Relevant Alternatives and Missed Clues: Redux

Preprint of a paper forthcoming in The Journal of Philosophy

Peter Hawke

peterhawke@ln.edu.hk *

Department of Philosophy, Lingnan University
Hong Kong Catastrophic Risk Centre, Lingnan University

February 10, 2023

Abstract

I construe Relevant Alternatives Theory (RAT) as an abstract combination of anti-skepticism and epistemic modesty, then re-evaluate the challenge posed to it by the *missed clue counter-examples* of Schaffer [2001]. The import of this challenge has been underestimated, as Schaffer's specific argument invites distracting objections. I offer a novel formalization of RAT, accommodating a suitably wide class of concrete theories of knowledge. Then, I introduce *abstract missed clue cases* and prove that every RA theory, as formalized, admits such a case. This yields an argument - in Schaffer's spirit - that resists easy dismissal.

Keywords: relevant alternatives theory; epistemic logic; fallibilism; missed

*Acknowledgements. This paper was largely written while I was a member of the project 'The Logic of Conceivability', funded by the European Research Council (ERC CoG), grant number 681404. Thanks to Francesco Berto and the LoC team. For helpful comments, thanks to Julien Dutant, Ben Lennertz, my anonymous referees, and audiences at the following venues: the 2018 meeting of the European Epistemology Network in Amsterdam, the Netherlands in June 2018; the 2019 Pacific Division Meeting of the American Philosophical Association in Vancouver, Canada in April 2019; the epistemology seminar at Arché at the University of St. Andrews, Scotland, in May 2019; the 2019 edition of the Formal Epistemology Workshop (FEW2019) in Turin, Italy, in June 2019.

clues; epistemic closure.

1 Introduction

Relevant Alternatives Theory (RAT) holds enduring appeal for epistemologists.[1] An RA theorist denies that 'a knows that φ at time t' is true only if a is positioned at t to eliminate every alternative to φ. Instead, she claims, a need only be positioned to eliminate the *relevant* alternatives to φ. 'Relevance' is here intended as a non-evidential constraint on the counter-possibilities.[2] So construed, RAT covers a vast family of concrete theories, ranging over explications of (i) alternative, (ii) the criteria of relevance, (iii) evidence and (iv) elimination. In particular, RA theorists flesh out (ii) diversely, usually as some combination of: psychological salience to either the attributor or subject of the knowledge ascription;[3] resemblance to actuality;[4] presupposition;[5] conversational relevance to the question or topic under discussion;[6] objective probability;[7] compatibility with the agent's beliefs;[8] non-evidential (e.g. pragmatic) rationality;[9] or practical stakes.[10]

Example: contextualists and subject-sensitive invariantists agree that the truth of 'a knows that φ' can vary as epistemic standards shift (though a's evidence doesn't). They disagree on whether the standards are partly determined by the attributor's discourse context (e.g. court vs. pub) or en-

[1]See Dretske [1970], Goldman [1976], Stine [1976], Cohen [1988], Lewis [1996], Kratzer [2012, Ch.6], Pritchard [2012, Part 2], Lawlor [2013], Bradley [2014], Hawke [2017], [Drucker, 2020, sect.4].

[2]That is: the relevance of a counter-possibility is independent of its compatibility with the agent's evidence. Some RA theorists nominally violate this dictum e.g. Cohen [1988, pg. 103] describes his 'internal' criterion of relevance in terms of (absent) evidential support. This is terminological: we take such criteria as bearing on *elimination* (by the evidence).

[3]See [Stine, 1976, pg. 256], Cohen [1988], Lewis [1996], [Pritchard, 2012, Sect. 6].

[4]See Dretske [1981], Lewis [1996], Heller [1989], [Pritchard, 2012, Sect. 6].

[5]Dretske [1970] suggests semantic presupposition. Blome-Tillmann [2014] considers pragmatic presupposition, embellishing [Goldman, 1976, pp. 776-777] and [Stine, 1976, pp. 255-256].

[6]See Schaffer [2005], Yablo [2014], Hawke [2016a].

[7]See [Cohen, 1988, pg. 102].

[8]See the 'rule of belief' in Lewis [1996].

[9]See [Stine, 1976, pg. 252], tracing to Goldman [1976]. Also: [Cohen, 1988, Sect. V], Wright [2014].

[10]See Hawthorne [2004] and Stanley [2005].

tirely by subject a's circumstances (e.g. her practical stakes). How to think about epistemic standards? A natural construal: an epistemic standard is a class of possibilities ('relevant alternatives'), with one stricter than another if it includes more possibilities.

Another example: an entitlement theory, following Wright [2014], posits that knowledge is grounded in a class of *hinge propositions* (e.g. that the senses are generally reliable). It's claimed that (i) rational inquiry is possible only if these hold, and so (ii) the possibility that they are false is rightly ignored ('irrelevant') for epistemic evaluations, despite its compatibility with any empirical evidence.

RAT's appeal lies in its joint accommodation of *epistemic modesty* and *anti-skepticism*, via epistemic fallibilism. Modesty counsels a theorist not to exaggerate our ordinary epistemic powers. Isn't it wishful thinking to take ordinary empirical evidence as ruling out radical skeptical possibilities? Such possibilities, one might think, are perfectly compatible with one's total basic empirical evidence. Similarly, it might seem like wishful thinking to take 'pure reflection' as a path to ruling out radical skeptical possibilities. A cognitively ideal brain-in-vat can, presumably, know *every* truth that is accessible by pure reflection: $2 + 2 = 4$, all bachelors are males etc. But, trivially, she can't know she isn't a brain-in-vat.

Meanwhile, fallibilism posits that the epistemic status expressed by ordinary knowledge ascriptions is compatible with certain possibilities of error. Cue RAT: though ordinary evidence is consistent with recalcitrant skeptical possibilities, it yields mundane knowledge. Skeptical possibilities are normally irrelevant.

However, Schaffer [2001] argues that RAT has a problem: *missed clue cases*, he alleges, serve as counter-examples to any well-motivated RA theory.

Call the following sort of situation a *Schafferian missed clue case* (SMC). Agent a knows proposition C in a mundane context where C is a 'clue' that P holds. By a 'clue', we here mean that C *objectively indicates* P: the possibility that C holds without P is abnormal, witnessed only in possible worlds 'far removed' from actuality. However, also by stipulation, a has no (subjective) reason to believe P. In particular, she has no reason to believe that if C then P.

Intuitively, a isn't positioned to know P: if she believes P on the basis of C, then she not only lacks knowledge of P, but is irrational. The clue is missed: C is discerned, but its significance for P is not. Schaffer [2001] claims that the RA theorist must here contradict intuition. As the possibility that C without P is distant, it's akin to a skeptical one. So RAT must count it as irrelevant. The only other alternatives to P are those where neither C nor P holds, but, where relevant, such possibilities are eliminated by a's evidence: if not, a would not know C after all. So, RAT, to its discredit, entails that a knows P.

Here is a concrete SMC:

> *Mica's missed clue.* Mica is browsing through The Bird Almanac, Italian edition. She doesn't read Italian but is enjoying the photos. She sees a photo of a bird with *red plumage* (the clue). Mica recognizes it as a canary. She knows that canaries are either wild or domestic. But she knows no general facts about their plumage. In particular, she doesn't know that wild canaries always have *yellowish-green* plumage while domestic canaries have diverse plumage - sometimes red.

Mica isn't positioned to know the depicted canary is domestic. Compare Professor Byrd, renowned ornithologist. If Byrd sees the photo, she knows the depicted canary is domestic. Byrd can *see* that the canary is domestic. Whatever Mica can see (e.g. that the plumage is red), she cannot see that the canary is domestic. According to Schaffer, RA theorists must disagree. Mica's perceptual evidence eliminates the possibility of a yellowish-green wild canary. As a red-plumed wild canary is a distant possibility, it's a skeptical one that counts as irrelevant in mundane contexts (even Byrd's evidence is consistent with freak possibilities like a mutant red-plumed wild canary). So much for the alternatives: by RA lights, Mica knows the canary is domestic.

Let *red* stand for 'the canary is red-plumed' and *domestic* for 'the canary is domestic'. The Schafferian objection from *Mica's missed clue* is explicated as:

P1. According to any RA theory (worth taking seriously), every relevant alternative to *domestic* is, in Mica's situation, a relevant alternative to *red*, despite her having no reason to connect them.

C1. So, according to any such RA theory, Mica can know *domestic* in virtue of knowing *red*.

P2. But if Mica has no reason to connect *red* and *domestic*, she is ignorant of any connection between them.

P3. Mica can know *domestic* in virtue of knowing *red* only if she knows of a connection between *red* and *domestic* (for otherwise she could not infer *domestic* from *red* in a knowledge-preserving way).

C2. So, Mica doesn't know *domestic* in virtue of knowing *red* (indeed, she has no grounds whatsoever for believing *domestic*).

C3. So, by C1 and C2, RAT is false.

Existing responses are predictable. Brueckner [2003] argues that the core problem isn't RAT specific, but applies to any theory of knowledge that undercuts reasonable belief. (Exhibit: simple process reliablism.) He recommends that RA theories complement the condition that every relevant alternative be eliminated with an independent reasonable belief condition, challenging the inference from P1 to C1. Black [2003], meanwhile, proposes that if reasonable belief is appropriately incorporated into the criteria of relevance (per the 'Rule of Belief' in [Lewis, 1996]), then 'distant' possibilities needn't be treated as equals in mundane contexts: that wild canaries could have *ordinarily* been red may be relevant despite the irrelevance of the possibility of a *mutant* wild canary. This challenges P1.

However, appeals to reasonable belief are worrying. Do modesty and anti-skepticism recommend an RA approach to reasonable belief? If so, Brueckner and Black merely kick the can down the road (nothing prevents one from constructing analogous missed clue counter-examples for an RA theory of reasonable belief). If not, RAT's motivation is dulled: skeptical worries lose their sting if ordinary agents have good reasons to deny radical skepticism. Furthermore, if reasonable belief depends on evidence, integrating it into the criteria of relevance violates the RA dictum that relevance be non-evidential.

But more importantly, the dialectic is puzzling. Schaffer [2001] aims for a universal source of counter-examples to RAT. Yet his construal of RAT is markedly narrow. The trouble from SMCs hinges on commitment to a specific criterion of relevance: if P objectively indicates Q, any possibility that P holds without Q is irrelevant. Further, Schaffer takes the elimination of every relevant alternative as *sufficient* for a true knowledge ascription by RA lights. Further, if a relevant alternative to Q is a not-P possibility, Schaffer takes that same possibility to be a relevant alternative to P by RA lights. This goes well beyond RAT's basic commitments. Schaffer invites an all-too-easy response: his considerations imperil simple RA theories, but leave nuanced ones untouched. Putting aside specifics, an RA theorist might join Brueckner in denying sufficiency or join Black in eschewing blinkered criteria of relevance.

All this obscures the deeper import of missed clue cases. It is better detected when focusing on the *basic structure* of SMCs, abstracting from specific criteria of relevance and other narrow constraints on RAT. Or so I will argue. The Schafferian argument from *Mica's missed clue* is indeed flawed (e.g. P1 is false). But the core considerations that drive it can be salvaged and generalized. In outline, I develop and study the claim that for every RA theory `T` (worth taking seriously), there exists scenario s (concerning *idealized* agent a) and aposteriori claims p and q such that the following reasoning is sound:

P1′. According to `T`, every uneliminated relevant alternative to q in s is a relevant alternative to p, though neither a's empirical evidence nor mere reflection positions her to know that p doesn't hold without q.

C1′. So, according to `T`, if a were to learn p (i.e., supplement her evidence in s with knowledge of p), she would know q.

P2′. If neither her empirical evidence nor mere reflection positions a to know that p doesn't hold without q, she is ignorant of a connection between p and q.

P3′. An agent would know q after learning p only if positioned to know of a connection between p and q.

C2′. So, a would not, in s, know q after learning p.

C3′. So, by C1′ and C2′, $\mathtt{T}$ is false.

This paper has two main goals. First, to precisify the above reasoning. Roughly, ‘idealized’ is taken as ‘logically ideal, relative to relevance constraints’. An RA theory is ‘worth taking seriously’ if it meets the basic goals of modesty and anti-skepticism. Knowing of ‘a connection’ between p and q is understood as knowing $p \supset q$ ($\supset$ is the material conditional).

The second goal is to identify plausible support for P1′-P3′ in full generality: a theorem yields P1′/C1′, while P2′ and P3′ are supported respectively by highly intuitive principles:

> ***Minimal Empiricism.*** If a claim cannot be known through reflection alone, then knowing it requires empirical evidence.
>
> ***Inferential Anti-dogmatism.*** If for all you know, p and not-q, then you cannot come to know q just by learning p and inferring q from p and things you already knew (cf. [Dorr et al., 2014, Sect.6]).

(These comport with the *perceptual dogmatism* of Pryor [2000] – see §3.2.)

We proceed as follows. §2 renders RAT with formal precision, studiously avoiding commitment to specific criteria of relevance or sufficient conditions for knowledge ascription. Formal techniques are apt: RAT is an *inherently abstract and structural* perspective on knowledge and knowledge ascription, yet is rarely discussed with appropriate generality and precision.[11]

Next, we put our formalism to philosophical work. (Proofs are largely in Appendix A.) §3 returns to SMCs and locates a core abstract structure, introducing *abstract missed clue cases* (AMCs). §3.1 proves that every RA theory admits an AMC. §3.2 shows that AMCs violate a formal closure principle – ‘basing closure’ – based on two intuitive principles: *No Cheap Knowledge* (**NoCK**) and *Minimal Connection* (**MinC**). The former precisifies *minimal empiricism*, the latter *inferential anti-dogmatism*. §3.3-3.5 bolsters their intuitive appeal. All this yields (§3.6) a compact and abstract argument in the *spirit* of Schaffer’s. In this sense, Schaffer [2001] is vindicated: suitably generalized, missed clue cases bring universal trouble for RAT.

[11] I build on Holliday [2015b], Hawke [2016a], Hawke [2016b].

§4 illuminates connections between the missed clue challenge, closure under deduction, and the paradox of the criterion. §5 explores avenues for response.

2 RAT with precision

2.1 *RAT clarified*

The core *RA thesis* says two things: first, that 'a knows that φ' is true, in discourse context c, only if (1) holds; second, that (2) is false.

(1) ***RA condition.*** E eliminates every relevant alternative to $\ulcorner\varphi\urcorner$ in c, where E is a's evidence.

(2) ***Universalism.*** For all φ and c, every alternative to $\ulcorner\varphi\urcorner$ is relevant in c.

We use P, Q, R, E for arbitrary propositions and φ, ψ, χ for arbitrary indicative sentences (we generally drop corner quotes).

For generality, we frame the RA thesis as semantic. Hence, we take ordinary knowledge attributions as data to accommodate (e.g. with a well-motivated error theory). We assume mundane knowledge attributions are *generally* unambiguous and accurate. A semantic lens is then harmless baggage for RA theories where relevance isn't discourse-sensitive. For expository smoothness, we often talk about what an agent knows when we should talk about which knowledge attributions are true in context.

The *alternatives to φ* are the possible worlds at which φ is false. Thus, for our purposes, the content of φ can be identified with its truth conditions and so a proposition identified with the set of possible worlds at which it is true. A world w is *eliminated* by proposition P when w isn't a member of P; otherwise w is *compatible* with P. (Our arguments transfer to *probabilistic* elimination: see §5.1.) P and Q are *inconsistent* if P eliminates every world in Q.

We are largely neutral on the 'possible worlds' at issue: metaphysical, conceptual, centered, or whatever. However, our worlds are *basic empirical possibilities*: eliminated only upon receipt of appropriate empirical information. Thus we class a proposition (or sentence expressing it) as *empirical* when it is compatible with some but not all possible worlds. φ is *strongly apriori* exactly when there are no empirical alternatives to φ: every possible world is a φ-world. Compare *weak apriority*: φ has no relevant alternatives in context.

Strongly apriori sentences are thus knowable independently of the available empirical information or relevant possibilities: a logically ideal agent knows these truths in every context. All standard examples of the apriori seem thus classified: mathematical truths, analytic truths. Such knowledge requires only the (alleged) tools of pure reflection: conceptual resources, computational resources, semantic competence, etc. In contrast, apriority without strong apriority is intuitively puzzling: if such sentences exist, they are knowable independently of both empirical information *and* pure reflection. We (suggestively) label strongly apriori sentences as *reflective* and sentences that are both apriori and empirical (if such there be) as *cheap*.

Some criteria of relevance demand that the 'relevant alternatives to φ' be *propositions* inconsistent with φ: propositions are presupposed by a conversation, believed for pragmatic reasons, and objectively probable. However, if the relevant alternatives to φ are primarily a set of propositions, one easily generates a set of 'relevant alternative worlds' by taking their union. If P is eliminated by evidence E just in case no world is a member of both P and E, then, abstractly, it makes no difference whether (1) is understood in terms of eliminating the 'relevant alternative propositions' or the (generated) 'relevant alternative worlds'.

Condition (1) is an 'evidence-centric' construal of RAT (cf. Dretske [1981], Lewis [1996]). Another approach frames the RA condition in terms of possibilities the *agent* can eliminate.[12] This is consistent with our discussion: we assume an agent eliminates an alternative only if her evidence eliminates it.

We understand 'evidence' in a minimal way that illuminates RAT's basic motivations: we take the agent's evidence E to be the *total empirical information* that she has *directly received* via her senses:

[12]See [Goldman, 1976, pg. 774], [Pritchard, 2012, p.65] and §1.

Empirical Information. E is a true empirical proposition.

Received. The agent knows E.

Theorists that countenance non-propositional evidence may take our notion of 'evidence' as derivative: the total empirical information that is immediately knowable purely in virtue of the agent's total evidence. Or E can be given a negative formulation: E represents the agent's total knowledge state (given her actual evidence) were *every* alternative relevant for *every* claim. Anyway, RA theorists characteristically assume that E is ordinarily *modest*:

Modesty. E is compatible with skeptical possibilities.

RA theorists differ on whether (1) is *sufficient* for the truth of a knowledge ascription in context. Lewis [1996] embraces sufficiency. Alternatively, (1) can be developed as a mere component of the meaning of knowledge ascriptions. We maintain neutrality: we theorize about (1) in isolation, as a purported characterization of the knowledge of a *logically ideal* agent.

RA-knowability. 'φ is knowable relative to evidence E' is true in context c iff E eliminates every relevant alternative to $\ulcorner\varphi\urcorner$ in c. (And (2) is false.)

If a's evidence E makes 'φ is knowable relative to E' true, we say that a is *positioned to know* φ. So construed, (1) bears on what knowledge can be ascribed *in principle* given E. The quality of the agent's information is emphasized over her contingent psychology: cognitive biases, computational limitations, memory failure, conceptual limitations. More carefully: our knowability ascriptions are intended to reflect what agent a with evidence E *would* know *were* she ideal along every cognitive dimension that is independent of the criteria of relevance and she *knows* she is ideal in this sense (e.g., if facts about *attention* bear on relevance, our idealization doesn't bear on attention).

2.2 *RAT formalized*

We now develop a framework for expressing our coming abstract argumentation. In particular, we propose a mathematically precise definition of ‘RA theory (that can be taken seriously)’, exploiting the basic motivation for an RA theory: modesty plus anti-skepticism (§1). It follows immediately that an RA theory, so defined, entails the core RA thesis (§2.1).[13]

Our formal language $\mathcal{L}$ includes: atomic proposition letters including distinguished atom e (p,q,r denote arbitrary atoms); boolean connectives; a one-place necessity modality $\blacksquare$; a one-place apriority operator A; and a two-place conditional knowability operator $\Rightarrow$ (we will define unconditional knowability derivatively). Think of atoms as mundane predications, except: e expresses the agent’s evidence at the actual world. $\blacksquare\varphi$ is intended to express *strong apriority* (‘epistemic necessity’, ‘knowable by reason alone’). $A\varphi$ expresses that φ is knowable apriori, i.e., without empirical evidence. Read the conditional $\varphi \Rightarrow \psi$ subjunctively: ‘if the agent’s evidence were supplemented with knowledge of φ, then ψ would be knowable’. Loosely: ‘knowing φ assures that ψ is knowable (relative to the evidence)’.

A *basic circumstance of evaluation* is a tuple $\langle W, @, \mathsf{E}\rangle$. W is the class of all possible worlds. @ is the actual world. E is a function that takes a world w and returns a non-empty proposition $E_w \subseteq W$: the agent’s evidence at w. An *abstract epistemic frame* $\mathcal{F}$ is a basic circumstance of evaluation plus a *relevancy function* R. Given w and φ, R returns a set of worlds $R_w(\varphi)$: the relevant alternatives to φ at w.

A *discourse context* $\mathcal{C}$ (‘context’) is an enriched frame $\langle W, @, \mathsf{E}, \mathsf{R}, \vDash\rangle$ with $w \vDash \varphi$ indicating φ is true at w, according to recursive constraints (we write X^c for the complement of $X \subseteq W$, i.e., for $W \backslash X$):

$$w \vDash e \text{ iff } w \in E_@$$

$$w \vDash \neg\varphi \text{ iff } w \nvDash \varphi$$

$$w \vDash \varphi \wedge \psi \text{ iff } w \vDash \varphi \text{ and } w \vDash \psi$$

$$w \vDash \blacksquare\varphi \text{ iff } u \vDash \varphi \text{ for every } u \in W$$

$$w \vDash A\varphi \text{ iff } R_w(\varphi) = \varnothing$$

$$w \vDash \varphi \Rightarrow \psi \text{ iff } R_w(\psi) \subseteq R_w(\varphi) \cup E_w^c$$

[13] Set-theoretic notation: $\in$ denotes membership; $\subseteq$ denotes subset; $\cap$ and $\cup$ denote intersection and union; $\backslash$ denotes set minus; $\varnothing$ denotes the empty set.

Atom e expresses the agent's evidence at @. $\varphi \Rightarrow \psi$ is true when every relevant alternative to ψ is either a relevant alternative to φ or eliminated by E_w. Intuitively, if $p \Rightarrow q$ holds at w, supplementing E_w with enough evidence to eliminate every relevant alternative to p assures that *both* p and q are knowable.

A context $\mathcal{C}$ *satisfies* φ (written $\mathcal{C} \models \varphi$) iff $@ \models \varphi$. Notation: $[\varphi] = \{w \in W : w \models \varphi\}$ is the *truth set* of φ, relative to a salient context $\mathcal{C}$; $\varphi \vee \psi$ is shorthand for $\neg(\neg\varphi \wedge \neg\psi)$; $\varphi \supset \psi$ means $\neg(\varphi \wedge \neg\psi)$; $\varphi \equiv \psi$ means $(\varphi \supset \psi) \wedge (\psi \supset \varphi)$.

A *theory of knowability* $\mathtt{T}$ is a sub-class of contexts. If $\mathcal{C} \in \mathtt{T}$, we say that $\mathcal{C}$ is *admissible* relative to $\mathtt{T}$. Intuitively, a theory demarcates 'legitimate' contexts (via, for instance, constraints on the relevancy function, or structure of possible worlds). For example, a context is *infallibilist* iff $R_w(\varphi) = [\neg\varphi]$ for every w and φ. An *infallibilist theory* admits only infallibilist contexts. Thus, infallibilism makes R redundant: every alternative to φ is relevant, always.

Relative to theory $\mathtt{T}$, φ is a *logical truth* (written $\models_{\mathtt{T}} \varphi$) iff $\mathcal{C} \models \varphi$ for every $\mathcal{C} \in \mathtt{T}$. ψ is a *logical consequence* of $\varphi_1 \ldots \varphi_n$ (written $\varphi_1, \ldots, \varphi_n \models_{\mathtt{T}} \psi$) iff: for every $\mathcal{C} \in \mathtt{T}$, if $\mathcal{C} \models \varphi_i$ for every $i = 1 \ldots n$, then $\mathcal{C} \models \psi$.

We only work with *sensible* contexts. A context *has veridical evidence* iff evidence is always true: $w \in E_w$. A context *displays contrast*[14] iff the alternatives to φ are always $\neg\varphi$ worlds: $R_w(\varphi) \subseteq [\neg\varphi]$. A context *makes actuality relevant* iff: $w \models \neg\varphi$ implies $w \in R_w(\varphi)$. A context is *sensible* iff it has veridical evidence, displays contrast, and makes actuality relevant. A *sensible theory* admits only sensible contexts. Call the class of all sensible contexts the *universal theory* ($\mathtt{U}$).

Displaying contrast has met some resistance as a requirement on a theory of knowledge/knowability: see Holliday [2015b], Vogel [2014]. By embracing contrast, we stick for now to a straightforward construal of 'RA theory': φ is knowable exactly when the 'relevant alternatives to φ' are eliminated, where a 'relevant alternative to φ' must genuinely be an *alternative* to φ, i.e., a scenario where φ is false. (§5.2 explores the 'RA heresy' of dropping contrast.)

[14]Compare 'contrast/enough' in Holliday [2015b].

As contrast is displayed, we have $R_w(\top) = \varnothing$ in every sensible context, where $\top$ is some tautology (i.e., $[\top] = W$). We can now define two critical notions: we write $K\varphi$ as shorthand for $\top \Rightarrow \varphi$ and $E\varphi$ for $K\varphi \wedge \neg A\varphi$. Thus:

- $w \vDash K\varphi$ iff $R_w(\varphi) \subseteq E_w^c$, i.e., $R_w(\varphi) \cap E_w = \varnothing$
- $w \vDash E\varphi$ iff $R_w(\varphi) \subseteq E_w^c$ and $R_w(\varphi) \neq \varnothing$

$K\varphi$ is true when the evidence eliminates all relevant alternatives to φ. Read $K\varphi$ as 'φ is knowable relative to the agent's evidence'. We also write 'the agent is positioned to know φ' or 'φ is knowable on the available evidence'. $E\varphi$ expresses that φ is knowable aposteriori, i.e., knowable in virtue of the available empirical evidence. We also say: E_w renders φ knowable.

Proposition 1 (Crucial Universal Validities)

(i) $\vDash_{\mathtt{U}} Ke$

(ii) $K\varphi \vDash_{\mathtt{U}} \psi \Rightarrow \varphi$

(iii) $K\varphi \vDash_{\mathtt{U}} \varphi$

(iv) $E\varphi \vee A\varphi \vDash_{\mathtt{U}} K\varphi$

(v) $\varphi \Rightarrow \psi, \psi \Rightarrow \chi \vDash_{\mathtt{U}} \varphi \Rightarrow \chi$

(vi) $\blacksquare\varphi, \blacksquare(\varphi \supset \psi) \vDash_{\mathtt{U}} \blacksquare\psi$

(vii) $\vDash_{\mathtt{U}} \neg E(e \supset \varphi)$

(viii) $\vDash_{\mathtt{U}} \neg E((e \wedge \varphi) \supset \psi)$

Notably, the proofs for (i), (ii), (vii), and (viii) require contrast to be displayed.

We assume there is a unique *true* and *complete* sensible theory $\mathtt{R}$: *reality*.

Relative to context $\mathcal{C}$, we define: p is *empirical* iff $[p] \neq W$ and $[p] \neq \varnothing$.

Now, we define an *RA context* to be a sensible context $\mathcal{C}$ for which p exists such that: (i) $@ \vDash Ep$ and yet (ii) there is a world $\$_p$ such that $\$_p \vDash e \wedge \neg p$. Call $\$_p$ the *skeptical alternative to p* in $\mathcal{C}$. Note that RA contexts, so defined, showcase both modesty ($E_@$ doesn't eliminate a skeptical alternative) and anti-skepticism ($E_@$ renders empirical p knowable). We define an *RA theory* as a theory that admits only sensible contexts and at least one RA context.

We thus define RAT *purely formally*, as delivering basic RA goals: modesty plus the possibility of mundane empirical knowledge. No substantive assumption about the criteria of relevance is thereby entailed. So, an RA theory admits a context where $\$_p \in [\neg p]$ and $\$_p \notin R_@(p)$. No RA theory is infallibilist.

3 Missed clues redux

We have tools to develop our abstract Schafferian argument. §3.1 introduces AMCs as an abstraction of SMCs and proves that every RA theory admits an AMC. §3.2-§3.5 identify and bolster plausible general epistemic principles that are jointly incompatible with AMCs. §3.6 finalizes our abstract Schafferian argument and discusses an application.

3.1 *Abstract missed clue cases*

Mica's missed clue (§1) spells trouble for some RA theories. To extract the troublesome structure, we precisify the sort of theory Schaffer has in mind.

Consider *normality theory*: $R_w(\varphi) = N_w \cap [\neg\varphi]$, where N_w is the set of 'sufficiently normal' worlds, relative to w. We say a world is *nearby* to w iff it falls within N_w ($w \in N_w$ for all w). So, the relevant alternatives to φ are the nearby worlds at which φ is false. Let @ be actuality, including Mica's situation. Let $E_@$ be Mica's evidence after seeing the photo; r be that the bird is a red-plumed canary; d be that it's domestic. As it's an empirical matter whether only domestic canaries have red plumage, there are worlds where wild canaries have red plumage (i.e., $[r \wedge \neg d] \neq \varnothing$). But these worlds fall beyond $N_@$: the wild canary worlds in $N_@$ (i.e., $N_@ \cap [\neg d] \neq \varnothing$) aren't red plumage worlds (i.e, $R_@(r \supset d) = N_@ \cap [r \wedge \neg d] = \varnothing$ and $N_@ \cap [\neg d] \subseteq N_@ \cap [\neg r]$). Further, Mica's evidence eliminates all worlds in $N_@$ in which the bird doesn't have red plumage (i.e., $R_@(r) = N_@ \cap [\neg r] \subseteq E^c_@$). So normality theory judges that Mica is positioned to know the canary is domestic, as her evidence eliminates every relevant world where the bird is wild (i.e., $R_@(d) = N_@ \cap [\neg d] \subseteq E^c_@$).

In short, normality theory delivers some intuitively satisfying verdicts:

- $E_@$ renders r knowable (i.e., $R_@(r) \neq \emptyset$ and $R_@(r) \subseteq E^c_@$).
- $r \supset d$ is empirical, not reflective (i.e., $[r \wedge \neg d] \neq \emptyset$).
- $r \supset d$ isn't knowable aposteriori (as $R_@(r \supset d) = \emptyset$).

However, it also counter-intuitively delivers:

- Knowing r assures d is knowable (as $R_@(d) \subseteq R_@(r)$).

Symbolically:

(3) Er

(4) $\neg\blacksquare(r \supset d)$

(5) $\neg E(r \supset d)$

(6) $r \Rightarrow d$

Hence a *Schafferian worry*: surely learning exactly r doesn't position Mica to know d unless pure reasoning and/or her evidence excludes $r \wedge \neg d$? If so, no *correct* theory of knowledge would yield (3)-(6) simultaneously. Thus, (6) helpfully precisifies conclusion C1 (plus premise P1) in the Schafferian argument from *Mica's missed clue* (§1), while (4) and (5) helpfully precisify the implicit rationale for P2.

(3) is inessential and distracting. Consider *prior-Mica*, right before seeing the photo. Let $E_@$ now denote prior-Mica's evidence. Though (3) isn't true of prior-Mica, (4), (5), and (6) are: she isn't positioned to determine a connection between red plumage and domestic canaries, but *were* she to learn the canary has red plumage, she *would* be positioned to know that it's domestic. Intuition disagrees: prior-Mica is *not* so primed.

Hence the Schafferian worry departs from a worry about *closure of empirical knowledge under deduction*, though Mica's case also highlights it. If, as normality theory predicts, Mica is positioned to know d (via Er and $r \Rightarrow d$, per (iv) and (v) in Proposition 1), then, intuitively, she is positioned to know $\neg(r \wedge \neg d)$, via inference from d. But, again, normality theory predicts $E(r \supset d)$ is *false*. That (3) holds is crucial for *this* problem.

Other RA theories may avoid objectionable judgments about Mica in particular. But is a core worry showcased? We abstract. Define an *abstract* *missed clue case* (AMC) to be a context $\mathcal{C}$ for which φ and ψ exist such that (cf. (4)-(6)):

i. $\mathcal{C} \models \neg\blacksquare(\varphi \supset \psi)$

ii. $\mathcal{C} \models \neg E(\varphi \supset \psi)$

iii. $\mathcal{C} \models \varphi \Rightarrow \psi$

That is: $\varphi \supset \psi$ can't be known by reason alone and isn't knowable aposteriori, yet knowing φ would assure ψ is knowable in this context. In terms of relevancy sets:

i. $[\varphi \wedge \neg\psi] \neq \varnothing$;

ii. $R_@(\varphi \supset \psi) = \varnothing$ or $R_@(\varphi \supset \psi) \nsubseteq E^c_@$;

iii. $R_@(\psi) \subseteq R_@(\varphi) \cup E^c_@$.

An AMC $\mathcal{C}$ is *manifest* if $\mathcal{C} \models K\varphi$ (so, $\mathcal{C} \models K\psi$). It is *latent* if $\mathcal{C} \models \neg K\psi$ (so, $\mathcal{C} \models \neg K\varphi$). Thus, normality theory egregiously judges that Mica's situation is a manifest AMC and prior-Mica's situation is a latent AMC.

Our framework (§2.2) now entails key facts.

Proposition 2 *[No Escape]*

i. Every RA context $\mathcal{C}$ is a manifest AMC. So, every RA theory admits an AMC.

ii. Suppose that RA theory $\mathtt{T}$ *is such that:*

(a) $\models_{\mathtt{T}} (\varphi \wedge \psi) \Rightarrow \psi$

(b) $\mathtt{T}$ *admits $\mathcal{C}$ such that $@ \models p \wedge q \wedge (p \Rightarrow q) \wedge \neg Kq$ and $\mathcal{C}$ includes a world \$ such that $\$ \models e \wedge p \wedge \neg q$.*

Then $\mathtt{T}$ *admits a latent AMC.*

Proof:

i. By the definition of 'RA context', there exists p such that $\mathcal{C} \models Ep$. So, by (ii) and (iv) in Proposition 1, $\mathcal{C} \models e \Rightarrow p$. Also, as $\mathcal{C}$ includes a skeptical alternative to p, $\mathcal{C} \models \neg\blacksquare(e \supset p)$. By (i) and (vii) in Proposition 1,

$\mathcal{C} \models Ke$ and $\mathcal{C} \models \neg E(e \supset p)$. Altogether:

$$\mathcal{C} \models \neg\blacksquare(e \supset p) \wedge \neg E(e \supset p) \wedge (e \Rightarrow p) \wedge Ke$$

ii. By (b), $\mathcal{C} \not\models \blacksquare((e \wedge p) \supset q)$. By (viii) in Proposition 1, $\mathcal{C} \not\models E((e \wedge p) \supset q)$. By (a), $\mathcal{C} \models (e \wedge p) \Rightarrow p$. Since $\mathcal{C} \models p \Rightarrow q$, it follows by (v) in Proposition 1 that $\mathcal{C} \models (e \wedge p) \Rightarrow q$. Altogether:

$$\mathcal{C} \models \neg\blacksquare((e \wedge p) \supset q) \wedge \neg E((e \wedge p) \supset q) \wedge ((e \wedge p) \Rightarrow q) \wedge \neg Kq$$

For an RA theorist, the assumption in part ii is modest: (a) says that coming to know a conjunction assures its conjuncts are knowable; for (b), think of p as further empirical evidence that, if collected, would render q knowable, despite a skeptical possibility that p holds without q.

The proof of i suggests a prototype abstract Schafferian argument (cf. §1): every RA theory $\mathtt{T}$ admits $\mathcal{C}$ for which empirical p exists such that:

P1*. In $\mathcal{C}$, knowing e – the agent's total empirical evidence – renders p knowable, despite $e \supset p$ being neither strongly apriori nor knowable aposteriori.

P2*. According to reality, $e \supset p$ isn't knowable if $e \supset p$ is neither strongly apriori nor knowable aposteriori in that context.

P3*. According to reality, knowing e renders p knowable only if $e \supset p$ is knowable.

C1*. So, $\mathcal{C}$ isn't admitted by reality.

C2*. So, $\mathtt{T}$ isn't true.

It remains to clarify the support for P2* and P3*.

3.2 *Basing closure*

Admitting an AMC violates the conjunction of two intuitive principles:

(7) Minimal Connection (**MinC**): $\varphi \Rightarrow \psi \models K(\varphi \supset \psi)$

(8) No Cheap Knowledge (**NoCK**): $K\varphi, \neg\blacksquare\varphi \models E\varphi$

In terms of relevancy sets:

(9) **MinC**: if $R_@(\psi) \subseteq R_@(\varphi) \cup E^c_@$ then $R_@(\varphi \supset \psi) \subseteq E^c_@$

(10) **NoCK**: if $R_@(\varphi) \subseteq E^c_@$ and $[\neg\varphi] \neq \varnothing$ then $R_@(\varphi) \neq \varnothing$

MinC says: if supplementing one's evidence $E_@$ with knowledge of p assures that q is knowable, it must be knowable relative to just $E_@$ that either p isn't the case or q is the case.

NoCK says: if p is knowable relative to the available evidence and is empirical, then p is knowable *in virtue of* the evidence, not *apriori*. Empirical knowledge isn't cheap.[15]

Picturesquely, violating either **MinC** or **NoCK** allows for miraculous knowledge, appearing from nothing. The former allows that, sometimes, coming to know exactly p secures knowledge that q despite one's prior knowledge leaving it open that p holds without q. The latter allows a strong epistemic status to be achieved without expending any epistemic resources.

Now, Proposition 2 entails:

Corollary 1 *No RA theory validates both **MinC** and **NoCK**.*

To see this, note that (7) and (8) jointly entail a closure principle[16] for conditional knowability that is obviously inconsistent with the possibility of AMCs:

(11) Basing Closure (**BC**:) $\varphi \Rightarrow \psi, \neg\blacksquare(\varphi \supset \psi) \models E(\varphi \supset \psi)$

3.3 *No cheap knowledge*

NoCK's pre-theoretic appeal is plain. Per (10), violating **NoCK** admits weakly apriori knowledge: knowability claims that hold *purely* in virtue of facts about relevance. It is familiar enough to *presuppose* knowledge for conversational purposes; or behave *as if* one knows for purely pragmatic reasons; and so on, for whatever criterion of relevance. But it seems supremely odd to defend a knowledge claim as *true* on the mere basis of conversational presupposition, pragmatic reasons, or whatever.

[15] This echoes principle 'NoVK' in Holliday [2015b].

[16] A principle of the form: $O(\varphi_1, \ldots, \varphi_n)$ entails $O(\psi_1, \ldots, \psi_n)$ on condition $C(\varphi_1, \ldots, \varphi_n, \psi_1, \ldots, \psi_n)$, where O is an n-place operator.

Rejecting **NoCK** is a doubly costly response to AMCs. Consider an AMC:

$$\mathcal{C} \vDash \neg\blacksquare(p \supset q) \wedge \neg E(p \supset q) \wedge (p \Rightarrow q)$$

The **NoCK** denier explains $\mathcal{C}$ as a context where $p \supset q$ is known *cheaply*. Thus, the strategy claims, counter-intuitively, that an agent can come to know q on the basis of learning only p, despite lacking epistemic resources for establishing $p \supset q$. But it *also* claims, counter-intuitively, that despite this lack, the agent is *positioned to know* $p \supset q$. Apply the strategy to *Mica's missed clue*: it posits that prior-Mica can know the canary is domestic by observing its red plumage, despite lacking any evidence that bears on their connection. It *also* posits that prior-Mica is positioned to *know* that there is a connection, despite a dearth of evidence. A lot of bullets to bite.

3.4 *Minimal connection*

MinC's intuitive appeal is multi-faceted. As §4.1 will clarify, **MinC** is closely related to the closure of knowability under disjunction introduction. This is particularly evident when $\varphi \Rightarrow \psi$ holds in virtue of $K\psi$ holding, for then disjunction introduction yields $K(\neg\varphi \vee \psi)$. So, **MinC** shadows a principle that is widely seen as intuitive.

However, it can also be motivated *independently*, via a distinct set of intuitions about inferential knowledge. It minimizes distractions to focus on **MinC**'s *diachronic* import:

(12) $\neg K\psi, \varphi \Rightarrow \psi \vDash K(\varphi \supset \psi)$

Or equivalently:

(13) $\neg K\psi, \neg K(\varphi \supset \psi) \vDash \neg(\varphi \Rightarrow \psi)$

(12) is the semi-converse of the principle that knowability is closed under knowable implication:

(14) $\neg K\psi, K(\varphi \supset \psi) \vDash \varphi \Rightarrow \psi$

This says: if p isn't knowable on $E_@$, but it's knowable that p implies q, then supplementing $E_@$ with knowledge of p assures q is knowable.

(13) precisifies the principle of *inferential anti-dogmatism* [Dorr et al., 2014, Sect.6]: if for all you know, p and not-q, then you cannot come to know q just by learning p and inferring q from p and things you already knew. Bacon [2020] and Dorr et al. [2014] affirm its common sense appeal.[17]

Common sense judgment about concrete cases bolsters (12)/(13). Suppose I don't know whether Mary arrived late for work. What's more, for all I know, Mary arrived late and didn't receive a sanction ($l \wedge \neg s$). (She often sneaks in without the boss noticing.) So learning merely that she arrived late (l) wouldn't position me to infer/know she received a sanction (s).

Or suppose Jim and I are yet to see the new edition of *The Times*. Jim claims that if we were to there read that Clinton won the election (t), we would thereby know she won (w). But he rightly retracts this after I play up his fears about mainstream reporting, leading him to judge that we don't know that *The Times* hasn't *erroneously* reported that Clinton won ($t \wedge \neg w$).

Examples are easily multiplied, yielding two arguments for **MinC**. One may simply generalize. Or one may posit **MinC** as the best explanation for this robust pattern of judgment.

More principled arguments for (12) are afoot. Assume ideal agent a doesn't know ψ, but *would* if she were to supplement her current evidence e with knowledge that φ. Presumably, this is because coming to know φ would enable her to *infer* ψ. So a can deploy the following procedure: first, she supposes that φ holds; then she competently infers ψ; then she discharges her supposition, coming to know $\varphi \supset \psi$ by conditional introduction.

For a third argument for (12), reason by cases: suppose that φ supports ψ as a matter of logic (broadly construed). Then $\varphi \supset \psi$ is apriori, and (12) is immediate. Suppose otherwise: an argument with φ as sole premise and ψ as conclusion isn't cogent. Then if a learns φ but *isn't* positioned to know $\varphi \supset \psi$, she cannot competently come to know ψ by an inference using φ. (Perhaps a has *some* reason to believe $\varphi \supset \psi$, but positing knowledge based

[17]Curiously, both observe a trilemma between denying inferential anti-dogmatism, positing an implausible excess of prior knowledge to ordinary reasoners, and denying certain kinds of non-deductive knowledge (based, respectively, on chancy events and simple enumeration). The current paper dashes any hope that RA tools can help, as a structurally analogous conundrum emerges for RAT.

on beliefs that aren't knowledge embraces epistemic alchemy.) For, if a can't rule out $\varphi \wedge \neg\psi$, she isn't positioned to know a conditional of the form 'if φ, ψ', and knowledge of additional premises that establish such a conditional is *exactly* what cogent reasoning from φ to ψ requires.

3.5 *Dispelling confusions*

Some philosophers reject **MinC** or **NoCK**. An immodest foundationalist *à la* Chisholm [1966] thinks that one can know one has hands merely via knowing it *appears* one has hands, without prior knowledge that appearances aren't misleading. A reliabilist might say one can know q by inference from known p merely in virtue of the inferential mechanism being reliable; the epistemic status of $p \supset q$ is irrelevant. These plausibly translate into a rejection of **MinC**. Contextualists like Lewis [1996] explicitly allow contexts where no substantive requirements are placed on the truth of certain knowledge attributions. This rejects **NoCK**. However, such opponents usually concede that their proposals balk at pre-theoretic intuition, as a *cost to bear* when facing skeptical paradox. §5 addresses whether this is a sensible trade-off; for now, we require only that rejecting **MinC** or **NoCK** is troubling and counter-intuitive.

It bears emphasis: **MinC** and **NoCK** are relatively *weak*. For example, **MinC** shouldn't be confused with the claim that having knowledge presupposes having meta-evidential knowledge. Access externalists deny that meta-evidential knowledge (connecting φ to ψ) is required for being positioned to know ψ upon learning φ. Suppose that if X were to learn that there are exactly ten apples, she would come to know (via inference) that there are an even number. It doesn't follow, access externalists argue, that X knows that these propositions stand in the relation of entailment, or knows that her cognitive faculties are sound. Arguably, common sense concurs.[18] A **MinC** supporter can agree. She insists only that, under the circumstances, X must (be positioned to) know that if there are exactly ten apples, there is an even number. This isn't meta-evidential knowledge (it's about the apples, not about any agent or evidence), and it sounds comparatively odd to deny that x has it.

[18] Alston [1988] surveys various meta-evidential principles.

Similarly, **MinC** is distinct from the claim that knowledge acquisition requires knowledge of 'authenticity conditions'. Dogmatic modest foundationalists (i.e., perceptual dogmatists) claim that it *seeming* to agent X that there is an apple provides, by itself, *prima facie* warrant for her to believe there is an apple [Pryor, 2000],[Lyons, 2017, Sect. 3.4]. Suppose that, absent defeat, it bestows *knowledge*. The dogmatist denies, then, that basic perceptual knowledge must partly derive its warrant from prior knowledge of 'authenticity conditions' that connect seemings to reality. Arguably, common sense concurs. A **MinC** supporter can agree. **MinC** concerns the relationship between pieces of potential knowledge. It says nothing about the structure of warrant transmission or foundational knowledge. Likewise, the dogmatist needn't deny **MinC** [Dorr et al., 2014].[19] To see this, the interesting case, on dogmatic thinking, is when X's prior evidence $E_{@}$ and the prevailing conditions are such that: were it to seem to X that p, then X would have warrant sufficient for knowing p. Thus, were X to come to *know* that it seems to X that p, then it would seem to her that p (by factivity), and so she would be positioned to know p. Does it follow, as **MinC** requires, that X is positioned to know that p is materially implied by it seeming to her that p? If the dogmatist accepts both that knowability is closed under disjunction introduction and that a seeming state (and its absence) is luminous, they can reason as follows. If it seems to X that p, then X can know p, and so come to know the material conditional by disjunction introduction: either it doesn't seem that p, or p. Alternatively, if it *doesn't* seem to X that p, then luminosity allows X to *know* it doesn't seem so, and so X can again deploy disjunction introduction to know: either it doesn't seem that p, or p. Either way, **MinC** is respected.

3.6 *Schaffer abstracted, refined, bolstered*

We can now finalize our main argument.

Schaffer Abstracted

A1. Every RA theory admits an AMC.

[19] Pryor [2013, Sect. 7] and Comesaña [2014] observe a similar point, with respect to a probabilistic analogue to **MinC** (bearing on propositional justification): the principle that the conditional probability of p given q is bounded by the unconditional probability of $p \supset q$.

$\therefore$ No RA theory validates both **MinC** and **NoCK**.

A2. **NoCK** is a true principle (valid for *reality*).

A3. **MinC** is a true principle (valid for *reality*).

C. $\therefore$ RAT is false.

A1 is a theorem; intuitive considerations ground **A2** and **A3**. Thus, an RA theory, whatever its assessment of prior-Mica, must oppose one of two epistemic principles that together explain our intuitive revulsion at (e.g.) the claim that (prior-)Mica would know the canary is domestic if she were to see the red plumage: prior-Mica has no empirical evidence to support the empirical claim that only domestic canaries have red plumage; thus, by **NoCK**, she cannot know this claim; thus, by **MinC**, she cannot come to know the canary is domestic simply by observing its plumage.

Going beyond normality theory, we illustrate the dilemma for entitlement theory. Assuming claim p – that there's a legal duty to read contracts – is true, one is typically positioned to know p by learning a:

a: law professor X (expert in the subject area) confidently **a**sserted p

But what of $a \supset p$? Its alternatives (e.g. X is a pathological liar, or has a rare blind spot) are skeptical scenarios compatible with typical evidence E. Nor can apriori reasoning establish the falsity of $a \wedge \neg p$. But if some $\neg p$ scenarios aren't eliminable via reason or evidence, how is p knowable? An entitlement theorist has a ready explanation: absent specific defeat, one has non-evidential entitlement to accept $a \supset p$, rooted in *practical* rationality. (Reflecting, perhaps, a dominant strategy for inquiry [Wright, 2014].) Thus, $a \wedge \neg p$ worlds are irrelevant for epistemic evaluation. It follows, however, that empirical claim $a \supset p$ is cheaply knowable. **NoCK** is violated. But it strains credulity to posit that merely practical/strategic considerations let one *know*, without evidence, that if professor X asserts something then it's true.

The entitlement theorist can retreat, insisting that knowledge of empirical claims requires evidence. For instance, she might claim that entitlement is *question-sensitive* (cf. Schaffer [2005]): if the question is whether p is true, then one is entitled to ignore $a \wedge \neg p$ worlds; if the question is whether $a \supset p$

is true, then one is *not* so entitled. But since the $a \wedge \neg p$ worlds are skeptical, e doesn't eliminate them. Thus, $a \supset p$ isn't known. **MinC** is violated. Our revised entitlement theorist seems to endorse the following bizarre dialogue as a normal and accurate expression of speaker knowledge in a fixed context:

A. Agreed, professors like X generally know a lot about contract law. Do you also agree, then, that if X says there's a legal duty to read contracts, then X is right?

B. Not at all. For all I know, X said this and is wrong, despite being *generally* well-informed.

A. I've got news: X *did* say there's a legal duty to read contracts. Did you learn *anything* useful from this news?

B. Of course! I learned that there's a legal duty to read contracts.

4 Connections and Contrasts

Schaffer Abstracted is intimately related to two classic epistemic issues.

4.1 *Closure under deduction*

Consider an intuitive closure principle:

(15) **ClosOR**: $K\varphi \vDash K(\psi \vee \varphi)$

Proposition 2 grounds another important argument against RAT:

Corollary 2 *No RA theory validates both **ClosOR** and **NoCK**.*[20]

Proof: Suppose T validates **ClosOR** and **NoCK**. Suppose T admits $\mathcal{C}$ where $\mathcal{C} \vDash Kp \wedge (p \Rightarrow q) \wedge \neg\blacksquare(p \supset q)$. By (v) in Proposition 1, $\mathcal{C} \vDash Kq$. By **ClosOR**, $\mathcal{C} \vDash K(p \supset q)$. So, by **NoCK**, $\mathcal{C} \vDash E(p \supset q)$. So, $\mathcal{C}$ isn't a manifest AMC. Hence, T doesn't admit any manifest AMCs. So, by Proposition 2, T isn't an RA theory.

So an RA theorist committed to **NoCK** must reject **MinC** and **ClosOR** as a package. How close is their relationship? To avoid a superficial analysis,

[20]Cf. Proposition 1 in [Holliday, 2015b, §2.5].

we need some definitions. A context $\mathcal{C}$ is *structurally primed for **ClosOR*** iff: $R_@(\psi \vee \varphi) \subseteq R_@(\varphi)$ for all φ and ψ. Context $\mathcal{C}$ is *structurally primed for **MinC*** iff: for all φ and ψ, if $R_@(\psi) \subseteq R_@(\varphi) \cup E^c_@$ then $R_@(\varphi \supset \psi) \subseteq E^c_@$.

If $\mathcal{C}$ is structurally primed for **ClosOR** then the property $@ \models K\varphi \supset K(\psi \vee \varphi)$ reflects the structure of $\mathcal{C}$'s relevancy function. If $\mathcal{C}$ is structurally primed for **MinC**, the property $@ \models (\varphi \Rightarrow \psi) \supset K(\varphi \supset \psi)$ similarly reflects $\mathcal{C}$'s relevancy function.

Proposition 3 *i. There is an RA theory* `T` *that validates **ClosOR** but not **MinC**.*

*ii. Every sensible context that is structurally primed for **ClosOR** is structurally primed for **MinC**.*

*iii. There is an RA context $\mathcal{C}$ that is structurally primed for **MinC** but not **ClosOR**.*

iv. If `T` *only admits contexts that are structurally primed for **MinC** and where $R_@(\neg\varphi \supset \psi) = R_@(\varphi \vee \psi)$ for all φ and ψ, then* `T` *validates **ClosOR**.*

Thus, the interaction between **ClosOR** and **MinC** is close, but nuanced.

4.2 *The paradox of the criterion*

The paradox of the criterion proceeds as follows (cf. [Cohen, 2010, pg. 141]). Empirical knowledge requires knowledge that one's perceptual faculties are reliable. But one can only know this on the basis of empirical knowledge. So, anti-circularity precludes empirical knowledge.

Let's make this precise. Let q be that Clinton lost the 2016 US election and p be that the morning newspapers reported so. Given normal background evidence $E_@$ (expressed by e), learning the morning newspapers reported Clinton lost would position one to know she lost.

(16) $(e \wedge p) \Rightarrow q$

But this requires knowledge that $e \wedge p$ isn't misleading on the question of q: given $E_@$, one can exclude that $e \wedge p$ holds without q.

(17) $((e \wedge p) \Rightarrow q) \supset K((e \wedge p) \supset q)$

But conspiracy theories in which the newspapers are systematically misleading are conceivable and consistent with normal background evidence.

(18) $\neg\blacksquare((e \wedge p) \supset q)$

So, $(e \wedge p) \supset q$ isn't knowable apriori: if known, it's on the basis of $E_@$.

(19) $\neg\blacksquare((e \wedge p) \supset q) \supset \neg A((e \wedge p) \supset q)$

(20) $K((e \wedge p) \supset q) \supset (E((e \wedge p) \supset q) \vee A((e \wedge p) \supset q))$

But can $e \wedge p \wedge \neg q$ be excluded on the basis of $E_@$? It seems objectionably circular to take one's total evidence as a basis for knowing that one's total evidence isn't misleading. Indeed, $E_@$ only excludes $\neg e$-possibilities.

(21) $\neg E((e \wedge p) \supset q)$

But (16)-(21) are jointly inconsistent. (16), (17) and (20) entail either $E((e \wedge p) \supset q)$ or $A((e \wedge p) \supset q)$ is true. (18), (19) and (21) entail neither is true.

This is a paradox for everyone. To escape is to somewhere oppose intuition. A *skeptic* rejects (16). An *irrationalist* rejects (17): gaining knowledge, on her view, doesn't require (potential) foreknowledge that its source isn't misleading.[21] A *rationalist* rejects (18): she holds that it is a reflective truth that $e \wedge p$ isn't misleading on q. This encompasses infallibilists that count $(e \wedge p) \supset q$ as metaphysically necessary and reflective, and fallibilists that count $(e \wedge p) \supset q$ as metaphysically contingent and reflective.[22] A *liberal aprioritist* denies (19): she denies that the reflective truths exhaust the category of apriori truths, as it's knowable apriori that $e \wedge p$ isn't misleading on q. An *empiricist* rejects (21), allowing virtuous circularity: sources that provide knowledge of their own reliability.[23]

§3.6 shows that RAT is locked into distinctive options for resisting the paradox. Anti-skepticism precludes denying (16). The proof of Proposition 2 part ii shows that rejecting (18) or (21) is off the table: RAT classifies the interaction between $e \wedge p$ and q as an AMC. So an RA theorist must

[21]Compare: Pryor [2000]; [Weatherson, 2007, pg. 174] on rejecting conditionalization; Black [2008]; Sosa [2009, Ch. 7] on 'animal knowledge'; Sharon and Spectre [2017].

[22]For infallibilism, see [Williamson, 2000, Ch.9]: if knowledge and evidence are identical and P is knowable on evidence E then E entails P. For fallibilism: compare Hawthorne [2002]; Weatherson [2005]; White [2006]; Cohen [2010]; Wedgwood [2013].

[23]Compare Sosa [2009, Ch.9]; Neta [2013]; Weatherson [2014] on rejecting principle 'LDLC'; Vogel [2014] on rejecting the 'Entailment Principle'.

be an irrationalist (denying **MinC**) or a liberal aprioritist (denying **NoCK**). Unlike other irrationalists, RA theorists *must* reject closure under deduction, per §4.1.

5 Strategies for response

How to respond to *Schaffer Abstracted* (§3.6)? Rejecting **A1** requires rejecting our formalization. One refinement incorporates *probabilistic* elimination: §5.1 shows that the core dilemma of §3.6 persists. An alternative refinement liberalizes our framework to allow so-called *junk knowledge*: §5.2 considers whether our dilemma here just mutates into a trilemma.

Rejecting **A2** or **A3** is to reject **NoCK** or **MinC**: the intuitive basis for the Schafferian worry is ruled as misleading. §4.2 provided momentum in this direction: rejecting **NoCK** or **MinC** answers the paradox of the criterion and *every* theorist (RA or otherwise) must choose their counter-intuitive poison. However, §5.3 casts doubt on arguments for **NoCK** rejection; §5.4 does so for **MinC** rejection.

5.1 *Probabilistic RAT*

A *probabilistic context* $\langle W, @, \mathsf{E}, \mathsf{Pr}, \tau, \mathsf{R}, \models\rangle$ enhances a sensible discourse context with a probability function Pr (defined on the subsets of W) and a *threshold for elimination* τ between 0 and 1. We leave open the origin of Pr or τ. A proposition P is now *eliminated by* Q iff $\mathsf{Pr}(P|Q) < \tau$, where $\mathsf{Pr}(P|Q)$ denotes the conditional probability of P given Q. Then:

- $w \models K\varphi$ iff $\mathsf{Pr}(R_w(\varphi)|E_w) < \tau$, i.e. E_w eliminates the relevant alternatives to φ
- $w \models \varphi \Rightarrow \psi$ iff $\mathsf{Pr}(R_w(\psi)|R_w^c(\varphi) \cap E_w) < \tau$, i.e. excluding the relevant alternatives to φ eliminates ψ's (outstanding) relevant alternatives
- $w \models A\varphi$ iff $\mathsf{Pr}(R_w(\varphi)) < \tau$
- $w \models E\varphi$ iff $\mathsf{Pr}(R_w(p)|E_@) < \tau$ and $\mathsf{Pr}(R_w(p)) \geqslant \tau$

Again, e expresses the actual evidence $E_@$ in $\mathcal{C}$.

A *probabilistic RA context* is a sensible $\mathcal{C}$ for which a claim p exists where:

(1) *Probabilistic Anti-skepticism.* $E_@$ renders p knowable, i.e., $@ \vDash Ep$

(2) *Probabilistic Modesty.* $\mathcal{C}$ includes a non-empty 'skeptical' proposition S_p such that:

- $\$_p \vDash e \wedge \neg p$ for every world $\$_p \in S_p$;
- $\mathsf{Pr}(S_p) \geqslant \tau$, i.e., S_p isn't eliminated reflectively;
- $\mathsf{Pr}(S_p|E_@) \geqslant \tau$, i.e., S_p isn't eliminated by the actual evidence.

This widens our scope for modeling RA theories. Cohen [2002, pg. 103] describes his 'internal criterion of relevance' as: "[The agent] lacks sufficient evidence (reason) to deny h, i.e. to believe not-h". Since we reserve 'relevance' for non-evidential constraints, let's regard this as an account of *elimination*. What's 'sufficient evidence'? Standard formalizations appeal to probability.

Now consider probabilistic RA context $\mathcal{C}$, with $@ \vDash Ep$ and skeptical proposition S_p. A simple fact about the probability calculus:[24] if $P \subseteq E$ then $\mathsf{Pr}(P|E) \geqslant \mathsf{Pr}(P)$. Since, $R_@(e \supset p) \subseteq E_@$, we have

$$(\dagger)\ \mathsf{Pr}(R_@(e \supset p)|E_@) \geqslant \mathsf{Pr}(R_@(e \supset p))$$

As $S_p \subseteq [e \wedge \neg p]$, we have $\tau \leqslant \mathsf{Pr}(S_p) \leqslant \mathsf{Pr}([e \wedge \neg p])$. So, $e \supset p$ isn't reflective.

Suppose that $\mathsf{Pr}(R_@(e \supset p)|E_@) < \tau$. Then, $\mathsf{Pr}(R_@(e \supset p)) < \tau$, by $(\dagger)$. So in this case, $e \supset p$ is apriori, despite not being reflective.

On the other hand, suppose that $\mathsf{Pr}(R_@(e \supset p)|E_@) \geqslant \tau$. In this case, $@ \nvDash K(e \supset p)$. Nevertheless, $@ \vDash e \Rightarrow p$, as $@ \vDash Ep$.

So our core dilemma is recovered: a probabilistic RA theorist must (essentially) reject **NoCK** or reject **MinC**.

5.2 *Junk knowledge*

Following [Hawthorne, 2004, Sect. 2.2], say that knowledge of $\varphi \supset \psi$ (i.e., $\neg\varphi \vee \psi$) is *junk* if it doesn't prime the agent for modus tollens (i.e., dis-

[24] See [Weatherson, 2007, pp.173–174] for a proof.

junctive syllogism). Suppose an agent's evidence strongly supports ψ, but not $\neg\varphi$. Thus, if her evidence for ψ is overridden by counter-evidence, she doesn't come to know $\neg\varphi$, but loses knowledge of $\varphi \supset \psi$.

A heretical RA theorist can appeal to the notion of junk knowledge to reject our proof for Proposition 2, casting doubt on a key premise in *Schaffer Abstracted*. Recall that the proofs for part i and part ii of Proposition 2 hinge (per Proposition 1) on the universal validity of:

I. $\neg E(e \supset p)$

II. $\neg E((e \wedge p) \supset q)$

The framework in §2.2 delivers these validities by constraining the relevant counter-possibilities to $e \supset p$ and $(e \wedge p) \supset q$ to e-worlds (respectively, $\neg(e \supset p)$-worlds and $\neg((e \wedge p) \supset q)$-worlds). These cannot be eliminated by $E_{@}$. But consider a context where p is knowable on the evidence, i.e, Kp. According to the heretic, $e \supset p$ is here rendered knowable as *junk*, even if no e-world (and so no $\neg(e \supset p)$-world) is eliminated. So to take junk knowledge seriously, one must depart from the orthodox RA idea that the knowability of φ only ever requires the elimination of relevant *alternatives to* φ, where the alternatives to φ are the $\neg\varphi$-scenarios.

To accommodate this heresy, our framework can be liberalized, following Holliday [2015b]. First, one drops the assumption of a *unique* relevancy set for every φ (a relevancy set for φ being a set of worlds the elimination of which renders φ knowable). This yields the *multi-path picture of knowledge*, as Holliday [2015b] calls it. Second, one drops the assumption that sensible contexts display contrast (§2.2), i.e., that a relevancy set for φ must contain only $\neg\varphi$-worlds. To illustrate, consider a context where Ep holds. Thus, every world in some relevancy set R for p is eliminated by e. According to the junk-friendly heretic, R also counts as a relevancy set for $e \supset p$. But R only contains $\neg e$-worlds (hence their elimination). Thus, $e \supset p$ has a relevancy set without *any* $\neg(e \supset p)$-worlds. Of course, the heretic needn't take such relevancy sets as the *only* ones for $e \supset p$. For example, eliminating every e-world should also, in principle, render $e \supset p$ knowable.

More precisely: a *multi-path frame* $\mathcal{F}$ is like an abstract epistemic frame but the *relevancy function* R returns a non-empty *set of sets* of worlds $\mathcal{R}_w(\varphi)$:

the *relevancy sets* for φ at w. To assure the possibility of junk, we stipulate:

- $\mathcal{R}_w(\varphi) \cup \mathcal{R}_w(\psi) \subseteq \mathcal{R}_w(\varphi \vee \psi)$ for all φ, ψ and w

A *discourse context* $\mathcal{C}$ is now an enriched multi-path frame $\langle W, @, \mathsf{E}, \mathsf{R}, \vDash \rangle$, where $\vDash$ is as before, except:

- $w \vDash K\varphi$ iff there exists $R \in \mathcal{R}_w(\varphi)$ such that $R \cap E_w = \varnothing$
- $w \vDash \varphi \Rightarrow \psi$ iff for every $R_\varphi \in \mathcal{R}_w(\varphi)$ there exists $R_\psi \in \mathcal{R}_w(\psi)$ such that $R_\psi \subseteq R_\varphi \cup E_w^c$
- $w \vDash A\varphi$ iff $\varnothing \in \mathcal{R}_w(\varphi)$
- $w \vDash E\varphi$ iff $\varnothing \notin \mathcal{R}_w(\varphi)$ and $w \vDash K\varphi$

A context is *apt* iff it has veridical evidence and makes actuality relevant. A *relevancy context* is an apt context $\mathcal{C}$ for which p and $\$_p$ exist such that: (i) $@ \vDash Ep$; (ii) $\$_p \vDash e \wedge \neg p$. A *relevancy theory* admits only apt contexts and at least one relevancy context.

Does the heretic successfully defuse the threat of missed clues? Even if our liberalized framework allows a relevance theorist to simultaneously validate **MinC** and **NoCK**, an initial worry is that it merely transforms the intuitively troubling dilemma of *Schaffer Abstracted* into an intuitively troubling *trilemma*: Proposition 2 can be evaded at the cost of rejecting the validity of I and II. As suggested in §4.2, I and II formalize the 'anti-circularity' intuition that an agent's empirical evidence cannot by itself render *a posteriori* knowledge of the reliability of that very evidence. More basically, *displaying contrast* has pre-theoretic appeal: certainly, it sounds odd to assert knowledge of a claim while also admitting that one's evidence doesn't exclude any scenario where it is false. Put differently, our liberalized framework opens a new front for relevance theorists to resist the paradox of the criterion: (21) may be rejected instead of (17) or (19). But embracing virtuous circularity is no more *intuitive* than other strategies.

Anyway, incorporating multiple relevancy sets and junk knowledge doesn't by itself *eradicate* AMCs (unlike failures of **ClosOR**): one easily constructs simple relevancy contexts that are AMCs. So, a heretic still owes independent motivation for excluding such contexts from her theory. For example, consider $\mathcal{C}$:

$$W = \{@, \$_r, w_1, w_2\} \qquad E_@ = \{@, \$_r, w_1\}$$

$$@ \models p \wedge q \wedge r \qquad \$_r \models p \wedge \neg q \wedge \neg r$$
$$w_1 \models \neg p \wedge \neg q \wedge r \qquad w_2 \models \neg p \wedge \neg q \wedge \neg r$$

$$\mathcal{R}_@(p) = \{R_p\} = \{\{w_1, w_2\}\} \quad \mathcal{R}_@(\neg p) = \{R_{\neg p}\} = \{\{@\}\}$$
$$\mathcal{R}_@(q) = \{R_q\} = \{\{w_1\}\} \qquad \mathcal{R}_@(r) = E^c_@ = \{w_2\}$$

$$\mathcal{R}_@(p \supset q) = \{R_{\neg p}, R_q, [p \wedge \neg q]\} = \{\{@\}, \{w_1\}, \{\$_r\}\}$$

(The set $R_{\neg p} \cap R_q$ is typically a natural choice for a relevancy set for $p \supset q$, but here $R_{\neg p} \cap R_q = \varnothing$. As $p \supset q$ is empirical, we exclude $R_{\neg p} \cap R_q$ from $\mathcal{R}_@(p \supset q)$ to honor **NoCK** and disallow weakly apriori claims.)

Then:

- $@ \models \neg\blacksquare(p \supset q)$, as $\$_r$ is a $p \wedge \neg q$-world
- $@ \models \neg K(p \supset q)$, as there is no $R \in \mathcal{R}_@(p \supset q)$ s.t. $R \subseteq E^c_@ = \{w_2\}$
- $@ \models \neg E(p \supset q)$, as $@ \not\models K(p \supset q)$
- $@ \models p \Rightarrow q$, as $R_q \subseteq R_p$

So, $\mathcal{C}$ is a (latent) AMC. Admitting it violates **MinC**.

5.3 *Cheap knowledge*

Rejecting **NoCK** allows cheap knowledge: weakly apriori empirical claims. This has popular precedent: contextualists such as Stine [1976], Cohen [1988] and Lewis [1996] classify the denials of radical skepticial hypotheses as cheap knowledge. The typical motivation for rejecting **NoCK** is *negative*: knowability, it is held, is obviously closed under deduction; but an RA theorist can only preserve closure by allowing cheap knowledge (cf. §4.1); so, cheap knowledge must be admitted.

This assumes closure under deduction is sacrosanct. But this shouldn't be assumed. First, an RA theorist shouldn't endorse closure *simply* because it's intuitive on its face. This ignores the dilemma: **NoCK** is also intuitive on its face. Second, plausible independent explanations for closure failure have been proposed, appealing to the relationship between knowledge and evidence [Sharon and Spectre, 2017] or topicality [Yablo, 2014]. These invite dispute [Comesaña, 2017, Yablo, 2017], but deserve serious attention. Third, closure rejection defuses a striking range of epistemic paradoxes, so has clear theoretical utility [Dretske, 1970, Cohen, 2002, Sharon and Spectre, 2010, 2017, Yablo, 2017]. Fourth, closure rejection needn't be identified with rejecting deductive knowledge entirely, or rejecting universally uncontroversial epistemic principles (e.g. knowing a conjunction entails knowing each conjunct). Closure deniers typically endorse carefully *restricted* closure principles.[25]

So, a *positive* case is needed for rejecting **NoCK**. Resources are scarce. Are **NoCK** violations evidenced by mundane cases? Apparently not. **NoCK** deniers emphasize basic epistemic principles like 'My senses aren't systematically deceptive' as paradigmatic cheap knowledge. Pre-theoretic judgment doesn't align: hesitation in ascribing such knowledge is easily provoked by classic epistemic paradoxes. Indeed, rather than exploiting pre-theoretic judgment, **NoCK** deniers expend considerable energy explaining it away.[26] As they recognize, positing knowledge based *only* on 'factors of relevance' provokes an 'incredulous stare'.

A **NoCK** denier might argue: (i) there are convincing arguments that certain deeply contingent truths are knowable apriori; (ii) it's difficult to describe a purely reflective mechanism by which these could be known; (iii) no such mechanism is required if this knowledge is cheap; thus, there is cheap knowledge. But it isn't clear that adherents of (i) should accept (ii). Consider a key argument for (i) discussed by Hawthorne [2002], White [2006], Cohen [2010] and Wedgwood [2013]. Start with the premise that **(*)** is a reflective truth.

(*) P is established if one has a perceptual experience as of P (without defeat).

Now deploy suppositional reasoning: suppose one has a perceptual experience as of P (without defeat). Use (*) to conclude P. Now exit suppositional reasoning with material conditional introduction, concluding 'either I will not have a perceptual experience as of P (without defeat) or P'. De-

[25] See Cohen [2002], Black [2008], [Yablo, 2014, Ch.7], Holliday [2015a], Hawke [2016a].

[26] Some contextualists posit that confusing context-shifts explain ordinary hesitation in ascribing knowledge that skeptical possibilities don't hold. See Cohen [1988], DeRose [1995], Lewis [1996].

spite this claim being deeply contingent (its negation is metaphysically possible), the foregoing reasoning plausibly renders it knowable - using only reflective tools. If so, premise (i) in the argument for **NoCK** denial is true, but (ii) is false. Of course, one could deny the reasoning is entirely reflective by denying that (*) is reflective, positing instead that it is cheap. However, this merely begs the question if the issue is: does cheap knowledge exist?

5.4 *Inferential dogmatism*

Is RAT's best bet to reject **MinC**? Dorr et al. [2014, Sect.6] field a candidate mundane counterexample to **MinC**. Suppose you know a digital thermometer's reading is 45 degrees. Presumably, you can thereby know, by inference, that the temperature is between 40 and 50 degrees. But there was a nonzero objective chance that the temperature fluctuated to 55 degrees while an internal fluctuation in the thermometer held its reading at 45 degrees. It presumably follows that, for all you knew *before* inferring, the thermometer was reading 45 despite the temperature not being between 40 and 50.

But this judgment is tentative at best. The example is more paradoxical than decisive. Inferential anti-dogmatism is intuitively attractive; that consulting a thermometer issues knowledge about the temperature is intuitively attractive; that one cannot rule out events with non-zero objective chance is intuitively attractive. Which to reject? Furthermore, rejecting **MinC** while preserving **NoCK** commits the RA theorist to rejecting the closure of knowability under deduction. This shouldn't be casually embraced.

References

W. P. Alston. An Internalist Externalism. *Synthese*, 74(3):265–283, 1988.

A. Bacon. Inductive Knowledge. *Noûs*, 54(2):354–388, 2020.

T. Black. The Relevant Alternatives Theory and Missed Clues. *Australasian Journal of Philosophy*, 81:96–106, 2003.

T. Black. Solving the Problem of Easy Knowledge. *The Philosophical Quarterly*, 58(233):597–617, 2008.

M. Blome-Tillmann. *Knowledge and Presuppositions*. Oxford University Press, 2014.

D. Bradley. A Relevant Alternatives Solution to the Bootstrapping and Self-Knowledge Problems. *Journal of Philosophy*, 111(7):379–393, 2014.

A. Brueckner. What Missed Clue Cases Show. *Analysis*, 63(4):303–305, 2003.

R. M. Chisholm. *Theory of Knowledge*. Prentice-Hall, 1966.

S. Cohen. How to Be a Fallibilist. *Philosophical Perspectives*, 2:91–123, 1988.

S. Cohen. Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2):309–329, 2002.

S. Cohen. Bootstrapping, Defeasible Reasoning and A Priori Justification. *Philosophical Perspectives*, 24:141–159, 2010.

J. Comesaña. There is No Immediate Justification. In Steup, Turri, and Sosa, editors, *Contemporary Debates in Epistemology*, pages 222–234. Wiley Blackwell, second edition, 2014.

J. Comesaña. On Sharon and Spectre's Argument Against Closure. *Philosophical Studies*, 174:1039–1046, 2017.

K. DeRose. Solving the Skeptical Problem. *Philosophical Review*, 104(1): 1–52, 1995.

C. Dorr, J. Goodman, and J. Hawthorne. Knowing Against the Odds. *Philosophical Studies*, 170:277–287, 2014.

F. I. Dretske. Epistemic Operators. *Journal of Philosophy*, 67(24):1007–1023, 1970.

F. I. Dretske. The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40:363–378, 1981.

D. Drucker. The Attitudes We Can Have. *The Philosophical Review*, 129(4): 591–642, 2020.

A. Goldman. Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73:771–791, 1976.

P. Hawke. Questions, Topics and Restricted Closure. *Philosophical Studies*, 73(10):2759–2784, 2016a.

P. Hawke. Relevance in Epistemology and Logic. In Redmond, Pombo Martins, and Nepomuceno Fernàndez, editors, *Epistemology, Knowledge and the Impact of Interaction*, volume 38 of *Logic, Epistemology, and the Unity of Science*, pages 207–235. Springer, 2016b.

P. Hawke. *The Problem of Epistemic Relevance*. PhD thesis, Stanford University, 2017.

J. Hawthorne. Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2):247–269, 2002.

J. Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004.

M. Heller. Relevant Alternatives. *Philosophical Studies*, 55:23–40, 1989.

W. H. Holliday. Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. *Journal of Philosophical Logic*, 44(1):1–62, 2015a.

W. H. Holliday. Fallibilism and Multiple Paths to Knowledge. In T. Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, volume 5, pages 97–144. Oxford University Press, 2015b.

A. Kratzer. *Modals and Conditionals*. Oxford University Press, 2012.

K. Lawlor. *Assurance: An Austinian View of Knowledge Claims*. Oxford University Press, 2013.

D. Lewis. Elusive Knowledge. *Australalasian Journal of Philosophy*, 74(4):549–567, 1996.

J. Lyons. Epistemological Problems of Perception. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2017 edition, 2017.

R. Neta. Easy knowledge, transmission failure and empiricism. In

T. Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, volume 4, pages 166–84. Oxford University Press, 2013.

D. Pritchard. *Epistemological Disjunctivism*. Oxford University Press, 2012.

J. Pryor. The Skeptic and the Dogmatist. *Noûs*, 34(4):517–549, 2000.

J. Pryor. Problems for Credulism. In C. Tucker, editor, *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, pages 89–133. Oxford University Press, 2013.

J. Schaffer. Knowledge, Relevant Alternatives and Missed Clues. *Analysis*, 61(3):202–208, 2001.

J. Schaffer. Contrastive Knowledge. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology 1*, pages 235–271. Oxford University Press, 2005.

A. Sharon and L. Spectre. Dogmatism Repuzzled. *Philosophical Studies*, 148 (2):307–321, 2010.

A. Sharon and L. Spectre. Evidence and the Openness of Knowledge. *Philosophical Studies*, 174:1001–1037, 2017.

E. Sosa. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford University Press, 2009.

J. Stanley. *Knowledge and Practical Interests*. Oxford University Press, 2005.

G. Stine. Skepticism, Relevant Alternatives and Deductive Closure. *Philosophical Studies*, 29:249–261, 1976.

J. Vogel. e&¬h. In D. Dodd and E. Zardini, editors, *Scepticism and Perceptual Justification*, pages 87–107. Oxford University Press, 2014.

B. Weatherson. Scepticism, rationalism, and externalism. In T. Szabó Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, volume 1, pages 311–31. Oxford University Press, 2005.

B. Weatherson. The Bayesian and the Dogmatist. *Proceedings of the Aristotelian Society*, CVII:169–185, 2007.

B. Weatherson. Probability and scepticism. In D. Dodd and E. Zardini, editors, *Scepticism and Perceptual Justification*, pages 71–86. Oxford University Press, 2014.

R. Wedgwood. A priori bootstrapping. In A. Casullo and J. C. Thurow, editors, *The A Priori in Philosophy*, pages 226–46. Oxford University Press, 2013.

R. White. Problems for Dogmatism. *Philosophical Studies*, 131:525–557, 2006.

T. Williamson. *Knowledge and its Limits*. Oxford University Press, 2000.

C. Wright. On Epistemic Entitlement (II): Welfare State Epistemology. In Dodd and Zardini, editors, *Scepticism and Perceptual Justification*, pages 213–247. Oxford University Press, 2014.

S. Yablo. *Aboutness*. Princeton University Press, 2014.

S. Yablo. Open Knowledge and Changing the Subject. *Philosophical Studies*, 174:1047–71, 2017.

A Proofs

Proposition 1: Suppose $\mathcal{C}$ is admitted by $\mathtt{U}$.

(i) $@ \vDash Ke$ holds iff $E_@ \cap R_@(e) = \varnothing$. Since contrast is displayed: $R_@(e) \subseteq E^c_@$. Thus $@ \vDash Ke$ holds only if $E_@ \cap E^c_@ = \varnothing$, a triviality.

(ii) Assume $@ \vDash K\varphi$. So, $R_@(\varphi) \subseteq E^c_@$. So, $R_@(\varphi) \subseteq R_@(\psi) \cup E^c_@$. So, $@ \vDash \psi \Rightarrow \varphi$.

(iii) Assume $@ \nvDash \varphi$. As actuality is relevant: $@ \in R_@(\varphi)$. As evidence is veridical: $@ \in E_@$. Thus: $R_@(\varphi) \cap E_@ \neq \varnothing$. Thus: $@ \nVdash K\varphi$.

(iv) Assume $@ \vDash E\varphi$ or $@ \vDash A\varphi$. Either way: $R_@(\varphi) \subseteq E^c_@$. Thus: $@ \vDash K\varphi$.

(v) Assume $@ \vDash (\varphi \Rightarrow \psi) \wedge (\psi \Rightarrow \chi)$. Thus: $R_@(\chi) \subseteq R_@(\psi) \cup E^c_@$ and $R_@(\psi) \subseteq R_@(\varphi) \cup E^c_@$. So, $R_@(\chi) \subseteq R_@(\varphi) \cup E^c_@$. So, $@ \vDash \varphi \Rightarrow \chi$.

(vi) Assume $@ \vDash \blacksquare\varphi$ and $@ \vDash \blacksquare\varphi \supset \psi$. Thus, for every w: $w \vDash \varphi \wedge (\varphi \supset \psi)$. Thus, $w \vDash \psi$ for every w. So, $@ \vDash \blacksquare\psi$.

(vii) For reductio, assume $@ \models E(e \supset \varphi)$. So, $R_@(e \supset \varphi) \subseteq R^c_@ = [\neg e]$ and $R_@(e \supset \varphi) \neq \emptyset$. As contrast is displayed, $R_@(e \supset \varphi) \subseteq [e \wedge \neg\varphi] \subseteq [e]$. So, $R_@(e \supset \varphi)$ is a non-empty subset of $[e] \cup [\neg e]$. Contradiction.

(viii) For reductio, assume $@ \models E((e \wedge \varphi) \supset \psi)$. So, $R_@((e \wedge \varphi) \supset \psi) \subseteq R^c_@ = [\neg e]$ and $R_@((e \wedge \varphi) \supset \psi) \neq \emptyset$. As contrast is displayed, $R_@((e \wedge \varphi) \supset \psi) \subseteq [e \wedge \varphi \wedge \neg\psi] \subseteq [e]$. So, $R_@((e \wedge \varphi) \supset \psi)$ is a non-empty subset of $[e] \cup [\neg e]$. Contradiction.

Proposition 3:

i. Consider $\mathcal{T} = \{\mathcal{C}\}$, where $\mathcal{C}$ is an RA context:

$W = \{@, \$_a, w_1\}$	$E_@ = \{@, \$_a, w_1\}$
$@ \models p \wedge q \wedge e \wedge a$	$w_1 \models \neg p \wedge \neg q \wedge e \wedge a$
$\$_a \models p \wedge \neg q \wedge e \wedge \neg a$	
$R(p) = R(q) = \{w_1\}$	$R(p \supset q) = \{\$\}$
Otherwise: $R(\varphi) = \{@\}$ if $@ \not\models \varphi$	$R(\varphi) = \emptyset$ if $@ \models \varphi$

Note that every truth at @ is apriori, except:

$$@ \models \neg Kp \wedge \neg Kq \wedge \neg K(p \supset q)$$

As $@ \models K\varphi$ implies $@ \models \varphi$, if $@ \models K\varphi$ then $@ \models K(\psi \vee \varphi)$. Hence, T validates **ClosOR**.

Further, $@ \models p \Rightarrow q$ (as $R_@(q) \subseteq R_@(p)$) and $@ \not\models K(p \supset q)$ (as $R_@(p \supset q) \nsubseteq E^c_@$). Hence, T doesn't validate **MinC**.

ii. Assume $\mathcal{C}$ is structurally primed for **ClosOR**: so, $R_@(\varphi \supset \psi) \subseteq R_@(\psi)$ for all φ, ψ. Now assume $R_@(\psi) \subseteq R_@(\varphi) \cup E^c_@$. So, $R_@(\varphi \supset \psi) \subseteq R_@(\varphi) \cup E^c_@$. As contrast is displayed, $R_@(\varphi \supset \psi) \subseteq [\varphi]$ and $R_@(\varphi) \subseteq [\neg\varphi]$. So, $R_@(\varphi \supset \psi) \subseteq E^c_@$.

iii. Let $\mathcal{C}$ be a context where: $W = \{@, \$, w_1, w_2\}$; $E_@ = \{@, \$, w_1\}$; otherwise $E_w = \{w\}$. Truth assignment:

$@ \models a \wedge b \wedge c \wedge e$	$\$ \models \neg a \wedge \neg b \wedge \neg c \wedge e$
$w_1 \models a \wedge b \wedge \neg c \wedge e$	$w_2 \models \neg a \wedge \neg b \wedge \neg c \wedge \neg e$

Relevance assignment:

$R_@(a) = R_@(b) = R_@(e) = \{w_2\}$
$R_@(c) = \{\$\}$
$R_@(b \supset c) = \{w_1\}$

For every other φ and w, set $R_w(\varphi) = \varnothing$ when $w \models \varphi$ and $R_w(\varphi) = \{w\}$ when $w \not\models \varphi$. Thus, $\mathcal{C}$ is a RA context: sensible and $@ \models Ea$.

$\mathcal{C}$ isn't structurally primed for **ClosOR**: $w_1 \in R_@(b \supset c)$ but $w_1 \notin R_@(c)$.

Call the following relation between φ and ψ the ***MinC** relation*: if $R_@(\psi) \subseteq R_@(\varphi) \cup E^c_@$ then $R_@(\varphi \supset \psi) \subseteq E^c_@$. We show that $\mathcal{C}$ is structurally primed for **MinC** by showing, by cases, that the **MinC**-relation holds between arbitrary φ and ψ.

Case 1: $\varphi = b$ and $\psi = c$. So, $\varphi \supset \psi = b \supset c$. So, $R_@(\varphi) = \{\$\}$, $R_@(\psi) = \{w_2\}$ and $R_@(\varphi \supset \psi) = \{w_1\}$. So, $R_@(\psi) \nsubseteq R_@(\varphi) \cup E^c_@$ and $R_@(\varphi \supset \psi) \nsubseteq E^c_@$, as $E^c_@ = \{w_2\}$. So, the **MinC**-relation holds.

Case 2: $\varphi \neq b$ or $\psi \neq c$. So, $\varphi \supset \psi \neq b \supset c$ and $R_@(\varphi \supset \psi) = \varnothing$ or $R_@(\varphi \supset \psi) = \{@\}$.

Case 2a: $R_@(\varphi \supset \psi) = \varnothing$. So, $R_@(\varphi \supset \psi) \subseteq E^c_@$. So, the **MinC**-relation holds.

Case 2b: $R_@(\varphi \supset \psi) = \{@\}$. So, $@ \not\models \varphi \supset \psi$, i.e., $@ \models \varphi$ and $@ \not\models \psi$. So, $R_@(\psi) = \{@\}$. If $\varphi = c$, then $R_@(\varphi) \cup E^c_@ = \{\$, w_2\}$. If $\varphi = b \supset c$, then $R_@(\varphi) \cup E^c_@ = \{w_1, w_2\}$. Otherwise, $R_@(\varphi) \cup E^c_@ = \{w_2\}$. In each case, $R_@(\psi) \nsubseteq R_@(\varphi) \cup E^c_@$. So, the **MinC**-relation holds.

iv. Suppose that if $\mathcal{C}$ in $\mathtt{T}$ then: for all φ and ψ, if $R_@(\psi) \subseteq R(\varphi) \cup E^c_@$ then $R(\varphi \supset \psi) \in E^c_@$. Consider any such $\mathcal{C}$ where: $@ \models K\varphi$. So, $R(\varphi) \subseteq E^c_@$. So, $R(\varphi) \subseteq R(\neg\psi) \cup E^c_@$. So, $R(\neg\neg\psi \vee \varphi) \subseteq E^c_@$. So, $@ \models K(\psi \vee \varphi)$.