pettigrew (2010) show that

Theorem 2 (Leitgeb & Pettigrew (2010)). *If $\mathcal{V} = \mathcal{Q}$, then the solution to the maximization problem in ($\star$) is*

$$p(A \parallel E) \stackrel{def}{=} p(AE) + \frac{\|AE\|}{\|E\|} \cdot (1 - p(E))$$

18. Updating your degrees of belief from p to $p(- \parallel E)$ is not epistemically defensible (*cf.* Levinstein (2012)).

19. Levinstein: we should keep the norm ($\star$), but instead of the quadratic accuracy measure $\mathcal{Q}$, we should use the logarithmic $\mathcal{L}'$, where

$$\mathcal{L}'(c, w) \stackrel{\text{def}}{=} \ln [c(w)]$$

Theorem 3 (Levinstein (2012)). *If $\mathcal{V} = \mathcal{L}'$, then the solution to the maximization problem in* ($\star$) *is* $p(- \mid E)$.

(a) $\mathcal{L}'$ may vindicate conditionalization, but it cannot vindicate probabilism, since every probabilistic credence function is $\mathcal{L}'$-dominated.

(b) What *is* proper is $\mathcal{L}(c, w_i) = \sum_{w_j \in \mathscr{W}} \ln [| \, (1 - \delta_{ij}) - c(w_j) \, |]$.

(c) But, if $\mathcal{V} = \mathcal{L}$, then the solution to the maximization problem in ($\star$) will only be $p(- \mid E)$ if your prior p was the uniform distribution or $p(E) = 1$.

(d) The solution to ($\star$) when $\mathcal{V} = \mathcal{L}$ is no more epistemically defensible than updating your beliefs to $p(- \parallel E)$.

5 Epistemic Value Change

20. Leitgeb & Pettigrew give a model of rational belief with three components:
    (a) a credal state;
    (b) an epistemic value function; and
    (c) a dynamical law—rational credence travels in the direction of highest expected accuracy
21. If the epistemic value function is proper, then this model will always be in equilibrium.
22. So, if there is to be a rational *change* in belief, then there must be an exogenous change to one of these three components.
23. The accuracy-firster shouldn't say that it is an exogenous change to the credal state.
    (a) If we say there's an exogenous change to the credal state, then either the change is rationally evaluable or it is not.
        i. If it is, then there are rational norms which haven't been vindicated in terms of the rational pursuit of accuracy and accuracy alone.
        ii. If it's not, then it's not irrational to become certain that climate change is a Chinese-perpetrated hoax after a snowfall.
24. The accuracy-firster should also not say that it is an exogenous change to the credal dynamics.
    (a) To say this is to abandon the accuracy-first project and insist that, sometimes, rationality means adopting credences which are expected to be less accurate than the ones you currently hold.
25. This leaves one option remaining: an exogenous change to the epistemic value function.
    (a) In general, expected accuracy maximizers think that learning experiences can change the degree to which you take accuracy at certain worlds into account.
    (b) On the standard way of thinking about things, this happens because you *weight* accuracy at w by $p(w)$, and learning can change these weights.
    (c) A proposal: reverse the order of explanation. Not: you rationally stop valuing accuracy at $\neg E$ possibilities *because* you are certain of E. Rather: you become certain of E *because* you rationally stop valuing accuracy at $\neg E$ possibilities.

5.1 Conditionalization

26. Suppose that learning E rationalizes not caring at all about accuracy at worlds $w \notin E$.
27. Then, if '$\mathcal{V}^E$' is the epistemic value function which it is rational to hold after learning that E, we can say that

$$\mathcal{V}^E(c, w) = \begin{cases} \mathcal{V}(c, w) & \text{if } w \in E \\ \kappa_w & \text{if } w \notin E \end{cases}$$

    (a) κ_w is any constant.
    (b) So, at $w \notin E$, you value accurate credences just as highly as you value inaccurate credences; which is to say: you don't value accuracy at $w \notin E$.

(c) Then, when you have a (probabilistic) prior credence p and you learn that E, you will attempt to maximize

$$\begin{aligned}\mathcal{V}_p^E(c) &= \sum_{w\in\mathscr{W}} p(w)\cdot\mathcal{V}^E(c,w)\\ &= \sum_{w\in E} p(w)\cdot\mathcal{V}(c,w) \;+\; \sum_{w\notin E} p(w)\cdot\kappa_w\end{aligned}$$

(d) Since $\sum_{w\notin E} p(w)\kappa_w$ is just a constant, the above will be maximized when and only when

$$\sum_{w\in E} p(w)\cdot\mathcal{V}(c,w)$$

is maximized.

(e) And **Theorem 1** assures us that, so long as $\mathcal{V}$ is proper, the c which maximizes this will be $p(- \mid E)$.

28. Note: the updated $\mathcal{V}^E$ will no longer be proper.

29. This affords us the following vindication of Conditionalization:

P7. Theorem 1

P8. Your ur-prior epistemic value function should be a proper measure of accuracy.

P9. Upon learning E, it is rationally required to update your epistemic value function by not valuing accuracy at non-E possibilities.

C2. Upon learning E, you should be disposed to conditionalize on E.

5.2 Propriety

30. $\mathcal{V}^E$ is not a proper measure of accuracy; but neither do the arguments for propriety from §3 give us any reason to reject it. For, on the proposed account, there is nothing stopping us from accepting all the premises of those arguments but rejecting their conclusions. So the arguments are invalid.

31. Those arguments presuppose that rational epistemic values cannot change; so they give us no reason to worry about $\mathcal{V}^E$ as an (updated) epistemic value function.

32. Neither do they give us any reason for thinking that the *ur-prior* value function $\mathcal{V}$ should be proper.

33. There are arguments for holding that that the quadratic measure is the uniquely best measure of accuracy (*cf.* Pettigrew, 2016); these arguments aren't shown to be invalid on the current approach, and could serve its needs.

6 In Summation

34. On this proposal, rational belief is belief formed in the rational pursuit of rationally-valued accuracy.

35. Daniel is irrational because he has adopted beliefs which he should expect to be less accurate than other beliefs he could have held instead.

    (a) Daniel is either not valuing accuracy rationally or not pursuing accuracy rationally.

36. Melissa is more rational because she has adopted the beliefs she should have expected to be most accurate.

    (a) Melissa is both rationally valuing and rationally pursuing accuracy.

References

Joyce, James M. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In *Degrees of Belief*, F. Huber & C. Schmidt-Petri, editors, 263–97. Springer, Dordrecht. [2]

Leitgeb, Hannes & Richard Pettigrew. 2010. "An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy." *Philosophy of Science*, vol. 77 (2): 236–272. [2], [3], [4] Levinstein, Benjamin Anders. 2012. "Leitgeb and Pettigrew on Accuracy and Updating." *Philosophy of Science*, vol. 79 (3): 413–424. [3], [4]

Oddie, Graham. 1997. "Conditionalization, Cogency, and Cognitive Value." *British Journal for the Philosophy of Science*, vol. 48: 533–41. [2]

Pettigrew, Richard. 2011. "An Improper Introduction to Epistemic Utility Theory." In *EPSA Philosophy of Science: Amsterdam 2009*, Henk W. de Regt, Stephan Hartmann & Samir Okasha, editors. Springer. [2]

—. 2016. *Accuracy and the Laws of Credence*. Oxford University Press, Oxford. [5]

Predd, Joel, Robert Seiringer, Elliot H. Lieb, Daniel N. Osherson, H. Vincent Poor & Sanjeev R. Kulkarni. 2009. "Probabilistic Coherence and Proper Scoring Rules." *IEEE Transations on Information Theory*, vol. 55 (10): 4786–4792. [2]

A Technicalities

Theorem 1 *If $\mathcal{V}$ is proper, then, for any probabilistic p and any E,*

$$\arg\max_c \left\{ \sum_{w \in E} p(w) \cdot \mathcal{V}(c, w) \right\} = p(- \mid E)$$

Proof. $p(- \mid E)$ is a probability; so, if $\mathcal{V}$ is proper, then $c = p(- \mid E)$ maximizes

$$\sum_{w \in \mathscr{W}} p(w \mid E) \cdot \mathcal{V}(c, w) = \sum_{w \in E} p(w \mid E) \cdot \mathcal{V}(c, w)$$

And, if $c = p(- \mid E)$ maximizes this function, it will continue to do so if we multiply it by the constant $p(E)$, so $c = p(- \mid E)$ maximizes

$$p(E) \cdot \sum_{w \in E} p(w \mid E) \cdot \mathcal{V}(c, w) \;=\; \sum_{w \in E} p(w) \cdot \mathcal{V}(c, w)$$

□