


University Press Scholarship Online

Oxford Scholarship Online



Oxford Studies in Epistemology Volume 5

Tamar Szabó Gendler and John Hawthorne

Print publication date: 2015

Print ISBN-13: 9780198722762

Published to Oxford Scholarship Online: April 2015

DOI: 10.1093/acprof:oso/9780198722762.001.0001

Fallibilism and Multiple Paths to Knowledge

Wesley H. Holliday

DOI:10.1093/acprof:oso/9780198722762.003.0004

Abstract and Keywords

This chapter argues that epistemologists should replace a “standard alternatives” picture of knowledge, assumed by many fallibilist theories of knowledge, with a new “multipath” picture of knowledge. The chapter first identifies a problem for the standard picture: fallibilists working with this picture cannot maintain even the most uncontroversial epistemic closure principles without making extreme assumptions about the ability of humans to know empirical truths without empirical investigation. The chapter then shows how the multipath picture, motivated by independent arguments, saves fallibilism from this problem. The multipath picture is based on taking seriously the idea that there can be multiple paths to knowing some propositions about the world. An overlooked consequence of fallibilism is that these multiple paths to knowledge may involve ruling out different sets of alternatives, which should be represented in a fallibilist picture of knowledge. The chapter concludes by considering inductive knowledge and strong epistemic closure from this multipath perspective.

Keywords: knowledge, fallibilism, multipath picture, relevant alternatives, subjunctivism, tracking, safety, epistemic closure, skepticism, contextualism

If knowledge required the elimination of all logically possible alternatives, there would be no knowledge (at least of contingent truths).

Alvin Goldman [1976, 775]

There are always, it seems, possibilities that our evidence is powerless to eliminate If knowledge ...requires the elimination of all competing possibilities ...then, clearly we seldom, if ever, satisfy the conditions for applying the concept.

Fred Dretske [1981, 365]

1. introduction

Being a fallibilist isn't easy. A fallibilist about empirical knowledge, in Lewis's [1996] sense, holds that an agent can know a contingent empirical proposition *P*, even if she has not ruled out *every last way* that *P* could be false.¹ In this sense, it seems that most contemporary epistemologists are fallibilists, at least relative to some way of understanding what it is to "rule out" an alternative. And with good reason: if knowing a contingent empirical proposition *P* required ruling out *every last way* that *P* could be false, then we would have little if any empirical knowledge. Radical (p.98) skepticism would reign. Yet fallibilism, despite its promise for defending the possibility of knowledge, also faces problems. To borrow an analogy sometimes applied to philosophical projects, trying to fill in the details of a fallibilist theory of knowledge is like trying to install an unstretched carpet: flatten a problematic lump in one place and a new one appears elsewhere. But then again, the alternative of radical skepticism about knowledge is like having the rug pulled out from under your feet.

The primary goal of this paper is to argue that what I call the *standard alternatives picture*, assumed by many fallibilist theories, should be replaced by a new *multipath picture* of knowledge. In §4.2, I identify the problematic lumps in the standard picture: fallibilists working with this picture cannot maintain even the most uncontroversial (single-premise, logical) epistemic closure principles without having to make extreme assumptions about the ability of humans to know empirical truths without empirical investigation. In §4.3, I show how the multipath picture, motivated by independent arguments, saves fallibilism from this problem. The multipath picture is based on taking seriously the idea that there can be multiple paths

to knowing some propositions about the world. An overlooked consequence of fallibilism is that these multiple paths to knowledge may involve ruling out different sets of alternatives, which should be represented in our picture of knowledge. In §4.4, I consider inductive knowledge and strong epistemic closure principles from this multipath perspective.

In what follows, I presuppose familiarity with the kinds of *skeptical hypotheses* that motivate fallibilism about knowledge (see, e.g., Dretske 1970, 1981, 2005). For lack of space, I cannot review the standard examples here. Instead, I leave it to the reader's imagination to fill in abstract discussions of skepticism, fallibilism, and epistemic closure with specific scenarios and propositions. Lewis (1996, 549) said it best: "Let your paranoid fantasies rip—CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself—and soon you find that uneliminated possibilities of error are everywhere. Those possibilities of error are far-fetched, of course, but possibilities all the same. They bite into even our most everyday knowledge. We never have infallible knowledge."

1.1. Scenarios and Propositions

Let us begin with some preliminary points of terminology and notation used throughout.

We start with a set W of triples (w, a, t) where w is a *way the world could (or could not) be including agent a at time t* .² I use " w ", " v ", " u ", etc., for members of W , which I will call *scenarios*. For each scenario w , let W_w be the subset of W containing those scenarios that are metaphysically possible relative to w . Everything in this paper is compatible with the view that $W_w = W_v = W$ for all w and v , so that no scenarios are metaphysically impossible relative to any others, and compatible with the rejection of this view. I leave these as parameter choices for the reader. However, for simplicity I assume that W does not include any "logically impossible" scenarios (see below).

(p.99) Following standard set-theoretic notation, I use " \in " for the membership relation, " \notin " to deny the membership relation, " \subseteq " for the subset relation, " $\not\subseteq$ " to deny the subset relation, and " \subsetneq " for the *strict* subset relation ($A \subseteq B$ but $B \not\subseteq A$); for any sets A and B , $A - B = \{w \in A | w \notin B\}$ is the complement of B in A , $A \cup B$ is the union of A and B , and $A \cap B$ is their intersection; given a set X of sets, $\cup X$ (resp. $\cap X$) is the union (resp. intersection) of all members of X ; and given an indexed family $\{A_i\}_{i \in I}$ of sets, $\cup_{i \in I} A_i$ (resp. $\cap_{i \in I} A_i$) is the union (resp. intersection) of all the A_i sets.

My topic is knowledge of *propositions*. I use “P”, “Q”, “S”, etc., for propositions and “ ” for the set of all propositions under consideration. I assume that propositions are *true* or *false* at scenarios in W and that propositions can have truth-functional structure: if P is a proposition, so is the negation of P , denoted by “ $\neg P$ ” if P and Q are propositions, so is the disjunction of P and Q , denoted by “ $P \vee Q$ ”; and so on for other truth-functions.³ If P does not have the structure of a truth-function applied to one or more propositions, call it *TF-atomic*.⁴ As usual, an assignment of truth values to TF-atomic propositions determines a truth value for every proposition; and Q is a *TF-consequence* (resp. *TF-equivalent*) of P iff any such assignment makes Q true if (resp. iff) it makes P true.

For any proposition P , define $P = \{w \in W \mid P \text{ is true at } w\}$, the set of scenarios at which P is true.⁵ Given a classical understanding of negation, disjunction, conjunction, etc., and the ban on logically impossible scenarios, we have $\neg P = W - P$, $P \vee Q = P \cup Q$, $P \wedge Q = P \cap Q$, etc. Let us also define $P_w = P \cap W_w$, the set of scenarios metaphysically possible relative to w at which P is true. Relativeto w , P is metaphysically *necessary* (resp. *possible*) iff $P_w = W_w$ (resp. $P_w \neq \emptyset$), P is metaphysically *contingent* iff $\emptyset \neq P_w \neq W_w$, and P *metaphysically entails* Q (resp. is *metaphysically equivalent* to Q) iff $P_w \subseteq Q_w$ (p.100) (resp. $P_w = Q_w$). According to some non-structured proposition views (Stalnaker 1981, Lewis 1986), if for all scenarios w based on the way our world is, $P_w = Q_w$, then $P = Q$; but for propositions *qua* objects of knowledge, I do not make this strong assumption for standard reasons and for a reason specific to fallibilism, discussed in §4.4.2an.

Finally, I use “ C ” “ C ”, etc. for *contexts* of knowledge attribution or assessment. Nothing in what follows depends on what contexts *are*, beyond the assumption that contexts play a certain “functional role” (namely by being something to which the functions in §4.2.1 are relativized). Following DeRose [2009, 187], I say that an agent in a scenario w does or does not “count as knowing proposition P in context C ” or “relative to C ” Yet I intend all that follows to be consistent with invariantism as well as contextualism and relativism; invariantists can assume that there is only one constant context C .

2. The standard alternatives picture

In this section, I introduce a *standard alternatives picture* of knowledge, show how a family of fallibilist theories fit into this picture, and then argue that the picture is fundamentally flawed.

2.1. Relevancy Set and Uneliminated Set

The starting point of the standard alternatives picture is the idea that for each proposition to be known, there is “a set of situations each member of which contrasts with what is [to be] known ...and must be evidentially excluded if one is to know” (Dretske 1981, 373). Dretske proposes that we “call the set of possible alternatives that a person must be in an evidential position to exclude (when he knows that P) the *Relevancy Set*” (371). Similarly, let us call the set of alternatives for P that the person has *not* excluded the *Uneliminated Set*. According to this picture, there are two *functions* r and u , each of which takes as input a proposition P, scenario w , and possibly a context, and returns a set of *alternatives*, which I take to be *scenarios* (for reasons explained later):

- $r(P, w)$ = the set of (“relevant”) alternatives such that the agent in scenario w counts as knowing proposition P relative to context only if she has eliminated these alternatives;
- $u(P, w)$ = the set of (“uneliminated”) alternatives that the agent in scenario w has *not eliminated* as alternatives for P relative to context .

The reasons for relativizing these sets to a scenario and possibly a context are well known. First, since objective features of an agent’s situation in a scenario w may affect what alternatives are relevant in w and therefore what it takes to know P in w (see Dretske 1981, 377 and Derose 2009, 30f. on “subject factors”), we allow that $r(P, w)$ may differ from $r(P, v)$ for a distinct scenario v in which the agent’s situation is different. Second, if we allow—unlike Dretske—that features of the conversational context of those attributing knowledge to the agent (or the context of assessment of a knowledge attribution, in the sense

(p.101) of MacFarlane 2005) can also affect what it takes to count as knowing P in w relative to (see Derose 2009, 30f. on “attributor factors”), then we should allow that $r(P, w)$ may differ from $r'(P, w)$ for a distinct context'.

Similarly, if we allow that what counts as eliminating an alternative may vary with context (see Derose 2009, 30n29) or depend on the agent’s situation, then our u function should take in a context and scenario as well.

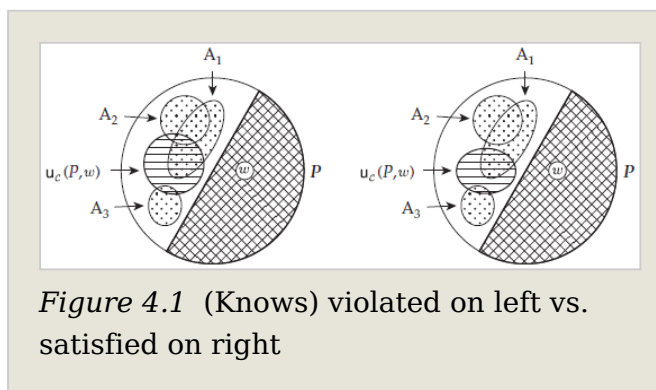


Figure 4.1 (Knows) violated on left vs. satisfied on right

According to the standard alternatives picture,⁶ an agent in scenario w counts as knowing P relative to context if and only if (or at least only if) (the agent believes P and) the following holds:

$$r_c(P, w) \cap u_c(P, w) = \emptyset$$

(Knows)

Fig. 4.1 shows the (Knows) condition violated vs. satisfied. Each of the large circles represents the set W of scenarios under consideration. The crosshatched region is the set P of scenarios in which the proposition P is true, including scenario w . The Relevancy Set and Uneliminated Set for P in w relative to context are shown in the ellipses with dots and horizontal lines, respectively, in the blank $\neg P$ -zone. If these sets overlap, as on the left, then the agent in w does not know P relative to c ; if they do not overlap, as on the right, then the agent in w knows P relative to c .⁷

In §2.3, I will show that a family of fallibilist theories fit into this picture as special cases, distinguished in part by the structural constraints they impose on the r and u functions. Some theories with more moving parts have another pair of functions r_c and u_c , also requiring $r_c(p, w) \cap u_c(p, w) = \emptyset$ for knowledge (see §4.2.3), but I will concentrate on theories with one pair of functions.

In virtue of what is an alternative in $r_c(p, w)$ or $u_c(p, w)$? For r , one can give “thick” or “thin” accounts of what it takes for a scenario to be in $r_c(p, w)$, depending on whether the account is independent of epistemic notions like *knowledge*. In §4.2.3, we will see some thick accounts with such independence, (p.102) but we already have a good thin account in the first bullet point above. Of course, this pushes us to the question about elimination and $u_c(p, w)$.⁸ But let us first consider the decision about what our “alternatives” are: *scenarios* or *propositions* or something else?

I take alternatives to be *scenarios*. What really matters is that the set of all $\neg P$ -alternatives in a context should form a nontrivial *partition* of the set of $\neg P$ -scenarios, so the alternatives are disjoint.⁹ (Recall the quotes from Stalnaker and Lewis in footnote 2.) We could call the cells in such a partition “Alternatives,” and let $r_c^A(p, w)$ and $u_c^A(p, w)$ be sets of Alternatives. But since I think of elimination in terms of scenarios, I take $r_c(p, w)$ and $u_c(p, w)$ to be sets of scenarios.¹⁰ This approach fits with what I consider the best-developed of previous fallibilist theories, discussed in §4.2.3. It also has other advantages, especially over taking the set of $\neg P$ -alternatives to be the set of *all*

propositions incompatible with P, which violates the disjointness of alternatives in a context.

For example, Vogel's (1999, 163) argument that *probability* cannot provide a sufficient condition for relevance of alternatives depends on assuming the proposition-based view of alternatives (see footnote 52 in §4.3.5). Moreover, the puzzling question (see Stine 1976, 258) of whether $\neg P$ is a relevant alternative to P —and if so, what it takes to “eliminate” $\neg P$ other than knowing P —suggests that the level of propositions might not be the best level at which to locate alternatives. It seems that one can give a more substantive account of what it is for a *scenario* to be (un)eliminated, since one may refer to the experiences or beliefs of the agent in that scenario, compared to those of the agent in another scenario. By contrast, accounts of what it is for a proposition to be eliminated seem not to take us very far from the idea of *knowing* the negation of the proposition.

According to Lewis [1996], “a possibility [v] is *uneliminated* iff the subject's perceptual experience and memory in [v] exactly match his perceptual experience and memory in actuality” (553). I will postpone discussion of whether such match is necessary.¹¹ All of the theories I consider seem to agree on at least this much: for $v \in W - P$, it is *sufficient* for $v \in u_c(P, w)$ that v and w are subjectively indistinguishable, appear the same way, etc., to the agent, given her total experience and memory, where this requires that the agent's (“narrow”) beliefs are the same in v and w. Many theorists would also agree that v and w are subjectively indistinguishable to the agent if she is in the same physical state in both, so this would provide another sufficient condition for $v \in u_c(P, w)$.

Given these sufficient conditions, it follows that for many contingent propositions P about the world external to the agent, $u_c(P, w) \cap (W_w - P) = \emptyset$. For given a scenario w , perhaps in which the agent believes P , there is another possible scenario v in which the agent is in the same physical state, or at any rate a scenario that is subjectively indistinguishable, but in which P is false, so $v \in u_c(P, w) \cap (W_w - P)$.¹² This is a reflection of the separation between mind and world.¹³

Given that $u_c(P, w) \cap (W_w - P) \neq \emptyset$ for so many empirical propositions P , radical skepticism about empirical knowledge follows from the “infallibilist” assumption that knowing a proposition P requires ruling out *all* possible $\neg P$ -scenarios, which in terms of Fig. 4.1 requires that

the dotted region covers the *entire* $\neg P$ -zone (at least within W_w , if W_w is a smaller circle than W):

infallibilism—for all w, P , and $C: W_w - P \subseteq r_c(P, w)$.

14

(p.104) It follows from infallibilism and $u_c(P, w) \cap (W_w - P) \neq \emptyset$ that $r_c(P, w) \cap u_c(P, w) \neq \emptyset$, so by (Knows), P is not known. Thus, in order to avoid radical skepticism, one must *at least* deny infallibilism:

fallibilism—for some w, P , and $C: W_w - P \subseteq r_c(P, w)$.

This is an extremely weak version of fallibilism: in effect, fallibilism about at least one possible case of knowledge. A stronger, but still extremely weak, version of fallibilism says that there is some proposition Q that is true in all of the relevant alternatives to P but not in all possible $\neg P$ -scenarios:

e-fallibilism—for some w, P, Q , and $C: r_c(P, w) \subseteq Q$ and $W_w - P \not\subseteq Q$.

Here “e-fallibilism” stands for *expressible fallibilism*, since it says that we can *express* with Q something that the relevant alternatives have in common with each other (and perhaps some other scenarios) but not with all possible $\neg P$ -scenarios. It would be a strange version of fallibilism that denied there was even one such proposition P for which we could express our fallibilism in this way. Note that if for every set of scenarios there is a corresponding proposition true in exactly those scenarios, then fallibilism is equivalent to e-fallibilism, taking Q to be the proposition corresponding to $r_c(P, w)$. Also note that e-fallibilism does not even require that the proposition Q be incompatible with P , i.e. $Q_w \subseteq W - P$. For that, one could assume what I will call *expressible contrast fallibilism*:

ec-fallibilism—for some w, P, Q , and $C: r_c(P, w) \subseteq Q$ and $Q_w \subseteq W_w - P$.

The reason for considering such weak principles will become apparent later. I will argue that being even a weak fallibilist is tricky, although not for the reasons that some philosophers think.¹⁵

(p.105) In addition to satisfying the above fallibilist conditions, all of the theories to be considered in the standard alternatives framework satisfy two further kinds of conditions. First, following Dretske’s characterization of the Relevancy Set for a proposition P as “a set of situations each member of which contrasts with what is [to be] known,” i.e. a set of $\neg P$ -scenarios, we have

contrast/enough— $r_c(P, w) \subseteq W - P$,

which says that the alternatives one must eliminate to know P are $\neg P$ -scenarios. (From now on I will leave the universal quantification over w , P , and c implicit.) A stronger version is

$$\text{M-contrast/enough} \rightarrow r_c(P, w) \subseteq W_w - P,$$

which says that the alternatives are all $\neg P$ -scenarios based on ways the world metaphysically *could be* (so an agent's ignorance cannot be witnessed by "impossible worlds"). Second, following Lewis's [1996] *Rule of Actuality*, that "actuality is always a relevant alternative" (554), we have

$$\text{r-RofA} \rightarrow w \notin P \text{ implies } w \in r_c(P, w),$$

which says that whenever w is a $\neg P$ -scenario, it is a relevant alternative that one must eliminate in order to know P in w . However, it is immediate from the sufficient condition for $v \in u_c(P, w)$ given above that an agent cannot eliminate her actual scenario:

$$\text{u-RofA} \rightarrow w \notin P \text{ implies } w \in U_c(P, w).$$

It follows from r-RofA and u-RofA together that if $w \notin P$, then $w \in r_c(P, w) \cap u_c(P, w) \neq \emptyset$, so by (Knows), P is not known. Hence only truths can be known.

In this framework we can also state necessary and sufficient conditions for epistemic closure. Let **R** be some relation that a sequence of propositions can bear to another proposition. Here is a general schema for an *empirical* epistemic closure principle with respect to **R**: if an agent knows propositions P_1, \dots, P_n , which together bear **R** to proposition Q, then, as MacFarlane (2014, 177) puts it, the agent "could come to know [Q] without further empirical investigation."¹⁶ This requires that

$$r_c(Q, w) \subseteq \bigcup_{1 \leq i \leq n} r_c(P_i, w),$$

to guarantee that if the agent has eliminated enough scenarios to know P_1, \dots, P_n , then she has eliminated enough to know Q. Note, though, that this guarantee assumes that if a scenario $v \in r_c(P_i, w) \cap r_c(Q, w)$ is eliminated as an (p.106) alternative for P_i , then v is also eliminated as an alternative for Q. In terms of *un* elimination:

$$\forall i \leq n: r_c(P_i, w) \cap r_c(Q, w) \cap u_c(Q, w) \subseteq u_c(P_i, w).$$

Together (1) and (2) imply that if (Knows) holds for P_1, \dots, P_n , then it holds for Q.¹⁷

As for specific closure principles, **R** could be the relation that the sequence P_1, \dots, P_n bears to Q iff Q is a TF-consequence of $\{P_1, \dots, P_n\}$, i.e. of $P_1 \wedge \dots \wedge P_n$. If $n = 1$, I call this *single-premise closure under TF-consequence*. If n is allowed to be arbitrary, I call this *multi-*

premise closure under TF-consequence. Sometimes I will not specify **R** explicitly, and I will write “ $(KP_1 \& \dots \& KP_n) \Rightarrow KQ$ ” to abbreviate the principle that if an agent knows P_1, \dots, P_n relative to c , then she could know Q relative to c without further empirical investigation. For example, *closure under conjunction elimination*, $K(P \wedge Q) \Rightarrow KQ$, says that if an agent knows $P \wedge Q(P_1)$, then she could come to know Q without further empirical investigation; *closure under known material implication*, $(KQ \& K(P \rightarrow Q)) \Rightarrow KQ$, says that if an agent knows $P(P_1)$ and $P \rightarrow Q(P_2)$ then she could come to know Q without further empirical investigation; and so on. Note that closure under known material implication is a multi-premise closure principle.¹⁸ In the next section, we will see a crucial pair of conditions that affect whether this principle holds.

2.2. The RS and RO Parameters

Fallibilists working with the standard alternatives picture face two questions. First, can one say whether a scenario v is simply “relevant” for the agent in a scenario w , independently of any proposition in question; or must one instead say that v is relevant in w as an alternative for a particular proposition Q , allowing that v may not be relevant in w as an alternative for a different proposition P ? Second, can one say whether v is simply “eliminated” by the agent in w , independently of any proposition in question; or must one instead say that v is eliminated in w as an alternative for a particular Q , allowing that v may not be eliminated in w as an alternative for a different P ?

Consider the first question. Dretske’s [1981] idea was that *for each proposition, there is a Relevancy Set for that proposition*, motivating the following definition of $RS_{\forall\exists}$ theories:

$RS_{\forall\exists}$ theories hold that for every context, for every scenario w , and for every (\forall) proposition P , there is (\exists) a set of *relevant (in w) $\neg P$ -scenarios*, $r_c(P, w) \subseteq W - P$, (p.107) such that in order to know P relative to the agent in w has to eliminate the scenarios in $r_c(P, w)$.

By contrast, Heller [1999] considers (and rejects) a version of the relevant alternatives (RA) theory in which “there is a certain set of worlds selected as relevant, and S must be able to rule out the not- p worlds within that set” (197), which suggests the following definition of $RS_{\exists\forall}$ theories:

$RS_{\forall\exists}$ theories hold that for every context and scenario w , there is (\exists) a set of *relevant (in w) scenarios*, $R_c(w)$, such that for every (\forall) proposition P , in order to know P relative to c the agent in w has to eliminate the $\neg P$ -scenarios in $R_c(w)$, i.e. the scenarios in $R_c(w) \cap (W - P)$.

As a simple logical point, every $RS_{\exists\forall}$ theory is a $RS_{\forall\exists}$ theory (take $r_c(P, w) = R_c(w) \cap (W - P)$), but not necessarily vice versa. From now on, when I refer to $RS_{\forall\exists}$ theories, I have in mind theories that are not also $RS_{\exists\forall}$ theories. As I will explain below, this distinction is at the heart of the disagreement about epistemic closure that pits Dretske [1970] and Nozick [1981], who defend $RS_{\forall\exists}$ theories, against Stine [1976] and Lewis [1996], who defend $RS_{\exists\forall}$ theories.

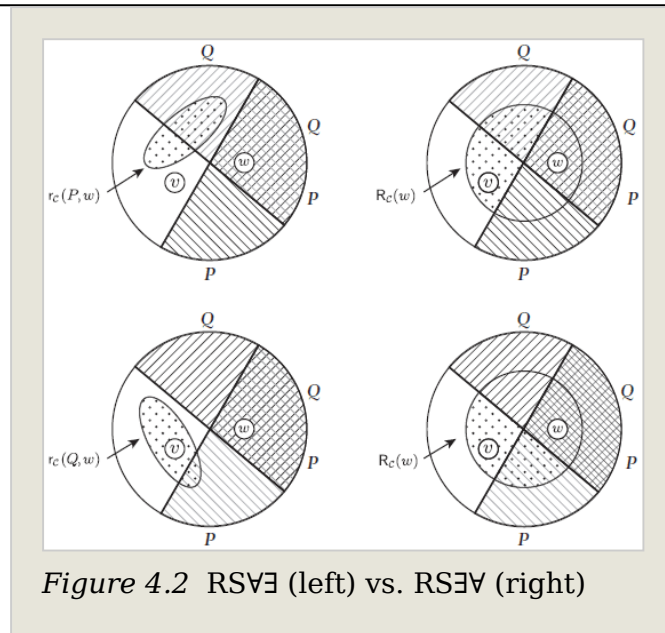
To be precise, let us define the following condition on r , of which $RS_{\forall\exists}$ is the denial:

$$RS_{\exists\forall} \text{---there is } (\exists) R_c(w) \subseteq W \text{ such that for all } (\forall) P:$$

$$R_c(P, w) = r_c(w) \cap (W - P).$$

In a *contextualist* $RS_{\exists\forall}$ theory, such as Lewis's (1996) RA theory, the set of relevant scenarios may change as context changes. Still, for any given context c , there is a set $R_c(w)$ of relevant (in w) scenarios, which does not depend on a particular proposition in question. The $RS_{\forall\exists}$ vs. $RS_{\exists\forall}$ distinction is about how theories view the relevant alternatives with respect to a *fixed context*.

Let us now return to the second question above: can one say, independently of any proposition in question, that v is eliminated by the agent in w ? According to Lewis's (1996) notion of elimination, the answer is "yes": whether there is exact match of experience and memory in v and w does not depend on any proposition in question. Hence for every scenario w , there is a fixed set of "uneliminated" scenarios $U_c(w) \subseteq W$ represent the fixed set, singled out independently of any proposition in question. However, as we shall see in §4.2.3, according to the notions of elimination implicit in *sensitivity* and *safety* theories of knowledge, the answer is "no"; it may be that v is eliminated as an alternative for a proposition P but not as an alternative for a proposition Q . Parallel to the definition of $RS_{\exists\forall}$ above, we define the following RO (for "ruling out") condition on u , of which $RO_{\forall\exists}$ is the denial:

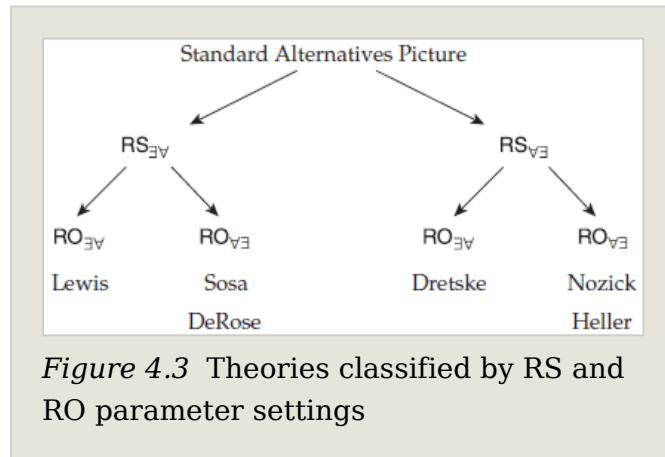


$RO_{\exists\forall}$ —there is $(\exists)U_c(w) \subseteq W$ such that for all $(\forall)P$:

$$U_c(P, w) = u_c(w) \cap (W - P).$$

(p.108)

Fig. 4.2 shows the difference between $RS_{\forall\exists}$ and $RS_{\exists\forall}$. Observe that v is a $\neg P$ -scenario and a $\neg Q$ -scenario. On the $RS_{\forall\exists}$ side (left), while v is a scenario that must be eliminated in order to know Q (where Q is the darker semicircle in the lower row), it is not a scenario that must be eliminated in order to know P (where P is the darker semicircle in the upper row). By contrast, on the $RS_{\exists\forall}$ side (right), where the inner circles represent the fixed set $R_c(w)$ of relevant scenarios, no such split-



decision on v is possible; so v is a scenario that must be eliminated in order to know P and in order to know Q . The pictures for $RO_{\forall\exists}$ vs. $RO_{\exists\forall}$ would be the same if we were to substitute u for r and U for R . As I will explain in §4.2.3, the theories of Lewis [1996], Sosa [1999], DeRose [1995], Dretske [1981], Nozick [1981], and Heller [1999] have the parameter settings in Fig. 4.3.

I claimed above that the distinction between $\forall\exists$ and $\exists\forall$ parameter settings is at the heart of the disagreement about epistemic closure. Assuming $RS_{\exists\forall}$ and $RO_{\exists\forall}$, the (Knows) condition beomes

$$\begin{array}{ccc} r_c(P, w) & \cap & u_c(P, w) & = \phi \\ \parallel & & \parallel & \\ R_c(w) \cap (W - P) & \cap & U_c(w) \cap (W - P) & = \phi, \end{array}$$

which is equivalent to

$$R_c(w) \cap U_c(w) \subseteq P.$$

(3)

(p.109) Now it is easy to see why $\exists\forall$ settings are to closure under known material implication. If the agent knows P and $P \rightarrow Q$, then as instances of (3) we have $R_c(w) \cap U_c(w) \subseteq P$ and $R_c(w) \cap U_c(w) \subseteq P \rightarrow Q$, which imply $R_c(w) \cap U_c(w) \subseteq Q$, so the agent has done enough elimination of scenarios to know Q . Indeed, this is why closure under known implication holds on Lewis's (1996) theory. By contrast, if we do not assume **RS** $_{\exists\forall}$ and **RO** $_{\exists\forall}$, then as shown in Fig. 4.2, a $(\neg P \wedge \neg Q)$ -scenario v that is relevant (or uneliminated) as an alternative for Q may not be relevant (or uneliminated) as an alternative for P , even if the agent knows the implication $P \rightarrow Q$, which opens up the possibility of a failure of closure under known implication (recall the end of §4.2.1). Indeed, this is why closure under known implication fails on Dretske's [1970] theory; an agent may know a mundane proposition P , because uneliminated skeptical scenarios v are not in $r_c(P, w)$, and yet fail to know Q , the *denial* of the skeptical hypothesis, because those v are in $r_c(Q, w)$.¹⁹

2.3. Unification

Let us see how some standard fallibilist theories are special cases of the standard alternatives picture. I will define the r and u functions according to each theory. With the exception of Lewis's (1996) theory, each theory requires belief for knowledge: if the agent in w does not believe P , then she does not know P ; if the agent in w does believe P , then, as the reader can verify, the (Knows) condition $r_c(P, w) \cap U_c(P, w) = \phi$ coincides with the knowledge condition of the theory. For **RS** $_{\exists\forall}$ theories, I will simply define $R_c(w)$, from which r is derived by $r_c(P, w) = R_c(w) \cap (W - P)$, and similarly for u . For example, for Lewis's (1996) RA theory, we have:

$R(w)$ = the set of scenarios that are not properly ignored in context c when attributing knowledge to the agent in scenario w ;

$U(w)$ = the set of scenarios in which the agent's perceptual experience and memory exactly match that of the agent in w .

(p.110) By contrast, for Dretske's (1981) RA theory (recall footnote 8) stated in terms of scenarios, we have:

$r_c(P, w)$ = the set of $\neg P$ -scenarios that the agent in w “must be in an evidential position to exclude” in order to know P (371);

$U_c(w)$ = the set of scenarios that the agent in w is not in an evidential position to exclude.

For Heller’s (1989, 1999) RA theory, we have the following definitions, “cashing out S ’s ability to rule out a not- p world in terms of her not believing p in that world” (1999, 198):

$r_c(P, w)$ = the set of closest $\neg P$ -scenarios according to an ordering²⁰ (dependent on the context) of scenarios “according to how realistic they are” [Heller, 1989, 25];

$u_c(P, w)$ = the set of $\neg P$ -scenarios where the agent²¹ believes P .

Thus, for Heller $r_c(P, w) \cap U_c(P, w) = \emptyset$ says that the agent does not believe P in any of the closest $\neg P$ -scenarios according to the ordering. For the similar *sensitivity* theories in the tradition of Nozick [1981] (without *adherence* and with counterfactuals understood following Lewis 1973²²) we have:

$r_c(P, w)$ = the set of closest $\neg P$ -scenarios according to an ordering (possibly dependent on the context) of scenarios for evaluating counterfactuals at w ;

$u_c(P, w)$ = the set of $\neg P$ -scenarios where the agent believes P (by the same method as in w).²³

(p.111) Theories that add an *adherence* condition use another pair r' and u' of functions such that $r'_c(P, w) = R'_c(w) \cap P$ where:

$R'(w)$ = the set of scenarios that are “close” or “nearby” to w (relative to);

$u'(P, w)$ = the set of P -scenarios where the agent does not believe P .²⁴

Thus, $r'_c(P, w) \cap U'_c(P, w) = \emptyset$ iff the agent believes P in all of the close P -scenarios.²⁵ Nozick’s [1981] full tracking theory adds this requirement to the sensitivity requirement above.²⁶

Finally, turning to *safety* theories in the tradition of Sosa [1999], we have:

$R_c(w)$ = the set of scenarios that are “close” or “nearby” to w (relative to);

$u_c(P, w)$ = the set of $\neg P$ -scenarios where the agent believes P (on the same basis as in w).

Thus, $r_c(P, w) \cap U_c(P, w) = \emptyset$ iff there are no close scenarios where the agent falsely believes P (on the same basis on which she believes P in w). Parallel to the fact that Nozick’s tracking theory requires sensitivity and adherence, DeRose’s [1995] “double safety” theory requires safety and adherence.

One can now check that the above definitions imply the classifications in Fig. 4.3.

It is important to realize that while safety theories are $RS_{\exists V}$ theories, which may lead one to think that they support full epistemic closure, they are also $RO_{\forall \exists}$ theories, so it is not at all obvious that they support full epistemic closure.²⁷ From the fact that in all close scenarios where the agent believes $P \wedge Q$, $P \wedge Q$ is true (and in all close scenarios where $P \wedge Q$ is true, the agent believes $P \wedge Q$), it obviously does not follow that in all close scenarios where the agent believes P , P is true. So an agent can have a (double) safe belief that $P \wedge Q$, even though she has an unsafe belief that P . But an agent who knows $P \wedge Q$ knows P , so safety theorists have some explaining to do.²⁸

(p.112) Now that we have definitions of r and u for each theory, we can investigate the properties of r and u implied by these definitions. For example, consider the theories according to which $r_c(P, w)$ is the set of closest $\neg P$ -scenarios according to some kind of ordering. We can extract a lot of information about r from this assumption. First, let us assume (cf. Lewis 1973, §2.3) that for each scenario w , there is a binary relation \leq_w^c on W_w that is a *total preorder*,²⁹ *weakly centered* on w ,³⁰ where we read $v \leq_w^c u$; as “ v is *at least as close* to w as u is.” Let us also assume that \leq_w^c is *well-founded*, which means that for every set $A \subseteq W$ of scenarios, if $A \cap W_w$ is nonempty, then

$$\text{Closest}_{\leq_w^c}(A) = \{v \in A \mid \forall u \in A \cap W_w : v \leq_w^c u\},$$

the set of closest scenarios to w among those in A , is also nonempty. This implies that for any proposition P , if P is possible relative to w ($P \cap W_w \neq \emptyset$), then there is a set of closest P -scenarios to w ($\text{Closest}_{\leq_w^c}(P) \neq \emptyset$), as epistemologists working with ordering-based theories typically assume. With this setup, we can completely characterize the properties of r for the ordering-based theories.

Theorem 1. Given a family $\{\leq_w^c\}_{w \in W}$ of orderings as above for each context c , the function r defined by $r_c(P, w) = \text{closest } \leq_w^c(W_w - P)$ satisfies all of the following conditions:

M-equiv—if $P_w = Q_w$ then $r_c(P, w) = r_c(Q, w)$;

M-contrast/enough— $r_c(P, w) \subseteq W_w - P$;

r-RofA— $w \notin P$ implies $w \in r_c(P, w)$;

noVK— $P_w \neq W_w$ implies $r_c(P, w) \neq \emptyset$;

alpha— $r_c(P \wedge Q, w) \subseteq r_c(P, w) \cup r_c(Q, w)$;

beta—if $P_w \subseteq Q$ and $r_c(P, w) \cap r_c(Q, w) \neq \emptyset$, then $r_c(Q, w) \subseteq r_c(P, w)$.

Conversely, given any function r satisfying these conditions, there is a family $\{\leq_w^c\}_{w \in W}$ of orderings for each c such that for all P and w , $r_c(P, w) = \text{closest } \leq_w^c(W_w - P)$.³¹

I omit the proof of Theorem 1, since it is essentially a variation on a well-known result of Arrow [1959], but here formulated using analogues of Sen's [1971] α and β conditions applied to r .³² With the possible exception of beta, all of the conditions should be self-explanatory. Most important for (p.113) our purposes in the next section will be the condition noVK for *no vacuous knowledge*.

2.4. The Problems of Vacuous Knowledge and Containment

All of the fallibilist theories developed so far in the standard alternatives picture have at least one of two serious problems, depending on whether they are $RS_{\exists V}$ theories or $RS_{\forall \exists}$ theories.

Assuming $RS_{\exists V}$, fallibilism implies that the set $R_c(w)$ of relevant/nearby scenarios is a *strict* subset of W_w . Thus, there can be *contingent* propositions Q ($Q_w \neq W_w$) true throughout $R_c(w)$ ($R_c(w) \subseteq Q$), as shown on the left of Fig. 4.4 after Proposition 1 in §2.5, where Q is the region with diagonal lines and $R_c(w)$ is the region with stars. But then $RS_{\exists V}$ implies $r(Q, w) = R(w) \cap (W - Q) = \emptyset$; and if $r(Q, w) = \emptyset$, then as long as the agent believes Q , she *knows* it, for *any* u function! No matter what (lack of) experience the agent has, and no matter what experience and beliefs the agent would have had under other circumstances, the agent supposedly knows the *contingent* proposition Q .³³

For example, according to Lewis's $RS_{\exists V}$ theory, even if the agent has never opened her eyes or ears, she knows any contingent Q that is true throughout the set $R(w)$ of relevant scenarios; and according to the $RS_{\exists V}$ safety theory, no matter how insensitive an agent's beliefs are to reality, she knows (or at least safely believes) any contingent Q that is true throughout the set $R(w)$ of nearby scenarios, provided

she believes it. Vogel [1999] recognizes this problem for some versions of the RA theory, observing that if we allow “for detailed empirical knowledge without evidence, then anyone who happens to arrive at the appropriate belief, no matter how, will enjoy that knowledge. This outcome is wrong; knowledge is dearer than that” (171f.). I call this problem the problem of *vacuous knowledge*, following Heller [1999], who also realizes that the $RS_{\exists v}$ assumption is to blame.

However, Heller and I view the problem differently. For Heller, the problem seems to be that when a contingent Q is true throughout $R(w)$, $RS_{\exists v}$ theories do not place a requirement on the agent to eliminate any $\neg Q$ -scenarios in order to know Q . In my view, the problem is that $RS_{\exists v}$ theories do not place on the agent *any requirement to eliminate any scenarios* in order to know Q . This distinction will come up again in the Answer to the First Reply below and in §4.3.3 and §4.4.1.

It will not help here to claim that Kripke [1980] has given examples of a priori knowable contingent truths. For one thing, we can take Q to be the set of scenarios v such that Q is true at v *considered as actual*, so $\phi?Qw?Ww$ means (p.114) that Q is *deeply* contingent (see Davies and Humberstone 1980). Then $RS_{\exists v}$ theories allow knowledge of deeply contingent truths with no requirement of eliminating scenarios. But even if one thinks there are some special counterexamples to Evans’s [1979] famous claim that “it would be intolerable for there to be a statement which is both knowable *a priori* and deeply contingent” (161), such examples are as much beside the point here as Kripke’s.

The main point is this: $RS_{\exists v}$ theories imply that *every proposition* Q with $R(w) \in Q$ is knowable with no requirement of eliminating scenarios, and there is no guarantee that every such Q fits the mold of one of the *recherché* examples of (deeply) contingent but a priori knowable propositions. Instead, $RS_{\exists v}$ theorists tell us that such Q may include the denials of skeptical hypotheses, not only what I call *self-side* skeptical hypothesis about how we are hooked up to the world (BIVs, etc.), but also *world-side* skeptical hypotheses about which objects there are and what they are like in particular locations in the external world (disguised mules, etc.). But if a theory implies that propositions about which objects there are and what they are like in particular locations in the external world are knowable with no requirement of eliminating scenarios—that’s intolerable.³⁴

As perhaps the first epistemologist to postulate the $RS_{\exists V}$ condition, Stine [1976] seemed to embrace the vacuous knowledge consequence that I take to be damning; but since then epistemologists have recognized that there appears to be a serious problem that must be addressed.³⁵ I will now consider three replies to the vacuous knowledge objection to $RS_{\exists V}$, answering each.

(p.115) First Reply—knowledge of deeply contingent empirical truths does require epistemic work, but this “epistemic work” may involve something less than *eliminating* scenarios. Vogel [1999, 159n12] considers and rejects something like this reply: the RA theory that assumes $RS_{\exists V}$ “is committed to the thesis that one can know that an irrelevant alternative is false even though one can’t rule it out The RA theorist might still require that you have *some* minimal evidence against irrelevant alternatives in order to know that they are false. However, holding onto this scruple will make it more difficult, if not impossible, for the RA theorist to resist skepticism.”

Answer—in addition to the problem of skepticism noted by Vogel,³⁶ there is another problem. While having “minimal evidence” may not require eliminating $\neg P$ -scenarios, where P is the proposition to be known, does it not require eliminating *some* scenarios, perhaps as alternatives to related propositions? (Cf. §4.4.1 on inductive knowledge.) If it does, then we must reject $RS_{\exists V}$, since it allows agents to know deeply contingent truths with *no* requirement of eliminating scenarios.

Second Reply—the “double safety” theory is an $RS_{\exists V}$ theory that avoids the problem of vacuous knowledge. For even if one’s belief that Q is *vacuously safe*, in virtue of the fact that Q is true throughout the set $R(w)$ of nearby scenarios, it is not vacuously *adherent*, since it is an epistemic achievement that in all of the nearby scenarios where Q is true, the adherent agent believes Q .³⁷

Answer—adherence doesn’t help. Kripke [2011] showed that if an agent’s belief that P is sensitive, then normally her belief that P *and I believe that* P will be both sensitive and adherent. Kripke rightly concludes that adherence “is almost without force, a broken reed. What can be the point of a condition whose rigor can almost always be overcome by conjoining ‘and I believe (via M) that p ’ ... ?” (184). A similar point applies to the adherence part of double safety. Suppose an agent’s belief that P is vacuously safe, since there are no $\neg P$ -scenarios among the nearby scenarios. It follows by an argument similar to Kripke’s that the agent’s belief that P *and I believe that* P will normally be double-safe, even if her belief that P is not. So on the

double safety theory, it is normally sufficient to know that *P* and *I* believe that *P* that one has a vacuously safe belief that *P*. Thus, double safety does not solve the problem of vacuous knowledge, but only relocates it.

Third Reply—allowing knowledge of deeply contingent empirical propositions with no requirement of eliminating scenarios may seem bad, but it's alright, because “[s]imply *mentioning* any particular case of this knowledge, aloud or even in silent thought, is a way to ...create a context in which it is (p.116) no longer true to ascribe the knowledge in question to yourself or others” (Lewis 1996).

Answer—there are a number of problems with this reply, three of which I will discuss:

First, there is a *motivation* problem. When Stine [1976] first posited $RS_{\exists\forall}$, the motivation was clear: defend closure from Dretske. But then when faced with the problem that $RS_{\exists\forall}$ leads to vacuous knowledge, Lewis [1996] appeals to a super-shifty version of contextualism, according to which whenever you try to claim the vacuous knowledge that is rightly yours according to closure, context change invariably prevents you from truly claiming it (so Lewis concedes that “Dretske gets the phenomenon right” (564) after all). Sure, you can endorse a fixed-context closure principle in the abstract, but be careful not to instantiate it with any *specific* propositions and trigger an instant, irresistible change in context! But with closure made impotent in this way, was it worth getting into this vacuous knowledge mess to defend it? As Dretske [2005, 19] observes of super-shifty contextualism, “it is a way of preserving closure for the heavyweight implications while abandoning its usefulness in acquiring knowledge of them,”³⁸ or rather, while abandoning its usefulness in reasoning about agents’ knowledge of them—a bad trade for the problem of vacuous knowledge. Moreover, if one wants to stick with super-shifty contextualism and fixed-context closure, one can do so without being committed to vacuous knowledge, using the *multipath picture* proposed below (see §4.4.2).

Second, there is a *mechanism* problem. Most contemporary contextualists do not think that sayings or thinkings invariably introduce relevant counter-possibilities as Lewis claims.³⁹ So it is unclear what general mechanism would prevent those who have vacuous knowledge from sometimes truly claiming that they do. If this is so, then Lewis’s “unclaimable knowledge” reply collapses.

Third, there is a *missing-the-point* problem. What is problematic about vacuous knowledge is not just that agents could *truly claim* to have it—which they probably could according to post-Lewisian contextualism—but rather that they could *have it* at all. Cohen [2000, 105] correctly sees this: “it looks as if the [RS??] contextualist is committed to the view that we have contingent *a priori* knowledge. And of course, these cases do not fit the structure of the reference-fixing cases called to our attention by Kripke. Of course, I am not entirely happy with this result.” Cohen concludes that this is a “bullet” he is “prepared to bite” (106). But contextualists need not bite this bullet if they (p.117) opt for a contextualist version of the multipath picture of knowledge to be introduced in §4.3.

So much for $RS_{\exists\forall}$ theories then. On to $RS_{\forall\exists}$ theories. $RS_{\forall\exists}$ theories that take $r(P, w)$ to be the set of closest $\neg P$ -scenarios according to some kind of ordering avoid the problem of vacuous knowledge. In fact, they satisfy the general noVK (no vacuous knowledge) principle in Theorem 1, which says that if P is (deeply) contingent, then knowing P requires eliminating some scenarios. This is one of Heller’s [1999] main arguments for his $RS_{\forall\exists}$ theory over $RS_{\exists\forall}$ theories.

Unfortunately, the ordering-based $RS_{\forall\exists}$ theories that avoid the problem of vacuous knowledge face what I call the *problem of containment*. While it may be a virtue that these theories invalidate controversial multi-premise closure principles like closure under known implication, it is not a virtue that they allow closure failures to spread far beyond those controversial principles, to uncontroversial single-premise closure principles. Nozick [1981, 228] was well aware that even such a weak closure principle as $K(P \wedge Q) \Rightarrow KP$ is invalid according to his theory. He resisted the idea that $KP \Rightarrow K(p \vee Q)$ is invalid, but his theory clearly invalidates it (see Holliday 2014a and the Appendix).

In Holliday 2014a, I systematically investigate this problem of containment for a family of what I call “subjunctivist-flavored” theories, including basic version of the RA, sensitivity/tracking, and safety theories. The main Closure Theorem gives an exact characterization of the closure properties of knowledge according to these theories. Surprisingly, it turns out that despite the differences within the family of subjunctivist-flavored theories, the valid epistemic closure principles are essentially the same for these different theories. The problem is that these theories allow egregious

failures of single-premise closure, failures of principles as weak as $K(P \wedge Q) \Rightarrow K(P \vee Q)$ and $(KP \& KQ) \Rightarrow K(P \vee Q)$.

The source of the problem with the ordering-based $RS_{\forall\exists}$ theories is that they do not satisfy a necessary condition for single-premise closure under TF-consequence (recall (1) in §2.1):

TF-cover—if Q is a TF-consequence of P , then $r_c(Q, w) \subseteq r_c(P, w)$,

which says that the empirical work needed to know P covers the empirical work needed to know Q . One can easily check that if $r_c(S, w)$ is always the set of closest $\neg S$ -scenarios according to an ordering, then r does not satisfy TF-cover, which explains the failures of single-premise closure.

Is there any way to avoid these problems of containment and of vacuous knowledge?

2.5. An Impossibility Result

In the standard alternatives picture, it is impossible to avoid both problems from §4.2.4, even if we restrict our attention to a limited domain of propositions. (p.118) Call a set Σ of propositions an *area* iff whenever $P \in \Sigma$ and Q is a TF-consequence of P , then $Q \in \Sigma$. Then we have the following result.

Proposition 1. For any scenario w , context c , and area Σ , the following principles are *inconsistent* in the standard alternatives picture:

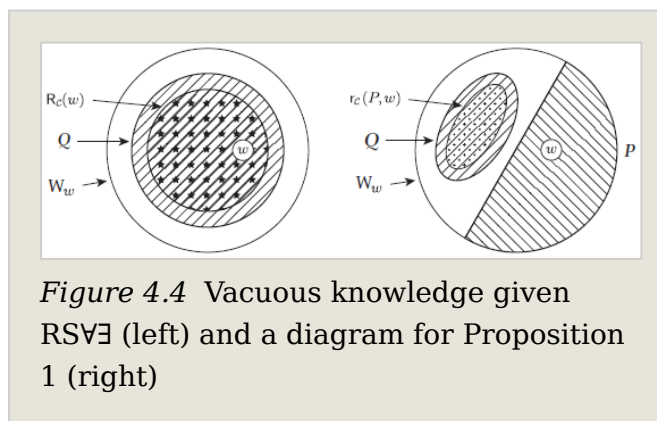


Figure 4.4 Vacuous knowledge given $RS_{\forall\exists}$ (left) and a diagram for Proposition 1 (right)

- contrast/enough₂— $\forall P \in \Sigma: r_c(P, w) \subseteq W - P$;
- e-fallibilism Σ — $\exists P \in \Sigma \exists Q \in P: r_c(P, w) \subseteq Q$ and $W_w - P \not\subseteq Q$;
- noVK Σ — $\forall P \in \Sigma: P_w \neq W_w$ implies $r_c(P, w) \neq \emptyset$;
- TF-cover₂— $\forall P, Q \in \Sigma$: if Q is a TF-consequence of P , then $r_c(Q, w) \subseteq r_c(P, w)$.

Here is the essence of the proof: by e-fallibilism Σ there are propositions P and Q as on the right side of Fig. 4.4 (where Q may overlap with P). Consider $P \vee Q$ and the set $P \vee Q$, which is the union of the two regions, P and Q , with diagonal lines. Where should we draw

$r_c(P \vee Q, w)$? Since $P \vee Q$ is a TF-consequence of P , TF-cover Σ requires that $r_c(P \vee Q, w)$ be a subset of $r_c(P, w)$; but contrast/enough Σ requires that $r_c(P \vee Q, w)$ be a subset of the *blank* region. The only way both can hold is if $r_c(P \vee Q) = \phi$. But this contradicts noVK, given that $P \vee Q$ does not include all of Ww .

Proof. By e-fallibilism Σ , there are propositions $P \in \Sigma$ and $Q \in P$ such that

$$(4) \quad r_c(P, w) \subseteq Q$$

and

$$(5) \quad W_w - P \not\subseteq Q.$$

Since $P \vee Q$ is a TF-consequence of P , $P \in \Sigma$ implies $P \vee Q \in \Sigma$, and TF-cover Σ implies

$$(6) \quad r_c(P \vee Q, w) \subseteq r_c(P, w),$$

which with (4) implies

$$(7) \quad r_c(P \vee Q, w) \subseteq Q \subseteq P \cup Q.$$

(p.119) However, contrast/enough Σ implies

$$(8) \quad r_c(P \vee Q, w) \subseteq W - (P \vee Q) = W - (P \cup Q),$$

which with (7) implies

$$(9) \quad r_c(P \vee Q, w) = \phi$$

Finally, (5) implies

$$(10) \quad (P \vee Q)_w \neq W_w,$$

which with noVK Σ implies

$$(11) \quad r_c(P \vee Q, w) \neq \phi,$$

which contradicts (9).

Note that Proposition 1 does not use the full strength of TF-cover, but only its instance $KP \Rightarrow K(P \vee Q)$. Also note that we could get the same result using $KP \Rightarrow ((P \vee Q) \wedge (P \vee \neg Q))$ and $K(S \wedge R) \Rightarrow KS$. In any case, I agree with Dretske [1970, 1009], Kripke [2011, 202], and Nozick [1981,

230n64] (not what his theory says, but what he says) that $KP \Rightarrow K(P \vee Q)$ should not fail.⁴⁰

The only principle in Proposition 1 that I have not yet defended is contrast/enough. All of the theories discussed in §4.2.3 satisfy this principle, but can we escape Proposition 1 by giving up contrast/enough? Not in the standard alternatives picture (but see §4.3.3). The reason is that in the standard alternatives picture, giving up contrast/enough means claiming that there are (p.120) some propositions P such that it is *necessary* in order to know P that one eliminate some P-scenarios. But if anything is sufficient for knowing P (as far as empirical work goes), it is eliminating *all* \neg P-scenarios, a kind of epistemic supererogation. Suppose I were to say, “I agree that you’ve ruled out every possible way in which P could be false, but that’s not enough for you to know that P is true; you also have to rule out such-and-such ways in which P could be true.” This seems absurd.⁴¹

I take the impossibility result in Proposition 1 to show that there is something seriously wrong with the standard alternatives picture. Remember that it is not enough to escape this result to argue that *there are* some cases in which knowing a deeply contingent empirical proposition imposes no requirement of empirically eliminating scenarios. Rather, to escape this impossibility result, one would have to argue that there is *no area* of propositions knowledge of which requires empirical investigation and with respect to which we are very weak fallibilists maintaining a very weak closure principle. This strikes me as an incredible claim. Until a credible argument for this claim appears, I conclude that fallibilists must seek a replacement for the standard alternatives picture.

3. The multipath picture of knowledge

In this section, I propose a new framework for fallibilism that solves the problems raised for the standard alternatives picture in §4.2. I call it the *multipath picture of knowledge*.

Recall the starting point of the standard alternatives picture: for each proposition to be known, there is “a [*single*] set of situations each member of which *contrasts* with what is [to be] known ... and must be evidentially excluded if one is to know” [emphasis added] [Dretske, 1981, 373]. Against these *single alternative set* and *contrast* assumptions, I will argue:

- In some cases, there is no set of situations all of which must be excluded if one is to know a proposition P; instead, there are

multiple sets of situations (scenarios), such that if one is to know P , one must exclude *all of the situations in at least one of those sets* (p.121) .

- In some cases, it is sufficient (as far as empirical investigation goes) for an agent to know a proposition P that she *only eliminates non-contrasting scenarios in which P is true*

A key observation will be that while the *single alternative set* and *contrast* assumptions may seem plausible for propositions that are “atomic” from a truth-functional or quantificational perspective (but see §4.4), fallibilists should reject these assumptions for logically complex propositions.

3.1. Against the Single Alternative Set Assumption

Suppose that an agent wants to know whether $P \vee Q$ is true, where P and Q are contingent empirical propositions. Further suppose that $P \vee Q$ is in fact true. Then there are at least three paths by which the agent could come to know it: she could start eliminating $\neg P$ -scenarios, and if she comes to know P , then she is done (at least with empirical investigation); or she could start eliminating $\neg Q$ -scenarios, and if she comes to know Q , then she is done (with empirical investigation); or she could come to know $P \vee Q$ without coming to know which disjunct is true, perhaps by eliminating all $(\neg P \wedge \neg Q)$ -scenarios without eliminating any $(\neg P \wedge Q)$ -scenarios or any $(P \wedge \neg Q)$ -scenarios. This is hardly a novel observation. But it raises the question of why any fallibilist should think that for a proposition like $P \vee Q$, there is a *single* set of scenarios that must be evidentially excluded if one is to know $P \vee Q$. It seems instead that there may be at least *three* sets of scenarios such that if one is to know $P \vee Q$, one must evidentially exclude all of the scenarios in at least one of those three sets, corresponding to the three paths to knowledge of $P \vee Q$ described above.

If we were *infallibilists*, there would be no need for these multiple “alternative sets” for $P \vee Q$. According to infallibilism, coming to know P requires eliminating all $(\neg P \wedge \neg Q)$ -scenarios; so does coming to know Q ; and so does coming to know $P \vee Q$ without coming to know which disjunct is true. Moreover, as argued in §4.2.5, eliminating all contrasting scenarios should be *enough* to know a proposition. Thus, infallibilists need only consider one alternative set for $P \vee Q$: to know $P \vee Q$ it is necessary and sufficient (as far as empirical work goes) that one eliminate all $(\neg P \wedge \neg Q)$ -scenarios.

But we are *fallibilists*. According to fallibilism, coming to know P might not require eliminating *all* $(\neg P \wedge \neg Q)$ -scenarios, at least not for every Q. Indeed, it might not require eliminating *any* $(\neg P \wedge \neg Q)$ -scenarios.⁴² But then since it is enough to know $P \vee Q$ that one eliminate all $(\neg P \wedge \neg Q)$ -scenarios, it is immediate that we need multiple alternative sets for $P \vee Q$, corresponding to the multiple paths to knowing $P \vee Q$ above: the scenarios that one must eliminate in order to know P may be different from those that one must eliminate in order to know Q, which may be different from those that one must eliminate in order to know $P \vee Q$ without knowing either disjunct. (p.122)

3.2. Multiple Alternative Sets

What §4.3.1 shows is that we should replace the r function of the standard alternatives picture, which assigns to each triple of a context c , proposition P , and scenario w , a single set $r_c(P, w) \subseteq W$ of scenarios, with a new “multipath” \mathbf{r} function that assigns to each such c , P , and w , a set

$$r_c(P, w) = \{A_1, A_2, \dots\}$$

of sets $A_i \subseteq W$ of scenarios. For example, for $P \vee Q$ we may have $r(P \vee Q, w) = \{A_1, A_2, A_3\}$, where A_1 is the set of scenarios to be eliminated in the path to knowing $P \vee Q$ via P ; A_2 is the set of scenarios to be eliminated in the path to knowing $P \vee Q$ via Q ; and A_3 is the set of scenarios to be eliminated in the path to knowing $P \vee Q$ without knowing either P or Q individually.

The foregoing points about disjunctive propositions apply to existential propositions as well. Assuming propositions can have quantificational structure as well as truth-functional structure, one could come to know $\exists x P(x)$ by coming to know $P(a)$, or by coming to $P(b)$, etc., or by coming to know $\exists x P(x)$ without coming to know $P(c)$ for any c . As a consequence of fallibilism, the alternative sets for these different paths to knowing $\exists x P(x)$ may be different. In this paper I concentrate on truth-functional structure, but a full treatment would include quantificational structure as well.

According to the *multipath picture of knowledge*, to know a proposition P , it is necessary and sufficient (as far as empirical elimination goes) that one eliminate all of the alternatives in *at least one* of the alternative sets for P (as on the right side of Fig. 4.5 with A_2):

$$\exists A \in r_c(P, w): A \cap U_c(P, w) = \emptyset, (\text{Knows})$$

where u is the same function as before.^{43, 44} As we shall see, this is just the generalization that fallibilists need in order to avoid the problems raised for them in the standard alternatives picture.

In explaining the multipath picture, I deliberately use the term “path to knowing” instead of “way of knowing.” There are often multiple “ways of knowing” a proposition in the sense that one can come to know it by eliminating a *single set* of alternatives in a number of ways: by sight, sound, smell, etc. I reserve the idea of multiple “paths to knowing” for the case in which for (p.123) a given proposition there are *multiple sets of alternatives* such that in order to know the proposition, it suffices to eliminate all of the alternatives in one of those sets, which one may often do in a number of *ways*.

The standard alternatives picture is equivalent to a special case of the multipath picture. Assuming

singlepath— $|r_c(P, w)|=1$,
 according to which each proposition has only one alternative set, we can move back and forth between the singlepath r function and multipath \mathbf{R} function as follows:

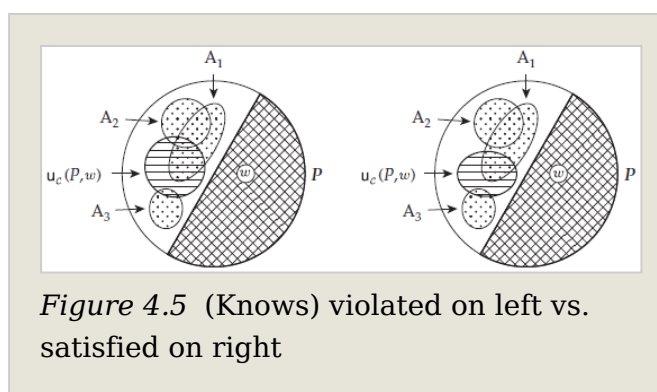


Figure 4.5 (Knows) violated on left vs. satisfied on right

$$r_c(P, w) = \{r_c(P, w)\}; \tag{12}$$

$$r_c(P, w) = \cup r_c(P, w). \tag{13}$$

It follows from these equations and **singlepath** that $r_c(P, w) \cap U_c(P, w) = \phi$ iff there is some $A \in r_c(P, w)$ such that $A \cap U_c(P, w) = \phi$, so (Knows) would be equivalent to (**Knows**).

Having rejected **singlepath** with the argument from disjunctive and existential propositions, let us consider multipath generalizations of singlepath principles. First, the singlepath principle

$$r\text{-Rof}A \text{---} w \notin P \text{ implies } w \in r_c(P, w)$$

from §4.2.1 generalizes to the multipath principle

$$r\text{-Rof}A \text{---} w \notin P \text{ implies } w \in \cap r_c(P, w),$$

which says that if P is false at w , then w is in every alternative set for P . As before, since w is always uneliminated for the agent in w , i.e. $w \in u_c(P, w)$ by u-RofA, only truths can be known.

Second, the singlepath principle

$$\text{ec-fallibilism—for some } P \text{ and } Q: r_c(P, w) \subseteq Q \text{ and } Q_w \subseteq W_w - P$$

from §4.2.1 generalizes to the multipath principle

$$\text{ec-fallibilism—for some } P, Q, \text{ and } A: A \in r_c(P, w), A \subseteq Q, \text{ and } Q_w \subseteq W_w - P,$$

(p.124) according to which there are propositions P and Q and a path to knowing P that only requires eliminating Q -scenarios, rather than all $\neg P$ -scenarios, giving us *expressible* fallibilism.

3.3. Against the Contrast Assumption

Next recall the contrast/enough assumption stated in the standard alternatives framework:

$$\text{contrast/enough—} r_c(P, w) \subseteq W - P.$$

In §2.5, I argued that the standard alternatives framework requires contrast/enough, because it should always be *enough* to know a proposition P that one eliminates *all* $\neg P$ -scenarios.

In the multipath alternatives framework, contrast/enough splits into two principles:

$$\text{contrast—} \forall A \in r_c(P, w): A \subseteq W - P;$$

$$\text{enough—} \exists A \in r_c(P, w): A \subseteq W - P.$$

As before, fallibilists should accept **enough**, which ensures that it is enough to know P that one eliminates all $\neg P$ -scenarios. Yet fallibilists should reject **contrast** and even the weaker principle

$$\text{semi-contrast—} \forall A \in r_c(P, w): A \neq \emptyset \text{ implies } A \cap (W - P) \neq \emptyset,$$

which says that every nonempty alternative set for P contains some $\neg P$ -scenario. Instead, we should allow one of the dotted alternative sets in Fig. 4.5 to overlap with or even be inside the crosshatched P -region.

The argument is simple. By **ec-fallibilism**,⁴⁵ there are propositions P and Q such that knowing P only requires eliminating Q -scenarios, where $Q_w \subseteq W_w - P$ and hence $(P \vee Q)_w \neq W_w$. But then since one path to knowing the contingent proposition $P \vee Q$ is via knowing P , and since knowing this P only requires ruling out Q -scenarios, which are of course $(P \vee Q)$ -scenarios, it follows that there is a path to knowing P

$\vee Q$ that only requires eliminating $(P \vee Q)$ -scenarios.⁴⁶ This contradicts **contrast** and **semi-contrast**.

(p.125) What may have fooled some fallibilists into assuming **contrast** for all propositions is that it may seem plausible when applied to logically atomic propositions. However, when we turn to the study of epistemic closure, we must consider logically complex propositions, for which universal **contrast** is not plausible from a fallibilist perspective. See the Appendix for further discussion of the relation between **contrast** and complex propositions.

In the disjunction counterexample to **contrast** assuming **ec-fallibilism**, one reason it makes sense for an alternatives set A for $P \vee Q$ to overlap with $P \vee Q$ (i.e. $A \cap P \vee Q \neq \emptyset$) is that A is also an alternative set for a *stronger* proposition, P , with which A does not overlap (i.e. $A \cap P = \emptyset$). One might think this is always the case when an alternative set for a proposition S overlaps with S :

$$\text{overlap.} - \forall A \in r_c(S, w): A \cap S \neq \emptyset \text{ implies } \exists P \in P: P_w \subseteq S, A \cap P = \emptyset, \text{ and } A \in r_c(P, w).$$

Nothing in my arguments turns on fallibilists accepting **overlap**, but it is noteworthy that **overlap** is consistent with all of the other principles I propose, as shown in the following section.

3.4. Problem Solved

Given the general arguments above, the multipath picture of knowledge should be attractive to all fallibilists. But I have yet to give one of the strongest arguments in its favor: it solves the problem represented by the impossibility result of Proposition 1.

First, observe that the singlepath principle

$$\text{noVK} - P_w \neq W_w \text{ implies } r_c(P, w) \neq \emptyset$$

from §§2.3-2.4 generalizes, following equation (12), to the multipath principle

$$\text{noVK} - P_w \neq W_w \text{ implies } \emptyset \notin r_c(P, w),$$

which also says that if P is (deeply) contingent, then knowing P requires eliminating scenarios.⁴⁷

Second, observe that the singlepath principle

$$\text{TF-cover} - \text{if } Q \text{ is a TF-consequence of } P, \text{ then } r_c(Q, w) \subseteq r_c(P, w)$$

from §2.4 generalizes to the multipath principle

$$\text{TF-cover} - \text{if } Q \text{ is a TF-consequence of } P, \text{ then } \forall A \in r_c(P, w) \exists B \in r_c(Q, w): B \subseteq A,$$

(p.126) which says that the empirical work for knowing P via *any* path covers the empirical work for knowing Q via *some* path. This principle is necessary for single-premise closure under TF-consequence.⁴⁸

It is now provable that all of the principles I have recommended for fallibilists are consistent in the multipath picture. Compare the negative Proposition 1 with the following positive result.

Proposition 2 (The Five Postulates). In the multipath picture, the following Five Postulates (for all w , P , and c) are *jointly consistent* with **ec-fallibilism**:

r-RofA— $w \notin p$ implies $w \in \cap r_c(p, w)$;

enough— $\exists A \in r_c(p, w): A \subseteq W - P$;

noVK— $P_w \neq W_w$ implies $\phi \notin r_c(p, w)$;

TF-cover—if Q is a TF-consequence of P , then
 $\forall A \in r_c(p, w) \exists B \in r_c(Q, W): B \subseteq A$;

overlap— $\forall A \in r_c(p, w): A \cap P \neq \phi$ implies $\exists Q \in P: Q_w \subseteq P, A \cap Q = \phi$, and
 $A \in r_c(Q, W)$.

Proof. The proposition holds as a corollary of Theorem 2 below, as explained in §3.5.

Although **ec-fallibilism** only says that we are fallibilists for at least one proposition, an **r**function can satisfy the Five Postulates while being fallibilistic for (infinitely) many propositions, as shown by Theorem 2 below. (Theorem 2 also shows that for **enough**, we could require $A \subseteq W_w - P$.)

It is important to understand why the multipath picture avoids an analogue of Proposition 1. Recall that the proof forced us to conclude in (11) that knowing the contingent proposition $P \vee Q$ does not require eliminating any scenarios; for if it did, then by contrast/enough it would require eliminating $(\neg P \wedge \neg Q)$ -scenarios; but that would contradict **TF-cover**, because knowing P did not require eliminating any $\neg Q$ -scenarios, by the very choice of Q as a proposition true in all of the relevant $\neg P$ -scenarios. Fortunately, the multipath picture does not lead to this contradiction. By **enough**, *one* path to knowing $P \vee Q$ is by eliminating all of the scenarios in some set of $(\neg P \wedge \neg Q)$ -scenarios, which is nonempty by **noVK**. But in line with **TF-cover**, *another* path to knowing $P \vee Q$ is via knowing P , which may involve eliminating only $(\neg P \wedge Q)$ -scenarios (note that here we use our rejection of both the

single alternative set and contrast assumptions). All of these paths require eliminating scenarios, so we respect **noVK**. We have no problem of vacuous knowledge.

Contrast this account with those of Nozick [1981] and Lewis [1996]. Let P be a mundane contingent proposition about the external world and S your favorite skeptical hypothesis incompatible with P . Recall that Nozick's tracking theory has the following problematic consequences: according to the theory, the logically weaker $P \vee \neg S$ may be much harder to know than the logically stronger P ; and the logically weaker $\neg S$ may be much harder to know than the logically stronger $P \wedge \neg S$. The reason is that on Nozick's theory, knowing P does not require eliminating skeptical ($\neg P \wedge S$)-scenarios, but knowing the weaker $P \vee \neg S$ does (where "elimination" for Nozick is understood as in §4.2.3); and knowing $P \wedge \neg S$ does not require eliminating skeptical S -scenarios, but knowing the weaker $\neg S$ does. This leads to the kind of extreme epistemic closure failures that illustrate the *problem of containment* from §4.2.4. As Vogel [2007, 76] explains:

It seems hard to deny that one's epistemic position with respect to a logically weaker proposition (X or Y) is at least as good as one's epistemic position with respect to a logically stronger proposition X The tracking condition T improperly inverts that relation by making the conditions for knowing (X or Y) more stringent than the conditions for knowing X [S]atisfying T with respect to (X or Y) can require that one is right over a greater region of logical space than is required to satisfy T with respect to X . Therefore, one's epistemic position with respect to (X or Y) may be inferior to one's epistemic position with respect to X .

While Nozick thereby makes knowing something like $P \vee \neg S$ too hard, Lewis [1996] makes it too easy. On Lewis's theory, there will be many contexts in which an agent can know the *contingent* $P \vee \neg S$ without any requirement of eliminating scenarios, simply because it is true throughout the fixed set of relevant possibilities (recall §4.2.4). Nozick and Lewis are pushed to these extreme positions by their assumption that for each proposition Q , there is only a single alternative set for Q , containing only contrasting $\neg Q$ -scenarios. By making such a single alternative set for $P \vee \neg S$ *nonempty*, Nozick avoids the problem of vacuous knowledge but saddles us with the problem of containment, whereas by making such a single alternative set for $P \vee \neg S$ *empty*, Lewis avoids the problem of containment but saddles us with the problem of vacuous knowledge.

We need not accept the Nozick-Lewis dilemma. In the multipath picture presented above, none of the alternative sets for the contingent $P \vee \neg S$ are empty, so there is no vacuous knowledge, and one of the alternative sets for $P \vee \neg S$ is from the path to knowing P via eliminating $(\neg P \wedge \neg S)$ -scenarios, so there is no problem of containment arising from $P \vee \neg S$. Nor is there a problem of containment arising from $P \wedge \neg S$. Like Lewis's theory but unlike Nozick's, in the multipath picture presented above, an agent who knows $P \wedge \neg S$ has done enough empirical work to know $\neg S$.

By establishing the consistency of fallibilism, **noVK**, **TF-cover**, and the other principles, Proposition 2 shows that by adopting the multipath picture of knowledge, fallibilists can avoid the problems raised in §4.2.5, a significant positive result. Of course, fallibilists who adopt the multipath picture must address the question: where do the possibly multiple alternative sets for a proposition come from? It may seem that fallibilists working with the standard alternatives picture have an easier time saying where their single set of alternatives for a proposition comes from, e.g. by using a relevance or similarity ordering of scenarios to pick out the set of close(st) scenarios where the (p.128) proposition is false. However, in §4.3.5 I show that the standard picture does not have an advantage in this respect.

3.5. From Singlepath to Multipath

The reason is that any standard alternatives function r determines a natural multipath alternatives function r^f ; and if r satisfies a few conditions, which are satisfied by any r based on orderings of scenarios as in §4.2.3, then r^f satisfies the Five Postulates of Proposition 2 and is fallibilistic in a way I will make precise. The alternative sets in $r^f_c(P, w)$ will depend on the *structure* of P . To keep things simple, I will first derive r^f from r for propositions in an easy-to-handle *normal form*; then we can immediately derive r^f for all propositions, using the fact that every proposition is TF-equivalent to one in normal form. To set this up, we need to review some basic logical concepts:

First, some notation and terminology. Assuming the structured propositions view of §4.1.1, let us write “ p ”, “ q ”, “ r ”, etc. for TF-atomic propositions. A *TF-basic* proposition is a TF-atomic proposition or the negation thereof. Let **basic-singlepath** and **basic-contrast** be the conditions **singlepath** and **contrast** from §4.3.2 and §4.3.3 applied to TF-basic propositions only. A *clause* is a disjunction of TF-basic propositions: e.g. $(p \vee \neg q \vee r)$. I assume that permutation and repetition of disjuncts does not matter, so “ $(p \vee \neg q \vee r)$ ” and “ $(\neg q \vee p \vee r)$ ” represent the same clause. A clause is *nontrivial* if it does not contain both p and $\neg p$ for any p . If a clause C' can be obtained by adding zero or more disjuncts to C , then C' is a *superclause* of C , and C is a *subclause* of C' : e.g. $(p \vee \neg q \vee r)$ is a superclause of $(p \vee \neg q)$ and a subclause of $(p \vee \neg q \vee \neg s \vee r)$. The set $\text{sub}(P)$ of TF-subpropositions of P is defined recursively: $\text{sub}(p) = \{p\}$ for p a TF-atomic proposition; $\text{sub}(\neg P) = \text{sub}(P) \cup \{\neg P\}$; $\text{sub}(P \# Q) = \text{sub}(P) \cup \text{sub}(Q) \cup \{P \# Q\}$ for any binary truth-functional connective $\#$, and so on for n -ary connectives. Finally, let $\text{at}(P)$ be the set of TF-atomic propositions in $\text{sub}(P)$.

Second, a fact: each proposition P (that is not a TF-tautology) is TF-equivalent to a proposition P' in *canonical conjunctive normal form* (CCNF), which is a conjunction of nontrivial clauses such that for each $q \in \text{at}(P')$, each clause in P' contains q or $\neg q$. Here is one way to calculate a CCNF of a proposition P . First, make a truth table for $\text{at}(P)$. Second, for each row of the truth table that makes the proposition *false*, write down a conjunction of TF-basic propositions describing that row; for example, the rows that make $p \wedge q$ false are described by: $(\neg p \wedge \neg q)$, $(\neg p \wedge q)$, and $(p \wedge \neg q)$. Third, write down a conjunction saying that we are not in any of those rows that make the proposition false: $\neg(\neg p \wedge \neg q) \wedge \neg(\neg p \wedge q) \wedge \neg(p \wedge \neg q)$. Finally, drive the negations inside: $(p \vee q) \wedge (p \vee \neg q) \wedge (\neg p \vee q)$. Thus, we obtain a CCNF

equivalent of $p \wedge q$. What is important for our purposes is that each proposition P (that is not a TF-tautology) is TF-equivalent to a P' in CCNF with $at(P) = at(P')$ that is *unique up to reordering of the conjuncts and disjuncts* (see Theorem 1.29 of Cori and Lascar 2000). Since order will not matter, let us associate with each (p.129) such P a unique CCNF(P) in CCNF. If P is a TF-tautology, let us stipulate that $CCNF(P) = (p \vee \neg p)$ for some atomic p .

Third, a definition using the notions above: for P in CCNF (not a TF-tautology), define $c_c(P)$ to be the set of all subclauses C of conjuncts in P such that every nontrivial superclause C' of C with $at(C') = at(P)$ is a conjunct of P . This implies that every conjunct of P is in $c_c(P)$, but there may be other clauses in $c_c(P)$. For example, if P is $(p \vee q) \wedge (p \vee \neg q)$, then $c_c(P) = \{(p \vee q), (p \vee \neg q), p\}$; and if P is the conjunction of $(p \vee q \vee r), (\neg p \vee q \vee r), (p \vee \neg q \vee r), (p \vee q \vee \neg r)$, and $(p \vee \neg q \vee \neg r)$, then $c_c(P)$ contains all of the conjuncts of P as well as $(p \vee q), (p \vee r), (q \vee r)$, and p . It turns out that $c_c(P)$ is the set of all nontrivial clauses C with $at(C) \subseteq at(P)$ that are TF-consequences of P .

Finally, some new notions related to fallibilism: a multipath function r is **fallibilistic** inat w with respect to P iff there is some $A \in r_c(P, w)$ with $A \subseteq W_w - P$; a standard alternatives function r is fallibilistic inat w with respect to P iff $r_c(P, w) \subseteq W_w - P$; and r is plurally fallibilistic inat w with respect to a set $\{P_1, \dots, P_n\}$ of clauses iff there are subclauses P_1', \dots, P_n' of P_1, \dots, P_n such that the union of all $r_c(P_i', w)$ sets is a strict subset of $W_w - (P_1 \wedge \dots \wedge P_n)$.

We are now ready to derive a multipath function rr from each standard alternatives function r . I will present the construction and the main result about the construction at the same time:

Theorem 2 (Multipath Theorem). Given a standard alternatives function r , define a multipath alternatives function rr as follows: for any clause C , define

$$r_c^r(C, w) = \{r_c(C, w) \mid C \text{ is a subclause of } C\} \quad (14)$$

for any CCNF conjunction $C_1 \wedge \dots \wedge C_n$ of clauses with $(C_1 \wedge \dots \wedge C_n) = \{D_1, \dots, D_m\}$, define

$$r_c^r(C_1 \wedge \dots \wedge C_n, w) = \left\{ A \subseteq W_w \mid \exists A_1 \in r_c^r(D_1, w) \dots \exists A_m \in r_c^r(D_m, w) : A = \bigcup_{1 \leq i \leq m} A_i \right\} \quad (15)$$

and if P is not in CCNF, define

$$r_c^r(P, w) = r_c^r(CCNF(P), w)$$

(16)

49

Then r^r satisfies **basic-singlepath** and **TF-cover**; if r satisfies **r-RofA**, then r^r satisfies **r-RofA**; if r satisfies contrast, then rr satisfies **basic-contrast**, **enough**, and **overlap**; if r satisfies noVK, then r^r satisfies **noVK**; for any clause P , if r is fallibilistic inat w with respect to P , then rr is **fallibilistic** inat w with respect to P ; and for any P in CCNF, if r is plurally fallibilistic inat w with respect to $c(P)$, then rr is **fallibilistic** inat w with respect to P .

Proof. See Appendix B of Holliday 2014a.

(p.130) The idea behind the definition of r^r_c is simple: (14) says that any path to knowing a subclause of a clause is a path to knowing the clause, a generalization of the idea that any path to knowing a disjunct of a disjunction is a path to knowing the disjunction; and (15) says that knowing a conjunction of clauses requires doing enough epistemic work to know each of the clauses that are TF-consequences of the conjunction. Note that for TF-basic propositions L , (14) implies $r^r_c(L, W) = \{r_c(L, W)\}$, so the derived multipath function rr differs from the input function r only for complex propositions.

For complex propositions P , it is not guaranteed that $r_c(P, w) \in r^r_c(P, w)$. To see this, one can check that for all $A \in r^r_c(p \wedge q, w)$, $r_c(p, w) \cup r_c(q, w) \subseteq A$, whereas there is no guarantee that $r_c(p, w) \cup r_c(q, w) \subseteq r_c(p, \wedge q, w)$, especially if r is based on orderings of scenarios. This is the source of the notorious problem for sensitivity theories that an agent may know that *the building is a barn and the building is red* ($p \wedge q$), despite not knowing that *the building is a barn* (p). By contrast, according to r^r , an agent knows a conjunction only if she has done enough to know each conjunct.

The definition of rr is best understood by example. Let us calculate the alternatives sets for $p \vee (q \wedge r)$. First, we calculate CCNF ($p \vee (q \wedge r)$) as above. The rows of the truth table for p , q , and r in which $p \vee (q \wedge r)$ is false are described by $(\neg p \wedge \neg q \wedge r)$, $(\neg p \wedge q \wedge \neg r)$, and $(\neg p \wedge \neg q \wedge \neg r)$, so

$$CCNF(p \vee (q \wedge r)) = (p \vee q \vee \neg r) \wedge (p \vee \neg q \vee r) \wedge (p \vee q \vee r)$$

One can then verify using the definition of c that

$$c(CCNF(p \vee (q \wedge r))) = \{(p \vee q \vee \neg r), (p \vee \neg q \vee r), (p \vee q \vee r), (p \vee r)\}$$

Call the members of this set D1–D5. By (15), every $A \in r^r_c(p \vee (q \wedge r))$ is of the form $A_1 \cup A_2 \cup A_3 \cup A_4 \cup A_5$, where by (14) each A_i is $r_c(C_i, w)$ for some subclause C_i of D_i .⁵⁰ The important choices of the subclauses C1–C5 of D1–D5 are: each C_i is p ; each C_i is q or r ; each C_i is $(p \vee q)$ or

$(p \vee r)$. These yield the following alternative sets in $r_c^r(p \vee (q \wedge r))$: $r_c(p, w)$ for the path via knowing p ; $r_c(q, w) \cup r_c(r, w)$ for the path via knowing $q \wedge r$; and $r_c(p \vee q) \cup r_c(q \vee r)$ for a path to knowing the disjunction without necessarily knowing either disjunct individually.

If r is based on orderings of scenarios, as in §4.2.3, then taking the $r_c(p \vee q) \cup r_c(p \vee r)$ path means eliminating the closest $(\neg p \wedge \neg q)$ -scenarios and the closest $(\neg p \wedge \neg r)$ -scenarios. By contrast, the singlepath picture with r says that there is only one path to knowing $p \vee (q \wedge r)$, by eliminating the closest $\neg(p \vee (q \wedge r))$ -scenarios, i.e. the closest $(\neg p \wedge (\neg q \vee \neg r))$ -scenarios. Note that each of these (p.131) scenarios is either a closest $(\neg p \wedge \neg q)$ -scenario or a closest $(\neg p \wedge \neg r)$ -scenario, so $r_c(p \vee (q \wedge r)) \subseteq r_c(p \vee q) \cup r_c(p \vee r)$.⁵¹

Let us now prove Proposition 2. Recall from Theorem 1 that if r is such that for all C, P , and w , $r_c(P, w)$ is the set of closest $\neg P$ -scenarios according to an ordering \leq_w^c as in §4.2.3, then r satisfies r -RofA, contrast, and noVK. Hence by Theorem 2, r^r satisfies the Five Postulates of Proposition 2. Moreover, r^r is highly fallibilistic if r is (not with respect to every proposition that r is, but with respect to those that meet the conditions in the theorem). To establish *expressible* contrast fallibilism, **ec-fallibilism**, let us make an extremely weak assumption about expressibility: for some TF-basic proposition L and proposition Q , Q entails $\neg L$, but not vice versa: $Q_w \subseteq W_w - L$. Then there exists an ordering \leq_w^c such that $\text{Closest } \leq_w^c(\neg L) = Q_w \subseteq W_w - L$; so by the fact that $\text{Closest } \leq_w^c(\neg L) = r_c(L, w) \in r_c^r(L, w)$, Q is a witness for the fact that r^r satisfies **ec-fallibilism**.

Although I have focused on the idea of deriving r^r from a function r based on the familiar qualitative orderings of scenarios, it is not necessary that the input function r be based on such orderings. If $w \notin p$, then we can assume $r_c(P, w) = \{w\}$, so r -RofA is satisfied. If $w \in P$, then perhaps $r_c(P, w)$ is some function of the probability, or cost-weighted probability, or other value of each $\neg P$ -scenario, so that the $\neg P$ -scenarios with relatively substantial probability, or cost-weighted probability, or whatever, relative to other $\neg P$ -scenarios, are in $r_c(P, w)$, where what “relatively substantial” means may depend on C, P , or w .⁵² These options would also satisfy contrast and noVK,⁵³ so the resulting (p.132) r^r would have the properties given by Theorem 2. I will not go into the details here. My main point in this section is that the multipath picture is not at a disadvantage relative to the

singlepath picture with respect to having more alternative sets for which to account.

In my view, the construction of rr from r above provides a kind of “lower bound” on what a multipath function should look like if derived from a singlepath function r : if $A \in r_c^r(p, w)$, then A should be an alternative set for P in w relative to according to any reasonable multipath function derived from r . In §4.4, I will consider whether a reasonable multipath function should provide even more alternative sets—even more paths—for knowing some propositions.

4. More Paths?

In §4.3, we saw what might be called the “conservative” version of the multipath picture. On the conservative version, the source of additional paths to knowledge of a proposition is the structure of the proposition itself; this is why the single alternative set and contrast assumptions are rejected for *complex* propositions. Let us now consider the questions: Are there additional paths to knowledge of a proposition that do not come from the structure of the proposition? Should the single alternative set and contrast assumptions be rejected in general, not just for complex propositions?

4.1. Inductive Closure

Recall that my motivating examples for the multipath picture in §4.3.2 involved cases where some of the multiple paths to knowing a complex proposition—such as a disjunctive or existential proposition—went via knowing logically *stronger* propositions—a disjunct or an instance. Might there be multiple paths to knowing a proposition via knowing logically *weaker* propositions? Anyone who thinks that *inductive* knowledge is possible is committed to an affirmative answer. Although so far I have concentrated on closure principles where the relation \mathbf{R} (recall §4.2.1) is a deductive relation, one can also consider closure with respect to inductive relations, asking whether an agent who knows the empirical premises of a “good” inductive argument has thereby done enough empirically to know the conclusion. Let us see how the multipath picture of knowledge bears on this issue.

To use a standard (oversimplified) example of enumerative induction, let $\mathbf{E} = \{e_1, \dots, e_n\}$ be the set of the first n emeralds, by distance, from some location; for any $e \in \mathbf{E}$, let G_e be the proposition that *e is green*; and let G be $\bigwedge_{e \in \mathbf{E}} G_e$, a conjunctive version of *all emeralds in \mathbf{E} are green*. According to some fallibilists, for large n one can come to know G by observing just the emeralds in some strict subset $E_1 \subseteq E$.

Since the proposition $\bigwedge_{e \in E_1} Ge$ is logically weaker than G , this answers the second part of the question above; but what about the multiplicity of paths? Presumably believers in inductive knowledge do not think (p.133) that G can only be known by observing the emeralds in just *one* set $E_1 \subseteq \mathbf{E}$; instead, there should be many sets $E_i \subseteq \mathbf{E}$ (with $E_i \not\subseteq E_j$ for $i \neq j$) such that if the agent observes all of the emeralds in one of them, she has done enough to know G . Hence for each such E_i , there will be an alternative set $A_i \in r_c(G, w)$ such that for every $e \in E_i$ there is some $B_e \in r_c(Ge, w)$ with $B_e \subseteq A_i$. Assuming $A_i \not\subseteq A_j$ for $i \neq j$, this gives us the multiple alternative sets, answering the first part of the question above. Indeed, this provides another reason to accept the multipath picture for fallibilists who wish to make room for the possibility of knowledge by enumerative induction.

Note that on the assumption of closure under single-premise TF-consequence, someone who comes to know G by observing the emeralds in some E_i should also be able to know that the so-far-unobserved emerald $b \in \mathbf{E}$ in my back pocket is green by observing those *other* emeralds in E_i ($b \notin E_i$). If this is correct, then it seems there may be multiple paths to knowing even the TF-atomic proposition G_b : by eliminating the $\neg G_b$ -scenarios in $r_c(G_b, w)$ or by eliminating the scenarios in some $A_i \in r_c(G, w)$. These paths will be genuinely distinct if $r_c(G_b, w) \not\subseteq A_i \not\subseteq r_c(G_b, w)$ (see footnote 44). If so, then basic-singlepath cannot hold for \mathbf{r} . Moreover, if A_i contains some G_b -scenarios, e.g. $(G_b \wedge \neg G_e)$ -scenarios for some $e \in E_i$, then **basic-contrast** cannot hold for \mathbf{r} either. Thus, fallibilists who wish to maintain the possibility of inductive knowledge and single-premise logical closure may be led to reject **basic-singlepath** and **basic-contrast**. One may then wonder whether such fallibilists can extend a singlepath function r to a multipath function \mathbf{r} as in §4.3.5. To do so, they must modify (14) in the construction for Theorem 2 in order to allow for some extra inductive paths to knowing some TF-basic propositions.⁵⁴ In the current example, to say just how many or which emeralds must be observed in the various paths to knowing G_b inductively, to determine the extra alternative sets in $r_c(G_b, w)$, is a topic for a theory of inductive knowledge. In general, presumably only propositions involving certain kinds of objects and (“projectible”) properties admit such extra inductive paths, which is again a topic for a theory of inductive knowledge to explain.⁵⁵

4.2. Metaphysical and Multi-Premise Closure

(p.134) Theorem 2 shows how closure under single-premise TF-consequence fits with the conservative view that additional paths to knowledge of a proposition come from the structure of the proposition itself. However, in order to guarantee more controversial closure principles, consistently with the Five Postulates of §4.3.4, one must go beyond the conservative view. This is easiest to see in the case of closure under single-premise metaphysical entailment, which requires the following assumption:

$$\text{M-cover} \text{---if } Pw \subseteq Q, \text{ then } \forall A \in rC(P, w) \exists B \in rC(Q, w): B \subseteq A$$

which is the multipath generalization of

$$\text{M-cover} \text{---if } Pw \subseteq Q, \text{ then } rC(Q, w) \subseteq rC(P, w)$$

M-cover says that if P entails Q as a matter of (deep) metaphysical necessity, then any path to knowing P covers a path to knowing Q , regardless of the structures of P and Q or what kinds of objects and properties they involve. Since the construction of rr in Theorem 2 only looks at the structure of P for extra paths to knowing P other than $r_c(P, w)$, it does not guarantee that **M-cover** will hold for r^x if **M-cover** does not hold for r . Moreover, by the impossibility result in Proposition 1, **M-cover** cannot hold for r together with the other conditions in the theorem, since **M-cover** implies **TF-cover**. In order to guarantee **M-cover**, along with the Five Postulates in §4.3.4, one must modify rr to allow extra paths to knowing P , not given by the structure of P or by $r_c(P, w)$.

Before discussing modifications, let us consider the desirability of the **M-cover** assumption. Dretske [1970, 2005] famously argues that it can take more epistemic work to know a Q metaphysically entailed by P than to know P itself, when Q has a “heavyweight” status compared to the “lightweight” status of P .⁵⁶ One of the Dretskean concerns is that **M-cover**/**M-cover** will lead to radical skepticism about knowledge. Let us try to derive this result in the standard alternatives picture and the multipath picture. As discussed in §4.2.1, for many empirical propositions P ,

$$uC(P, w) \cap (Ww - P) = \emptyset$$

so there are some uneliminated $\neg P$ -scenarios. Hence it is reasonable to assume that there is some proposition S (think of a “skeptical counter-hypothesis”) such that

$$\emptyset \neq S \subseteq uC(P, w) \cap (Ww - P)$$

If for every set of scenarios there is a proposition true in exactly that set of scenarios, then (18) is immediate from (17). Now let us suppose

that for at (p.135) least one of the propositions S as in (18), $\neg S$ is what could be called a *semi-contrast proposition*, in the sense that knowing $\neg S$ requires eliminating at least one S -scenario:⁵⁷

$$r_c(\neg S, w) \cap S \neq \phi$$

$$\forall B \in r_c(\neg S, w): B \cap S \neq \phi$$

It follows from $S \subseteq W_w - P$ in (18) that $P_w \subseteq \neg S$, so **M-cover**/**M-cover** implies

$$r_c(\neg S, w) \subseteq r_c(P, w)$$

$$\forall A \in r_c(P, w) \exists B \in r_c(\neg S, w): B \subseteq A$$

Together (21) and (19) imply $r_c(P, w) \cap S \neq \phi$ which with (18) implies $r_c(P, w) \cap U_c(P, w) \neq \phi$. Thus, by (**Knows**), the agent in w *does not know* P relative to. Similarly, (22) and (20) imply that for every $A \in r_c(P, w)$, $A \cap S \neq \phi$, which with (18) implies $A \cap u_c(P, w) \neq \phi$. Since this holds for *every* $A \in r_c(P, w)$, by (**Knows**) the agent in w *does not know* P relative to C . Since P , w , and C were arbitrary, we seem to be left with radical skepticism about empirical knowledge.

Essentially the same argument for skepticism can be given using other closure principles. I will demonstrate this in the multipath picture, leaving the singlepath case as an exercise for the reader. First, consider closure under *metaphysical equivalence* and the principle $K(P \wedge Q) \Rightarrow (KP \ \& \ KQ)$:⁵⁸

$$\mathbf{M-equiv} \text{--- if } P_w = Q_w, \text{ then } \forall A \in r_c(P, w) \exists B \in r_c(Q, w): B \subseteq A$$

$$\text{concover---} \forall A \in r_c(P \wedge Q) \exists B \in r_c(Q, w): B \cup B' \subseteq A$$

(p.136) **M-equiv** says that if P and Q are equivalent as a matter of (deep) metaphysical necessity, then any path to knowing P covers a path to knowing Q ; and **concover**, which follows from **TF-cover**, says that any path to knowing a conjunction covers paths to knowing each conjunct. It follows from $S \subseteq W_w - P$ in (18) that $P_w = (P \wedge \neg S)_w$, which with **M-equiv** and **concover** implies (22).⁵⁹ The rest of the skeptical argument goes exactly as before.

Finally, the argument can be given with closure under multi-premise TF-consequence:

$$\mathbf{Multi} \text{--- if } Q \text{ is a TF-consequence of } \{P_1, \dots, P_n\},$$

$$\forall A_1 \in r_c(P_1, w) \dots \forall A_n \in r_c(P_n, w) \exists B \in r_c(Q, w): B \subseteq \bigcup_{1 \leq i \leq n} A_i$$

so any paths to knowing P_1, \dots, P_n together cover a path to knowing Q . It follows from $S \subseteq W_w - P$ in (18) that $(\neg P \vee \neg S)_w = W_w$, so the disjunction $\neg P \vee \neg S$ is a (deeply) necessary truth. Some would conclude that knowing $(\neg P \vee \neg S)$ does not require empirically eliminating scenarios, i.e. $\phi \in r_c(\neg P \vee \neg S)$, but let us only make the

weaker assumption that there is a path to knowing $(\neg P \vee \neg S)$ that does not require eliminating alternatives for $\neg S$,⁶⁰ i.e. some $A_1 \in r_c(\neg P \vee \neg S)$ that does not overlap with any $B \in r_c(\neg S, w)$. Since $\neg S$ is a TF-consequence of $(\neg P \vee \neg S)$ together with P , **Multi** with $P_1 = (\neg P \vee \neg S)$ and $P_2 = P$ implies that for every $A_2 \in r_c(P, w)$, there is some $B \in r_c(\neg S, w)$ such that $B \subseteq A_1 \cup A_2$ and hence $B \subseteq A_2$ by the choice of A_1 ; and this implies (22), which leads to skepticism as before.

What are our options for avoiding this kind of argument for radical skepticism?

I have already mentioned the Dretskean option of denying closure under single-premise metaphysical entailment. The same considerations about lightweight propositions entailing heavyweight propositions suggest that Dretske would reject closure under metaphysical *equivalence* as well; for if $\neg S$ is a heavyweight proposition compared to the lightweight P , then surely $P \wedge \neg S$ is heavyweight as well. The construction in Theorem 2 is compatible with this view: without further assumptions about r or about how r^r arises from r , there may be an alternative set in $r_c^r(P, w)$ that does not cover any in $r_c^r(P \wedge \neg S, w)$,⁶¹ even if P and $P \wedge \neg S$ are metaphysically equivalent. If we had assumed that propositions are sets of metaphysically possible scenarios or worlds, then **M-equiv** would basically be unavoidable, but for generality I have not assumed such a view (recall §4.1.1).

(p.137) As for multi-premise closure, without further assumptions about r or about how r^r arises from r , Theorem 2 does not guarantee that r^r satisfies **Multior**, as a special case, $(KP \ \& \ KP') \Rightarrow K(P \ \& \ P')$. The reason is that someone who knows P according to r^r , so has done enough work to know every $C \in c(P)$,⁶² and knows P' according to r^r , so has done enough work to know every $C \in c(P')$, has not necessarily done enough work to know $P \wedge P'$ according to r^r , because there may be some $D \in c(P \wedge P')$ that is not a superclause of anything in $c(P)$ or $c(P')$; in Dretskean terms, D may be a new “heavyweight” TF-consequence of $P \wedge P'$, which neither P nor P' had individually. Theorem 2 does guarantee that if an agent knows two TF-atomic (or TF-basic) propositions p and p' , then she has done enough empirically to know $p \wedge p'$; every $D \in c(p \wedge p')$ is a superclause of something in $c(p)$ or $c(p')$, so the problem of new heavyweight consequences does not arise.⁶³ But if P and P' are complex, then the aggregation principle is not guaranteed without further assumptions, given the possibility of new heavyweight consequences coming from the combination of P and P' . On this view, it is not necessarily harmless to combine P and P' with

\wedge ; the impression that nothing more is required to know $P \wedge P'$ ("just put the \wedge in between!") may be an illusion induced by too much focus on syntax. The same points apply to closure under known implication:⁶⁴ p and $p \rightarrow q$ together have what may be a heavyweight TF-consequence, q , that neither has individually.⁶⁵

Much more could be said about views that limit the scope of closure. But let us change gears: is it possible to reject the skeptical argument while defending the strong closure principles? The only real way to do so is to maintain that for *every* S as in (18) for a known P , $\neg S$ can be known without a requirement of eliminating S -scenarios. In the standard alternatives picture, this would force defenders of strong closure to say that every such contingent $\neg S$ can be known without *any* requirement of eliminating scenarios, i.e. $r_c(\neg S, w) = \phi$, which is the problem of vacuous knowledge from §4.2.4. For if any scenarios had to be eliminated, they would be S -scenarios according to the contrast/enough condition that I have argued is built in to the standard alternatives picture (§4.2.5).

(p.138) However, in the multipath picture, defenders of strong closure can say that knowing $\neg S$ *does* require eliminating scenarios: a hard path to fulfilling this requirement is to eliminate some nonempty set of skeptical S -scenarios, in line with **enough** and **noVK** from §§4.3.3–4.3.4; but another path is to go via an ordinary proposition P that entails $\neg S$, eliminating all of the scenarios in some set $A \in r_c(P, w)$ of $(\neg P \wedge \neg S)$ -scenarios, rejecting **semi-contrast** for $\neg S$, but consistently with **overlap** from §4.3.3. That's why not just *anyone* gets to know the contingent $\neg S$, but someone who did the epistemic work to know an ordinary P that entails $\neg S$ can.⁶⁶ This is certainly an improvement over the vacuous knowledge story. What it shows, I think, is that the issue of how far closure holds ultimately comes down to the question of how far **contrast/semi-contrast** fails. In particular, since there is no guarantee that S will be complex, defenders of strong closure must reject **basic-contrast**, **basic-semi-contrast**, and **basic-singlepath**.

We have seen that defenders of strong closure benefit from the multipath picture. Can they also view a multipath function r^r as arising from a singlepath function r ? The simplest way to do so is to replace

$$r_c^r(C, w) = \{r_c(C, w) \mid C \text{ is a subclause of } C\}$$

from Theorem 2 with

$$r_c^r(P, w) = \{r_c(P, w) \mid P_w \text{ is a sub } r \text{ of } P\}$$

Then clearly r^f satisfies **M-cover**; if r satisfies r -RofA, then r^f satisfies **r-RofA**; if r satisfies contrast, then r^f satisfies **enough** and **overlap**; if r satisfies noVK, then r^f satisfies **noVK**; and if r satisfies alpha (recall §4.2.3), then r^f satisfies the analogous multipath principle,

$$\text{alpha}—\forall A_1 \in r_c(P, w) \forall A_2 \in r_c(Q, w) \exists B \in r_c(P \wedge Q, w): B \subseteq A_1 \cup A_2$$

for $((KP \ \& \ KP) \Rightarrow K(P \wedge P))$. It follows that r^f guarantees full single and multi-premise closure.

I will not try to decide here between the two positions on closure outlined above. In essence, defenders of strong closure think that knowledge is easier to come by than do defenders of limited closure. Does the former camp make knowledge too cheap? Without answering this question, we can say (p.139) this much: at least in the multipath picture, no fallibilist need be committed to the cheapest knowledge of all—the vacuous knowledge of the standard alternatives picture in §4.2.4.

5. conclusion

There are multiple paths to knowing some propositions about the world. This sounds like a truism, but it has yet to be fully appreciated in the theory of knowledge. According to the *standard alternatives picture* assumed in so much fallibilist epistemology, knowing a proposition involves eliminating a single set of alternatives. Proposition 1 in §4.2.5 suggests that this picture is fundamentally flawed. In its place, I proposed a *multipath picture* of knowledge for fallibilists, according to which knowing a proposition involves eliminating all of the alternatives in one of the proposition's alternative sets, of which there may be many. Proposition 2 in §4.3.4 showed that this picture solves the problems raised by Proposition 1 for the standard alternatives picture. Moreover, the Multipath Theorem in §4.3.5 showed how the multiple alternative sets for a proposition may emerge out of the standard alternatives picture in a way that depends on the structure of the proposition. Unlike the standard alternatives picture, the multipath picture allows fallibilists to maintain uncontroversial (single-premise, logical) epistemic closure principles without having to make extreme assumptions about the ability of humans to know empirical truths without empirical investigation. It also offers benefits to those who endorse more controversial (multi-premise and metaphysical) closure principles, thereby taking a more liberal attitude about paths to knowledge. Hard questions remain about how far fallibilists should claim that closure goes. But nobody ever said being a fallibilist was easy.

Appendix: Negation, Contrast, And Closure

The rejection of the single alternative set assumption in §4.3.1 and the contrast assumption in §4.3.3 can help us make sense of an otherwise puzzling feature of Nozick and Dretske's views on closure, concerning the following closure principles:

$$KP \Rightarrow K \neg (\neg P \wedge S)$$

$$K \neg P \Rightarrow K \neg (P \wedge S)$$

Beginning with Nozick [1981, 228f.], he explicitly rejects (24): "it is possible for me to know p yet not know the denial of a conjunction, one of whose conjuncts is not- p . I can know p yet not know ...not-(not- p &SK) However, we have seen no reason to think knowledge does not extend across known logical equivalence." Only a page later Nozick [1981, 230] writes: "It seems that a person can track 'Pa' without tracking 'there is an x such that Px '. But this apparent nonclosure result surely carries things too far. As would the apparent result of nonclosure under the propositional calculus rule of inferring (p.140) ' p or q ' from ' p '..."⁶⁸ Let us write the latter principle as $KP \Rightarrow K(P \vee Q)$. What is interesting is that Nozick's views in these passages are inconsistent.⁶⁹ Surely Nozick knows that $P \vee \neg S$ is logically equivalent to $\neg(\neg P \wedge S)$, so given his endorsement of closure under known logical equivalence, if he knew $P \vee \neg S$ then he would know $\neg(\neg P \wedge S)$. But he says he does not know $\neg(\neg P \wedge S)$, so he must not know $P \vee \neg S$. But he also says he knows P and accepts $KP \Rightarrow K(P \vee Q)$, so he should know $P \vee \neg S$. (We can assume Nozick makes the relevant inferences.)

I do not think this inconsistency was simply a mistake. Instead, I suspect that it reflects an intuition that Nozick shares with others, including Dretske. While Nozick explicitly endorses $KP \Rightarrow K(P \vee Q)$ and explicitly rejects (24), Dretske explicitly endorses $KP \Rightarrow K(P \vee Q)$ and is committed to rejecting (25). First, Dretske [1970] says that "it seems to me fairly obvious that if someone ...knows that P is the case, he knows that P or Q is the case" (1009). Second, Dretske [1970, 1015–16] claims in his famous zebra that the zoo visitor does not know that the animal in the zebra cage is not a mule (M) disguised to look like a zebra (D): $\neg K \neg (M \wedge D)$. But I assume that as a strong fallibilist, Dretske will allow that in ordinary cases of observing zebras at the zoo, people who know the difference between zebras and mules know that the zebras are not mules: $K \neg M$. But together these commitments force Dretske to deny (25). Then since Dretske endorses $KP \Rightarrow K(P \vee Q)$, an instance of which is $K \neg M \Rightarrow K(\neg M \vee \neg D)$, Dretske must deny $K(\neg M \vee \neg D) \Rightarrow K \neg (M \wedge D)$.

Thus, Dretske must deny the “De Morgan” closure principle $K(\pm P \vee \pm Q) \Rightarrow K \neg(\mp P \wedge \mp Q)$,⁷⁰ and there is pressure for Nozick to do the same to resolve the inconsistency in his views.

In my view, (24) and (25) seem problematic because their consequents claim knowledge that something is *not* the case, and this *negation* brings with it the idea of **contrast** that I have argued fallibilists should not accept in general. In particular, I argued that **contrast** can fail for disjunctions like $P \vee \neg S$; for I agree with Dretske, Nozick, and Kripke that one path to knowing $P \vee \neg S$ is via knowing P , and I agree with fallibilists in general that coming to know P may not require eliminating $(\neg P \wedge S)$ -scenarios. But can one come to know $\neg(\neg P \wedge S)$ without eliminating $(\neg P \wedge S)$ -scenarios?

With the negated conjunction, I expect some people’s intuitions to shift in favor of **contrast**, perhaps because the processing of negations in non-epistemic contexts in natural language involves the construction of contrast classes (see Oaksford and Stenning 1992). In Dretske’s example, when considering a disjunction like $\neg M \vee \neg D$, one may recognize that knowing $\neg M$ provides a path to knowing the disjunction; but when considering the equivalent $\neg(M \wedge D)$, one might have a competing intuition in favor of **contrast** and the thought (p.141) that $(M \wedge D)$ -scenarios must be eliminated, which fallibilists would not insist on for knowing $\neg M$. (One might also have the mistaken intuition that $\neg D$ follows from $\neg(M \wedge D)$, so D -scenarios must be eliminated.)⁷¹ There are three ways the explanation might go from here, depending on the kind of significance assigned to these intuitions:

1. *Epistemic*: $K(\pm P \vee \pm Q) \Rightarrow K \neg(\mp P \wedge \mp Q)$ is not a valid principle even for a fixed context, because **contrast** may apply to $r(\neg(\pm P \vee \pm Q), w)$ without applying to $r_c(\neg(\mp P \wedge \mp Q), w)$.
2. *Pragmatic*: $K(\pm P \vee \pm Q) \Rightarrow K \neg(\mp P \wedge \mp Q)$ is valid, but when an attributor claims that an agent knows a negated proposition N , this has a tendency to pragmatically trigger the (mistaken) intuition that **contrast** must hold for $r_c(N, w)$.
3. *Contextual*: $K(\pm P \vee \pm Q) \Rightarrow K \neg(\mp P \wedge \mp Q)$ is valid for a fixed context, but when an attributor claims that an agent knows a negated proposition N , this has a tendency to change the context to one in which **contrast** holds for $r_c(N, w)$ (cf. DeRose 1995).

In my view, the pragmatic or contextual explanations are more plausible than the epistemic, although I will not argue for this here.⁷²

The point I wish to make is that the multipath picture has the potential to explain divergent intuitions concerning knowledge of disjunctions and knowledge of equivalent negated conjunctions in terms of the key ideas of §4.3.1 and §4.3.3.

References

Bibliography references:

Alspector-Kelly, Marc. 2011. Why Safety Doesn't Save Closure. *Synthese*, 183 (2):127-42.

Arrow, Kenneth J. 1959. Rational Choice Functions and Orderings. *Economica*, 26(102):121-7.

Barwise, Jon. 1981. Scenes and Other Situations. *Journal of Philosophy*, 78(7): 369-97.

Barwise, Jon and John Perry. 1983. *Situations and Attitudes*. MIT Press.

Blome-Tillmann, Michael. 2009. Knowledge and Presuppositions. *Mind*, 118 (470):241-94.

Bordes, Georges. 1976. Consistency, Rationality and Collective Choice. *The Review of Economic Studies*, 43(3):451-7. (p.142)

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

Cohen, Stewart. 1998. Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery. *Australasian Journal of Philosophy*, 76(2): 289-306.

Cohen, Stewart. 2000. Contextualism and Skepticism. *Philosophical Issues*, 10:94-107.

Comesaña, Juan. 2007. Knowledge and Subjunctive Conditionals. *Philosophy Compass*, 2(6):781-91.

Cori, René and Daniel Lascar. 2000. *Mathematical Logic: A Course with Exercises Part I*. Oxford University Press.

Davies, Martin and Lloyd Humberstone. 1980. Two Notions of Necessity. *Philosophical Studies*, 38(1):1-30.

DeRose, Keith. 1995. Solving the Skeptical Problem. *The Philosophical Review*, 104(1):1-52.

- DeRose, Keith. 2000. How Can We Know that We're Not Brains in Vats? *The Southern Journal of Philosophy*, 38(S1):121-48.
- DeRose, Keith. 2009. *The Case for Contextualism*. Oxford University Press.
- Dretske, Fred. 1970. Epistemic Operators. *The Journal of Philosophy*, 67(24): 1007-23.
- Dretske, Fred. 1971. Reasons, Knowledge, and Probability. *Philosophy of Science*, 38(2):216-20.
- Dretske, Fred. 1981. The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40(3):363-78.
- Dretske, Fred. 2005. The Case against Closure, in Matthias Steup and Ernest Sosa, editors, *Contemporary Debates in Epistemology*, pages 13-25. Wiley-Blackwell.
- Dretske, Fred. 2006. Information and Closure. *Erkenntnis*, 64:409-13.
- Evans, Gareth. 1979. Reference and Contingency. *The Monist*, 62(2): 161-89.
- Goldman, Alvin I. 1976. Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73(20):771-91.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Heller, Mark. 1989. Relevant Alternatives. *Philosophical Studies*, 55(1):23-40.
- Heller, Mark. 1999. Relevant Alternatives and Closure. *Australasian Journal of Philosophy*, 77(2):196-208.
- Holliday, Wesley H. 2012. *Knowing What Follows: Epistemic Closure and Epistemic Logic*. PhD thesis, Stanford University. Revised Version, ILLC Dissertation DS-2012-09.
- Holliday, Wesley H. 2014a. Fallibilism and Multiple Paths to Knowledge (Extended Version). URL <<http://philosophy.berkeley.edu/holliday>>.
- Holliday, Wesley H. 2014b. Epistemic Closure and Epistemic Logic II: A New Framework for Fallibilism. Manuscript.

- Holliday, Wesley H. Forthcoming. Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. *Journal of Philosophical Logic*. Published online <http://link.springer.com/article/10.1007/s10992-013-9306-2?sa_campaign=email/event/articleAuthor/onlineFirst>, DOI 10.1007/s10992 -013-9306-2.
- Ichikawa, Jonathan. 2011. Quantifiers and Epistemic Contextualism. *Philosophical Studies*, 155(3):383–98. (p.143)
- King, Jeffrey C. 2007. What in the World Are the Ways Things Might Have Been? *Philosophical Studies*, 133(3):443–53.
- Kripke, Saul. 1980. *Naming and Necessity*. Harvard University Press.
- Kripke, Saul. 2011. Nozick on Knowledge. In Saul Kripke, editor, *Collected Papers*, volume 1, pages 162–224. Oxford University Press.
- Lawlor, Krista. 2013. *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford University Press.
- Lewis, David. 1973. *Counterfactuals*. Basil Blackwell.
- Lewis, David. 1979. Attitudes *De Dicto* and *De Se*. *The Philosophical Review*, 88(4):513–43.
- Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, 74 (4):549–67.
- Luper-Foy, Steven. 1984. The Epistemic Predicament: Knowledge, Nozickian Tracking, and Scepticism. *Australasian Journal of Philosophy*, 62(1):26–49.
- MacFarlane, John. 2005. The Assessment-Sensitivity of Knowledge Attributions. *Oxford Studies in Epistemology*, 1:197–233.
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
- Murphy, Peter. 2005. Closure Failures for Safety. *Philosophia*, 33:331–4.
- Murphy, Peter. 2006. A Strategy for Assessing Closure. *Erkenntnis*, 65(3):365–83.
- Nozick, Robert. 1981. *Philosophical Explanations*. Harvard University Press.

- Oaksford, Mike and Keith Stenning. 1992. Reasoning with Conditionals Containing Negated Constituents. *Journal of Experimental Psychology*, 18(4):835–54.
- Perry, John. 1986. From Worlds to Situations. *Journal of Philosophical Logic*, 15 (1):83–107.
- Perry, John. 1989. Possible Worlds and Subject Matter. In Sture Allén, editor, *Possible Worlds in Humanities, Arts and Sciences: Proceedings of Nobel Symposium 65*, pages 124–37. Walter de Gruyter.
- Roush, Sherrilyn. 2005. *Tracking Truth: Knowledge, Evidence, and Science*. Oxford University Press.
- Roush, Sherrilyn. 2010. Closure on Skepticism. *The Journal of Philosophy*, 107(5): 243–56.
- Roush, Sherrilyn. 2012. Sensitivity and Closure. In Kelly Becker and Tim Black, editors, *The Sensitivity Principle in Epistemology*, pages 242–68. Cambridge University Press.
- Rysiew, Patrick. 2001. The Context-Sensitivity of Knowledge Attributions. *Noûs*, 35(4):477–514.
- Schiffer, Stephen. 2004. Skepticism and the Vagaries of Justified Belief. *Philosophical Studies*, 119(1/2):161–84.
- Sen, Amartya K. 1971. Choice Functions and Revealed Preference. *The Review of Economic Studies*, 38(3): 307–17.
- Sherman, Brett and Gilbert Harman. 2011. Knowledge and Assumptions. *Philosophical Studies*, 156(1):131–40.
- Sosa, Ernest. 1999. How to Defeat Opposition to Moore. *Noûs*, 33(13): 141–53. (p.144)
- Stalnaker, Robert. 1986. Possible Worlds and Situations. *Journal of Philosophical Logic*, 15(1):109–23.
- Stanley, Jason. 2005. Fallibilism and Concessive Knowledge Attributions. *Analysis*, 65(2):126–31.
- Stine, G. C. 1976. Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies*, 29(4):249–61.

Vogel, Jonathan. 1987. Tracking, Closure, and Inductive Knowledge. In S. Luper-Foy, editor, *The Possibility of Knowledge: Nozick and His Critics*, pages 197–215. Rowman and Littlefield.

Vogel, Jonathan. 1999. The New Relevant Alternatives Theory. *Noûs*, 33(s13): 155–80.

Vogel, Jonathan. 2007. Subjunctivitis. *Philosophical Studies*, 134(1): 73–88.

White, Roger. 2006. Problems for Dogmatism. *Philosophical Studies*, 131(3): 525–57.

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

Wright, Crispin. 2004. Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary Volume*, 78(1):167–212.

Wright, Crispin. 2014. On Epistemic Entitlement (II): Welfare State Epistemology. In Dylan Dodd and Elia Zardini, editors, *Scepticism and Perceptual Justification*, pages 213–47. Oxford University Press.

Yablo, Stephen. 2011. Achieving Closure. Kant Lecture delivered at Stanford University on May 20.

Yablo, Stephen. 2012. Aboutness Theory. Online paper accessed August 26, 2013. URL <<http://www.mit.edu/~yablo/home/Papers.html>>.

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

Notes:

(1) The term “fallibilism” means many different things to many different people. I explain in more detail what I mean by “fallibilism” in §2.1.

(2) In possible-worlds parlance, W would be a set of “centered possible (or impossible) worlds” (see Lewis 1979 on centered worlds and King 2007 on impossible worlds), but this need not be a context-independent “intended standard model of super-reality” (Stalnaker 1986: 122). As Stalnaker remarks, “The formalism of possible worlds semantics assumes that possible states of the world are disjoint alternatives, and that everything that can be said *within a given context* can be said by distinguishing between these alternatives Nothing in the formalism of possible worlds semantics, or in the intuitive conception of a way things might be, or a possible state of the world, excludes an interpretation in which possible worlds are alternative states of some limited subject matter. Possible worlds must be complete, relative to the distinctions that can be made within the given interpretation, but they might be quite partial relative to another interpretation, or relative to an external intuitive commentary on the interpretation” (118–19). Compare Lewis (1996, 552): “we needn’t decide whether they must always be maximally specific possibilities, or whether they need only be specific enough for the purpose at hand. A possibility will be specific enough if it cannot be split into subcases in such a way that anything we have said about possibilities, or anything we are going to say before we are done, applies to some subcases and not to others. For instance, it should never happen that proposition P holds in some but not all subcases; or that some but not all sub-cases are eliminated by S ’s evidence.” For simplicity, I will not relativize the set W to contexts, but these remarks should be kept in mind. The framework developed here can also be generalized to include what Perry (1986) calls *partial ways* the world could be (see Holliday 2014b).

(3) See King 2011 for a survey of views of structured propositions.

(4) I deliberately use the term “TF-atomic” instead of “atomic.” A proposition that has a complex structure may count as TF-atomic, because it does not have the structure of a truth-function applied to one or more propositions.

(5) As explained in §2.4, one may take $P\{w \in w \mid P \text{ is true at } w \text{ considered as actual}\}$.

(6) By calling this picture “standard,” I am not claiming that *all* contemporary views of knowledge fit into it.

(7) The sizes of the various regions in the diagram are not intended to reflect the sizes of the corresponding sets, and the locations of the regions are not intended to reflect the “distance” of scenarios from w .

(8) Dretske [1981] gives thin accounts of both r and u in terms of knowledge: “let us call the set of possible alternatives that a person must be in an evidential position to exclude (when he knows that P) the *Relevancy Set* (RS). In saying that he must be in a position to exclude these possibilities I mean that his evidence or justification for thinking these alternatives are *not* the case must be good enough to say he *knows* they are not the case” (371). Lawlor [2013] gives a thicker account of what makes an alternative one that must be eliminated for knowledge, i.e. an alternative in $r^c(P, w)$ it is “an alternative to p that a reasonable person would want ruled out by reasons or evidence before judging that S knows p ” (152), where the notion of a *reasonable person* is given a substantive independent characterization (Lawlor 2013, §5.1).

(9) Or more generally, the set of alternatives for P in a given context should form a nontrivial partition of W .

(10) Of course, this is the partition view where each cell contains only one scenario. Another option would be to use Alternatives, but only from partitions with the property that if one of the scenarios in an Alternative is in $u_c(P, w)$, then all of the scenarios in that Alternative are in $u_c(P, w)$, following the quote from Lewis in footnote 2. Then we could define u_A from u : an Alternative is in $u_A^c(P, w)$ iff all scenarios in that Alternative are in $u_c(P, w)$.

(11) Cf. Goldman’s (1976, 779–84) detailed discussion of the notion of a *perceptual equivalent* of a state of affairs, which does not require exact match of perceptual appearances (781).

(12) I am not claiming (what certain kinds of externalists about perception would deny) that given a scenario w in which the agent believes p , there is always another possible scenario v in which the agent has the same *type of experience* or the *same evidence*, but in which P is false; for I am not assuming that subjective indistinguishability entails the same type of experience or evidence. I am also not claiming that if w and v are subjectively indistinguishable, then $u_c(P, w) = u_c(P, v)$, i.e. that the agent in w has eliminated exactly the same alternatives as the agent in v .

(13) Examples abound in the literature on skepticism, but let us consider another. In the actual scenario w , Jones, who lives in the U.S., receives a postcard from Smith, who is visiting the U.K. The postcard is signed by Smith in his unique handwriting, stamped and dated by U.K. postal officials, and so on. Jones recognizes all of this, and he correctly takes Smith to be a perfectly reliable reporter of his

vacation whereabouts. According to everyone but radical skeptics, on the basis of receiving such a postcard, Jones can know that *Smith visited the U.K. some days ago* (P). Yet everyone must also admit that there are possible scenarios v in which everything appears the same to Jones (during his whole life up until now) as in w , but the postcard was not sent by Smith, and Smith never visited the U.K., so $v \in u_C(P, w) \cap (W_w - P)$. Some of these scenarios are ones in which skeptical hypotheses incompatible with P obtain: in some of them, *the postcard was forged by a team of deceivers* (SH_1); in others, *all the world and Jones's memories were created five seconds before he received the postcard* (SH_2); and so on. Of course, such deceptive possibilities arise for a tremendous number of other propositions that Jones believes about the external world.

(14) It is sometimes suggested that one has an “infallibilist” conception of knowledge if one accepts the following principle: if an agent knows that P , then her evidential probability for P is 1. According to the present conception of infallibilism and fallibilism, that suggestion is incorrect. As defined below, a fallibilist may hold that (i) an agent knows that P , so (ii) the agent's evidential probability for P is 1, even though (iii) there may be some scenarios that are subjectively indistinguishable from the agent's actual scenario—and in that sense are uneliminated—in which P is false. The fallibilism is in the conjunction of (i) and (iii). Dretske (1981; 1971) is such a fallibilist who holds that (i) implies (ii). (Note that such a view is not inconsistent with Dretske's denial of closure, because he does not hold that probability 1 is sufficient for knowledge. Some propositions will have probability 1, although they are not known, because they are entailed by other propositions with probability 1 that are known. See Dretske 2006.)

(15) Not, for example, for worries about *concessive knowledge attributions* of the form “I know that P , but it's possible that $\neg P$ ” or “I know that P , but it might be that $\neg P$ ” (Rysiew 2001, Stanley 2005). I see no reason why a fallibilist in one of the senses stated above should be committed to the felicity of such claims. Fallibilists hold that an agent can know P even if $u_C(P, w) \cap (W - P) \neq \emptyset$, but what does this have to do with the semantics/pragmatics of claims with the epistemic modals “possible” and “might”? According to Yalcin (2011, 309), an utterance of “it might be that $\neg P$ ” expresses (roughly) that there is a $\neg P$ -scenario V compatible with what the agent believes, which does not follow from there being a $\neg P$ -scenario $v \in U_C(P, w)$. Indeed, it is compatible with there being a $\neg P$ -scenario $v \in u_C(P, w)$ that the agent in w believes P with the utmost certainty. It is noteworthy in this

connection that Dretske [1981], a strong fallibilist, remarks: “it does seem reasonable to insist that if S knows that P , he does not believe that he might be wrong. In other words, if the bird-watcher really believes that the bird he sees might be a grebe, then he does not know it is a Gadwall” (378n8) (cf. footnote 14). Of course, fallibilists are committed to there being contexts in which it would be true to say (if it can be said without changing the context) “the agent knows P , but the agent has not eliminated all $\neg P$ -scenarios.” But here “eliminated” is a theoretical term, so we should not conclude that pre-theoretic intuitions about natural language pose any problem for fallibilism here.

(16) Something more may be required to know Q , such as “putting two and two together” and inferring Q from P_1, \dots, P_n , or simply *coming to believe Q as a result of the same experiences that make the agent believe P_1, \dots, P_n* , but no more empirical investigation of the world is required to know Q than to know P_1, \dots, P_n (assuming the agent has already had sufficient experience to enable her to grasp the concepts required for understanding Q).

(17) *Proof*: if **(Knows)** does not hold for Q , then there is some $v \in r_c(Q, w) \cap u_c(Q, w)$. Since $v \in r_c(Q, w)$, it follows by **(1)** that $v \in r_c(P_i, w)$ for some $1 \leq i \leq n$; then since $v \in r_c(P_i, w) \cap r_c(Q, w) \cap u_c(Q, w)$, it follows by **(2)** that $v \in u_c(P_i, w)$. Thus, $r_c(P_i, w) \cap u_c(P_i, w) \neq \emptyset$, so **(Knows)** does not hold for P_i .

(18) One could treat any multi-premise closure principle as a single-premise principle by loading the other premises into **R**, e.g. taking **R** to be the relation that P bears to Q iff the agent knows $P \rightarrow Q$, but this trick is not helpful.

(19) Recall the postcard example in footnote 13, and take Q to be $\neg(\text{SH}_1 \vee \dots \vee \text{SH}_n)$.

(20) Heller rejects the idea that $r_c(P, w)$ contains only the closest $\neg P$ -scenarios according to a Lewisian similarity ordering \leq_w^c (see below in the text for this notation), arguing that any “close enough” $\neg P$ -scenarios must be included as well. But since Heller [1999, 201f.] holds that the set of possible scenarios that are “close enough” to w *Close Enough_c(w)*, is independent of any proposition in question, Heller’s view is equivalent to the view that $r_c(P, w)$ is the set of *closest* $\neg P$ -scenarios according to a more coarse-grained ordering \leq_w^c , of which \leq_w^c is a refinement: define $v \leq_w^c$ iff $v \in \text{CloseEnough}_c(w)$ or $v \leq_w^c$. Then assuming that whenever $v \in \text{CloseEnough}_c(w)$ and $v \leq_w^c$, we have

$v \in \text{CloseEnough}_c(w)$, the claimed equivalence follows from the fact that for any set A of scenarios:

$$\text{Closest}_{\leq w}^c(A) = \text{Closest}_{\leq w}^c(A) \cup (\text{CloseEnough}_c(w) \cap A).$$

(21) Where $w = \langle w, a, t \rangle$, by “the agent” I mean a (recall § 1.1). Those who reject “trans-world identity” may substitute “a counterpart of a .” Nothing here turns on this subtlety, so I will ignore it in what follows.

(22) Nozick [1981, 680n8] tentatively proposes alternative truth conditions for counterfactuals. However, he also indicates that sensitivity may be understood in terms of Lewis’s semantics for counterfactuals. This has become the standard practice in the literature. For example, see Vogel 1987, Comesaña 2007, and Alsepector-Kelly 2011.

(23) Here I follow Luper-Foy’s [1984, 29] statement of the sensitivity condition with “methods,” which differs slightly from Nozick’s, which we could write down as well. For simplicity I omit “methods” for adherence below.

(24) One may not wish to call this a set of “uneliminated” scenarios, but there is nonetheless a structural analogy between r'_c and u'_c on the one hand and r_c and u_c on the other.

(25) Nozick [1981, 680n8] suggests interpreting adherence counterfactuals $P \square \rightarrow BP$ with true antecedents in such a way that the sphere over which $P \rightarrow BP$ needs to hold may differ for different propositions P . By contrast, I am interpreting adherence as a kind of $\exists \forall$ condition, in a sense that generalizes that of §2.2: there is a fixed set $R'_c(w)$ of scenarios such that for all propositions P , to know P one needs to meet an epistemic success condition in the P -worlds in $R'_c(>w)$. A $\forall \exists$ interpretation of adherence that, e.g., allows the adherence sphere for $P \vee Q$ to go beyond that of P , would create yet another source of closure failure in Nozick’s theory.

(26) Nozick used the term “variation” for what I call “sensitivity” and used “sensitivity” to cover both variation and adherence; but the narrower use of “sensitivity” is now standard.

(27) For those safety theorists who propose only necessary conditions for knowledge, see Remark 4.2 in Holliday Forthcoming on the relation between closure failures for necessary conditions for knowledge and closure failures for knowledge.

(28) For discussion of closure failures for safety, see Murphy 2005, 2006, Alspector-Kelly 2011, and Holliday Forthcoming.

(29) i.e. reflexive, transitive, and such that for all $u, v \in W_w$, either $u \leq_w^c v$ or $v \leq_w^c u$.

(30) i.e. $w \in W_w$ and for all $v \in W_w$, $w \leq_w^c v$.

(31) For all $x, y \in W_w$, define $x \leq_w^c y$ iff either (i) for all propositions P , $y \notin r_c(P, w)$, or (ii) there is some proposition Q such that $x \in r_c(Q, w)$ and $y \in W - Q$.

(32) I have written alpha in the equivalent form that Sen [1971, §9, n1] calls α^* and beta in the form given by Bordes [1976, §2]. My conditions look different than theirs at first because my \mathbf{r} function picks the “best” $\neg P$ -scenarios, whereas the economist’s *choice function* picks the best P -scenarios. Another minor difference is that the \mathbf{r} function takes in a proposition, whereas a choice function takes in a set. A proof of Theorem 1 in the case where the input to \mathbf{r} is a set is in Holliday 2012, §3.A, and the proof there can be easily adapted for the present setup.

(33) According to what Vogel [1999, 168] calls “Backsliding” RA theories, which blur the roles of the \mathbf{r} and \mathbf{u} functions, alternatives can become “irrelevant” when one has good evidence against them. A Backsliding RA theorist might claim that $r_c(Q, w) = \emptyset$ holds only if the agent has done a lot of empirical investigation. But this is not how the \mathbf{r} function works for any of the theories discussed here. See Vogel 1999 for arguments against Backsliding accounts.

(34) I am not objecting to the view that one may be *entitled*, in the sense of Wright [2004], to accept contingent empirical propositions, such as the negations of skeptical hypotheses, without doing empirical work for them—provided the view does not add that one thereby *knows* the propositions; Wright [2004] is careful not to make this further claim. Relatedly, I am not objecting to the view that one may *justifiably take for granted*, in the sense of Sherman and Harman [2011], contingent empirical propositions, such as the negations of skeptical hypotheses, without doing empirical work for them; Sherman and Harman are explicit that one cannot come to know a proposition just by justifiably taking it for granted. Turning from justified *taking for granted* to justified *belief*, White [2006, §9] argues that we have a priori “default” justification for believing, or that we are “entitled” to believe, the negations of skeptical hypotheses; but he does not claim that we have a priori *knowledge* of the negations of skeptical hypotheses. Contrary to White, Schiffer [2004, 178] argues that “There is nothing in the concept of a priori justified belief to warrant the claim that we’re a priori justified in disbelieving skeptical hypotheses.” Although Schiffer proposes a revised concept of justification* according to which we are a priori justified* in disbelieving skeptical hypothesis, he also does not claim that we have a priori *knowledge* of the negations of skeptical hypotheses. Thanks to an anonymous referee for posing the question of how what I claim is intolerable relates to the views of Wright, White, and Schiffer.

(35) About Dretske’s [1970] zebra case, Stine [1976, 258] writes: “[O]ne does know what one takes for granted in normal circumstances. I do know that it is not a mule painted to look like a zebra. I do not need evidence for such a proposition ... [I]f the negation of a proposition is not a relevant alternative, then I know it—obviously, without needing to provide evidence.” Cohen [1988, 99] responds: “Here, I think Stine’s strategy for preserving closure becomes strongly counter-intuitive. Even if it is true that some propositions can be known without evidence, surely this is not true of the proposition that S is not deceived by a cleverly disguised mule.” The key point is to consider the *kinds* of propositions that RE₃V theories imply can be known with no requirement of eliminating scenarios.

(36) And by Cohen [1988, 111]: “Radical skeptical hypotheses are immune to rejection on the basis of *any* evidence. There would appear to be no evidence that could count against the hypothesis that we are deceived by a Cartesian demon Radical skeptical hypotheses are

designed to neutralize any evidence that could be adduced against them.”

(37) Heller [1999, 207] considers and rejects this reply. By contrast, DeRose [2000, 135] endorses a similar position.

(38) To put it this way is misleading, since a closure principle is not something that agents use in acquiring knowledge (except about other agents' knowledge). It is something that we use in *reasoning about agents' knowledge*.

(39) Cohen [1998, 303n24] suggests that Lewis's Rule of Attention may need to be defeasible; Ichikawa [2011, §4] disavows it; and Blome-Tillman (2009: 246–7) argues that it is too strong. DeRose [2009, Ch. 4] suggests that members of a conversation may resist context changes by sticking to their own “personally indicated epistemic standards.”

(40) I am assuming that the agent grasps the concepts needed to understand the new disjunct Q (cf. Williamson 2000, 283). Also recall the meaning of the notation $KP \Rightarrow KP'$ from §2.1. Although I agree with Dretske, Kripke, and Nozick in endorsing $KP \Rightarrow K(P \vee Q)$ so understood, not everyone does. According to Yablo's (2011; 2012; 2014) view of closure, knowing $P \vee Q$ relative to a context C may require more empirical investigation than knowing P relative to C , since the *subject matter* of $P \vee Q$ is not in general included in that of P . Although I take Yablo's move to connect subject matter and epistemic closure to be deep and important (see footnote 58), I disagree with the specifics of his view of the connection (see Holliday 2014b). While I cannot do justice to his view here, I will briefly register points of disagreement. On one way of developing Yablo's view (based on “reductive truthmaking” and the relation of “content-parthood”), $K(P \wedge Q) \Rightarrow K(P \vee Q)$ does not hold—so an agent may know a conjunction and yet have more empirical work to do to know the disjunction of the conjuncts—and although $K(P \wedge Q) \Rightarrow KP$ holds when P and Q are TF-atomic, it does not hold in general for arbitrary P and Q —so an agent may know a conjunction without knowing the conjuncts—which is hard to swallow. (Note that the propositions in the consequents of those closure principles do not seem to “change the subject” relative to the propositions in the antecedents.) On another way of developing his view (based on “recursive truthmaking” and the relation of “inclusive entailment”), the principle $K(P \wedge Q) \Rightarrow K(P \vee \neg Q)$ does not hold, which is also hard to swallow. Indeed, so is the denial of $KP \Rightarrow K(P \vee Q)$, even though this principle—apparently unlike the previous ones—has

something new appear in its consequent. One must be careful to distinguish two ideas: first, the less radical idea that in a context C relative to which an agent has done enough empirical work to count as knowing P , someone's raising the issue of Q by bringing up $P \vee Q$ may change the context to a C' relative to which the agent has not done enough empirical work to count as knowing $P \vee Q$ (or P); second, the more radical idea that knowing $P \vee Q$ relative to a context C may require more empirical work than knowing P relative to C . The second idea, like the denial of the conjunctive closure principles, strikes me as unnecessary and undesirable (admittedly, we are in near-bedrock territory here).

(41) Some might think that Gettier cases show we should reject **contrast/enough**. For example: not having any idea what time it is, you check a clock that—unbeknownst to you—has been stopped for weeks on 5:43; as it happens, the time is now 5:43; but you do not come to know this from the stopped clock. Where F is the proposition that *the time* is 5:43 and S is the proposition that *the clock has stopped*, one might think this is a case in which knowing F requires ruling out $(F \wedge S)$ -possibilities, which would explain your ignorance of F (since you have not ruled those out) and violate **contrast/enough**. But this is a mistake. What explains your ignorance of F is that since you have only looked at a stopped clock, you have not ruled out various relevant scenarios in which F is *false* and the time is something other than 5:43. If by some other means you had ruled out every scenario in which F is false, then it would be absurd to say “I agree that you've ruled out every scenario in which the time is something other than 5:43, but you still don't know the time is 5:43 unless you rule out such-and-such scenarios in which the time is 5:43.”

(42) The claim that for every Q , knowing P requires eliminating some $(\neg P \wedge \neg Q)$ -scenario is essentially equivalent to infallibilism; and if for every set of scenarios there is a proposition true in exactly those scenarios, then the claim is exactly equivalent to infallibilism.

(43) For reasons of space, I cannot go into the theory of the **u** function here. For simplicity (but not in the final analysis), one may assume that for each w and C there is a set $U_c(w)$ of uneliminated scenarios such that for all P , $u_c(P, w) = U_c(w)$. Note that this is not the same as $RO_{\exists v}$ from §2.2, since it does not imply an analogue of **contrast** for **u**. This will be important given the argument of §3.3.

(44) An r function may contain some eliminable redundancy in the following sense. Given $r_c(P, w)$, define

$\neg r_c(P, w) = \{A \in r_c(P, w) \mid \text{there is no } A' \in r_c(P, w): A' \subseteq A\}$. $\neg r_c(P, w)$ may contain fewer alternative sets than $r_c(P, w)$, but **(Knows)** holds for r iff **(Knows)** holds for $\neg r$

(45) Or even a weaker **e-fallibilism** generalizing **e-fallibilism**.

(46) Here is a trickier argument using existential propositions. Start with a standard fallibilist view according to which there are many propositions P that Jones can know by eliminating $\neg P$ -scenarios, without having to eliminate the $\neg P$ -scenarios that a radical skeptic would raise against him, such as subjectively indistinguishable scenarios in which *Jones is a brain in a vat (BIV)*. One such proposition P that Jones can know, just by getting a good look at Smith's body, is that *Smith is not a BIV*. (Knowing that *someone else* is not a BIV is less of a problem!) Now where Q is the proposition that *someone is not a BIV*, surely if Jones knows P , then with a step of logic he can know Q . Finally, noticing that our initial assumption implies that Jones can know P by eliminating only scenarios in which Q is true (for they are scenarios in which *Jones is not a BIV* and hence *someone is not a BIV*), it follows that Jones can know Q by eliminating only scenarios in which Q is true. Notably, this argument is compatible with **no VK** in §3.4.

(47) Note that if $r_c(P, w) = \emptyset$, then P cannot be known, given the existential character of **(Knows)**.

(48) And it is *sufficient* together with certain assumptions on u , such as that of footnote 43.

(49) Note that neither (14) nor (15) depend on the order of the disjunct or conjuncts, so the particular choice of $CCNF(P)$ among equivalent but permuted CCNFs does not matter for $r_c(P, w)$

(50) In this case, there are $7^3 \times 3^2$ ways of choosing C_1 - C_5 . But this does not mean that there are $7^3 \times 3^2$ distinct sets in $r_c(P \vee (q \wedge r))$, since many ways of choosing the C_i may result in the same union $A_1 \cup A_2 \cup A_3 \cup A_4 \cup A_5$, and even distinct unions may be redundant because they contain others (see footnote 44).

(51) The reason is that each of the closest $\neg(C_1 \wedge \dots \wedge C_n)$ -scenarios is a closest $\neg C_i$ -scenario for some i :

$$\text{Closest}_{\leq w}(\neg(C_1 \wedge \dots \wedge C_n)) \subseteq \bigcup_{1 \leq i \leq n} \text{Closest}_{\leq w}(\neg C_i)$$

so

$$r_c(C_1 \wedge \dots \wedge C_m, w) \subseteq \bigcup_{1 \leq i \leq m} r_c(C_i, w)$$

and it is easy to check from (15) and (14) that there is some $A \in r_c(C_1 \wedge \dots \wedge C_m)$ such that

$$\bigcup_{1 \leq i \leq m} r_c(C_i, w) \subseteq A$$

(52) It would be more natural to think in terms of the probability of something more coarse-grained than scenarios, such as the *Alternatives* mentioned in §2.1 and footnote 10, but I skip over these details here. I also skip over the kind of probability in question, whether probability on the agent's evidence—in which case the suggestion in the text might blur the roles of the r and u functions (recall footnotes 33 and 14)—or a kind of objective probability or, in the spirit of contextualism, probability for the attributors. It is noteworthy here that Vogel [1999, 163] argues that probability cannot provide a sufficient condition for relevance. Roughly, the argument runs as follows: if a *proposition* S is an irrelevant “alternative” to P , but a *proposition* Q is probable enough to be a relevant “alternative” to P , then of course $Q \vee S$ is also probable enough to be a relevant “alternative” to P ; but then since on Vogel's view, ruling out $Q \vee S$ requires ruling out S , it follows that S is a relevant “alternative” that must be ruled out for knowledge of P after all. Contradiction. Of course, this argument assumes the propositional view of alternatives that I rejected in §2.1 because it violates the disjointness condition on alternatives in a context.

(53) They would satisfy **noVK** because there are always some $\neg P$ -scenarios with maximal probability, cost-weighted probability, or whatever, *relative to* other $\neg P$ -scenarios.

(54) Note that if for each $e \in \mathbf{E}$, the inductive path to knowing G_e by observing the emeralds in the set E_i is included in $r_c(G_e, w)$ by a modified version of (14), then the inductive path to knowing the conjunction G by observing the emeralds in the set E_i will be included in $r_c(G, w)$ by (15).

(55) It is important that Vogel's [1999, §4] arguments, according to which a certain version of the standard alternatives picture cannot handle inductive knowledge, do not apply to the multipath picture. In short, Vogel attacks a view according to which knowing a proposition like G involves eliminating a single set of $\neg G$ -worlds that resemble the actual world. On such a view, it seems difficult to explain why after observing a few emeralds, one has not eliminated the right $\neg G$ -worlds resembling the actual world, but after observing more emeralds, one

has. The solution in the multipath picture is to reject the view that the only path to knowing a proposition like G involves eliminating “close” $\neg G$ -scenarios; instead, an agent can take one of the A_i paths described above, which involves coming to know G_e for each $e \in E_i$, where this may involve eliminating close $\neg G_e$ -scenarios. One might object that this response just assumes inductive knowledge is possible, rather than deriving its possibility from first principles. But I do not see this as an objection. It is just an observation that the multipath picture by itself is not a theory of inductive knowledge.

(56) In one of Dretske’s [2005] examples, P is the proposition that *there are cookies in the jar*, and Q is the proposition that *idealism is false*. As Dretske quips, “Looking in the cookie jar may be a way of finding out whether there are any cookies there, but it isn’t—no more than kicking rocks—a way of refuting Bishop Berkeley” (15).

(57) It follows from (18) that $\neg S$ is not deeply necessary, $\neg S_w \neq W_w$, which with *noVK* implies $\phi \notin r_{\Delta}(\neg S, w)$, which with **semi-contrast** implies (20).

(58) A similar argument is given by Hawthorne [2004, 41], employing a principle similar to **M-equiv**, namely closure under a priori equivalence, though Hawthorne uses the argument for different dialectical purposes. See Sherman and Harman 2011 for an argument against Hawthorne's equivalence principle, which also applies to **M-equiv**. Another reason to worry about the equivalence principle and **M-equiv** is that these principles seem to commit the mistake of what Perry [1989] calls "losing track of subject matter": losing track of what propositions are *about*, considering only the possibilities in which they are true. Barwise and Perry [1983] have argued that losing track of subject matter leads to serious problems in semantics, and the same may be true in epistemology. If the range of alternatives that one must eliminate in order to know a proposition may depend not only on the structure of the proposition, as I have argued, but also on what the proposition is *about*, then it is not clear why knowing $P \wedge \neg S$, from the example below in the text, should not require eliminating more alternatives than knowing P —even if (it is a priori that) P metaphysically entails $P \wedge \neg S$. Given a typical skeptical hypothesis S , an ordinary proposition P does not, in the terminology of Barwise [1981, 395], *strongly imply* $P \wedge \neg S$, i.e. it is not the case that every *situation* that supports P supports $P \wedge \neg S$, since supporting $P \wedge \neg S$ requires supporting $P \wedge \neg S$, which brings in extra subject matter (nor, of course, is $P \wedge \neg S$ a truth-functional consequence of P). (By contrast, $P \wedge Q$ strongly implies P , and P strongly implies $P \vee Q$.) The move to a framework that includes partial situations is fully compatible with the multipath picture of knowledge, though I cannot go into details here (see Holliday 2014b). The general idea that the range of alternatives that one must eliminate in order to know some Q depends on what Q is *about*, contrary to Hawthorne's equivalence principle, is due to Yablo (2011; 2012; 2014). However, his specific view of the connection between closure and subject matter disagrees with some of the views about closure in this paper (recall footnote 40).

(59) This argument (like that of Hawthorne 2004, 41) reflects the fact, which should be obvious to students of modal logic, that together the principles $KP \& \Box(P \leftrightarrow Q) \Rightarrow KQ$ and $K(\wedge Q) \Rightarrow (KP \& KQ)$ (or $KP \Rightarrow K(P \vee Q)$) suffice to derive $(KP \& \Box(P \rightarrow Q)) \Rightarrow KQ$, where \Box is a normal modal operator.

(60) Epistemologists typically assume that knowing a conditional $P \rightarrow \neg S$, where P is an ordinary proposition and S is a metaphysically incompatible skeptical hypothesis, does not require eliminating skeptical S -scenarios.

(61) One can easily verify this by comparing the the CCNFs of p and $p \wedge \neg s$ for TF-atomic p and s .

(62) Here I mean $c(\text{CCNF}(P))$, but I will write “ $c(P)$ ” for convenience.

(63) This assumes that p and p' do not have further structure, ignored by the truth-functional analysis given here, such that $p \wedge p'$ has new heavyweight consequences.

(64) Of course, there is a close connection between the multi-premise principles of closure under known implication, $(KP \& K(P \rightarrow Q)) = KQ$, and $(KP \& KP') = K(P \wedge P')$. First, the former essentially guarantees the latter: by closure under known implication, an agent who knows P and the tautology $P \rightarrow (P' \rightarrow (P \wedge P'))$ has done enough empirical work to know $P' \rightarrow (P \wedge P')$, so if the agent also knows P' , then by closure under known implication again she has done enough empirical work to know $P \wedge P'$. Second, the latter guarantees the former assuming **TF-cover**: if $(KP \& KP') = K(P \wedge P')$ holds, then an agent who knows P and knows $P \rightarrow Q$ must have done enough empirical work to know $P \wedge (P \rightarrow Q)$, which by **TF-cover** requires that she has done enough empirical work to know Q . Thus, assuming single-premise logical closure, the two multi-premise principles stand or fall together.

(65) Can a *TF-consequence* Q of P have a “heavyweight” status compared to the “lightweight” status of P , requiring more epistemic work to know? I don't think so. See the Appendix for a related discussion.

(66) This is compatible with the super-shifty contextualist view (recall §4.2.4) that whenever we mention or think about S , we shift the context from C to a C' relative to which the agent does not count as knowing $\neg S$. The benefit to super-shifty contextualists of adopting the multipath picture is that they are no longer forced to say that relative to C , the agent counted as knowing $\neg S$ no matter what epistemic work she had done; instead, the reason she could count as knowing $\neg S$ relative to C is that she did the epistemic work required to know the P that entails $\neg S$.

(67) This way of achieving strong closure bears some resemblance to the more sophisticated *recursive tracking theory* of Roush (2005; 2012). So does the recursive definition in Theorem 2, although in Theorem 2 the alternative sets for a proposition P are determined by the structure of the proposition itself and the alternative sets for its

parts, rather than the alternative sets for all propositions that (are known to) entail P .

(68) The second quoted sentence is from the footnote to the first sentence.

(69) Kripke [2011, 199] also discusses the inconsistency, pointed out to him by Assaf Sharon and Levi Spectre.

(70) Notation: $\pm P$ is either P or $\neg P$; if $\pm P$ is P , then $\mp P$ is $\neg P$; if $\pm P$ is $\neg P$, then $\mp P$ is P .

(71) Wright (2014) also warns his reader not to confuse the likes of $\neg D$ and $\neg(M \wedge D)$: “we don’t have a visual warrant for thinking that *those animals have not been cleverly disguised in a visually undetectable way*, but we do, in the relevant circumstance, have a visual warrant for thinking that *those animals are not **mules** that have been so disguised*. Maybe we are confused by the operation of some kind of implicature here: maybe saying, or thinking, ‘It is not the case that those animals are cleverly disguised mules’ somehow implicates, in any context of a certain (normal) kind, that ‘Those animals have not been cleverly disguised’. But anyway, it doesn’t entail it: not-($P \wedge Q$), dear reader, does not entail not- Q !” (234–5).

(72) See Roush 2010 for arguments to the effect that if an agent knows $\pm P$, then she can know $\neg(\mp P \wedge S)$ for any S you like, and that intuitions to the contrary can be explained away.



Access brought to you by: New York University