

CHAPTER I

*Motivation, Strategy, and Definition***1.1 Closure as Axiom**

Knowledge closure is, roughly, the principle that any agent who knows *P* and recognizes that *P* implies *Q* knows – or is in a position to know – *Q*.¹ Although the principle has been challenged – most famously by Robert Nozick and Fred Dretske – most contemporary epistemologists remain committed to it.² For many, indeed, that commitment is so firm that they view a theory that rejects closure to be seriously undermined if not refuted outright for that reason alone.³

A set *S* is closed under an operation *O* when applying *O* to members of *S* delivers members of *S*. The set of natural numbers, for example, is closed under addition: adding two natural numbers always produces a natural number. This follows from the Peano axioms for the natural numbers and the definition of addition. The set of one's ancestors is closed under the parent-relation: any parent of one's ancestor is also one's ancestor. This is explicable by appeal to the fact that “*x* is an ancestor of *y*” means “*y* is a direct or indirect descendant of *x*.” And the set of true propositions is closed under deductive consequence: any consequence of a true proposition is also a true proposition. This is explicable by appeal to the nature of deductive consequence and a soundness proof. Closure principles typically have some sort of explanatory ground.⁴

¹ Refinements will come in §§1.4–1.7.

² The *loci classici* of closure denial are Dretske 1970 and Nozick 1981, chapter 3.

³ “[T]he idea that no version of [closure] is true strikes me, and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years.” (Feldman 1995, 487) “[I]f a philosopher advances a view that forces us to reject closure, that should be taken as a *reductio* of that philosopher's view.” (Fumerton 1995, 131) “Robert Nozick's counterfactual analysis of knowledge is famously inconsistent with intuitive closure, but that is usually taken as a reason for rejecting the analysis, not for rejecting closure.” (Williamson 2000b, 117)

⁴ Although not always. The Peano axioms themselves include two closure principles.

However, this does not seem to be the case for knowledge closure. There is no argument of the form:

- (1) Knowledge is like so . . .
- (2) Inference is like so . . .
- (3) Therefore, knowledge closure is true,

where the premises cite uncontroversial characteristics of knowledge and inference.⁵

Nor does every serious epistemological view imply closure. While the closure-denying sensitivity accounts of Nozick and Dretske face numerous objections, they nevertheless exert considerable intuitive pull. A belief is sensitive when, were the belief false, the agent would not believe it.⁶ Sensitivity is not closed under deductive entailment.⁷ Nevertheless, there is undoubtedly *something* unattractive about a belief's counting as knowledge when the agent would still have believed it if it were false. Perhaps this intuition is illusory in some way.⁸ But the intuition is there, and strong enough to motivate a variety of successor views that attempt to reconcile sensitivity with closure.⁹

So closure advocates cannot claim that every intuitively plausible epistemological view implies closure. But there is also no argument from relatively uncontroversial truisms about knowledge and deductive inference on the table. This isn't to say that there are no arguments for closure

⁵ I will discuss a possible exception in §1.8 and §7.2.

⁶ This is the simplest version of sensitivity. Actual accounts are more elaborate. Some relativize to method: if the belief were false and the agent were to employ the same method, then the agent would not believe it by that method (Nozick 1981). (Nozick also includes an "adherence" condition: were the belief true and the agent to employ the same method, the agent would believe it by that method.) Others attach the modality to the agent's reason for belief rather than the belief itself: if the belief were false, the agent would not have the reason she has for believing it (Dretske 1971). Finally, others couch sensitivity in probabilistic rather than modal terms: the probability that the belief is true given that the agent believes it is 1 (Dretske 1983) or high but not necessarily 1 (Roush 2005, who also includes a probabilistic version of Nozick's adherence condition). And there are many other versions. Nozick called sensitivity "variation," and called the conjunction of variation and adherence "sensitivity" (or "tracking"). Nevertheless, "sensitivity" is reserved for variation in the subsequent literature, a terminological tradition that I follow here.

⁷ "I have hands" is sensitive: if I were not to have hands — because of an unfortunate accident with a table saw, for example — then I would not believe that I do. But "I am not a handless brain in a vat (BIV) stimulated to have the very experiences I do have" is not sensitive: if I were a BIV, I would still believe that I am not a BIV (since my reasons for believing this, whatever they might be, would remain). Nevertheless, "I am not a handless BIV" follows from "I have hands."

⁸ Sosa 1999a suggests that the intuition results from a confusion of sensitivity with its contrapositive, safety. These are not equivalent since both incorporate subjunctive conditionals, which are not truth-preserving under contraposition.

⁹ See, for example, DeRose 1995, 1996, and 2010; Roush 2005 and 2012; Baumann 2012; Black 2008a; and Murphy & Black 2012.

at all; far from it. But they typically proceed by presenting considerations, not directly in favor of closure, but rather against its denial. Those arguments deserve serious attention.¹⁰ But it is *prima facie* surprising that this widely endorsed principle isn't derivable from undisputed characteristics of knowledge and inference.

Closure is instead typically treated as an independent epistemological axiom.¹¹ As such, it is thought to be warranted in the way that axioms often are: it is intuitively obviously correct.¹² But if closure is a primitive epistemological axiom, it's an unusually complex one. Like the parallel postulate in Euclidean geometry, whose complexity motivated attempts to derive it from the other simpler axioms, it stands out among seeming truisms concerning knowledge ("what you know you believe," "what you know is true," "what you know can't be accidentally true," and so on) as begging for derivation from simpler axioms concerning knowledge and deductive inference.

The parallel postulate proved not to be so derivable. It also proved to be eliminable in favor of alternatives, giving rise to non-Euclidean geometries. Euclidean geometry is both intuitive and reasonably accurate as a representation of local observable space. But the parallel postulate turned out to be a dispensable theoretical posit rather than an unassailable geometric primitive in the representation of the geometry of the physical universe overall.

I suggest that a similar situation holds with respect to closure. Closure is an expression of an undeniable truth: deductive inference is an excellent way to extend one's knowledge. But that undeniable truth is compatible with closure's strict falsehood as a universal feature of epistemic space; excellence is not undermined by failure under exceptional circumstances. Just as caution must be exercised when extrapolating from the apparent geometric characteristics of our local space to the universe overall, so caution must be exercised when extrapolating from the undeniable truth that closure reflects to its supposed status as a universal epistemic truth. Closure is also a dispensable theoretical posit rather than an immobile pivot around which the epistemological landscape must turn.

¹⁰ See Chapter 10.

¹¹ "This principle seems to me something like an axiom about knowledge." (Cohen 2002, 312).

¹² "That something like [closure] is true, I will be taking as a primitive epistemic fact. I'm unable to formulate an argument that [closure] is true, just as I cannot provide an argument that killing innocent children without cause is morally wrong. But just as I nevertheless take it to be obviously true that we shouldn't kill innocent children without cause in spite of my inability to argue for this truth, so I will be taking the truth of [closure]." (Dodd 2012, 341).

1.2 Why Care?

Many riches flow from the renunciation of closure. A prominent skeptical threat – that one can't know, for example, that one has hands unless one knows that one is not a handless brain in a vat – is defused in a way that respects our intuition that we don't know that skeptical hypotheses are false while preventing the spread of that ignorance to more pedestrian knowledge claims.¹³ There is no need to countenance highly unintuitive “easy knowledge” inferences. A plausible solution to the problem of “bootstrapping” becomes available.¹⁴ And there is no need to resort to various theoretical, semantic, or pragmatic maneuvers in order to articulate views that incorporate closure while at the same time conceding that most, if not all, of our knowledge is acquired from fallible sources.¹⁵

But fundamentally at stake, for me at any rate, is an untenable conception of the demands that an agent must satisfy in order to know. One knows by courtesy of internal and external conducive circumstances: you don't know where your car is by seeing it in the parking lot unless you remember what your car looks like, there's adequate lighting, light travels in a straight line, and so on. Call these *enabling conditions*. Such conditions must be in place for knowledge acquisition. But must the agent also know that they are in place?

It's hard to see why. There's no reason in general why an agent S's standing in a particular relation R to some fact, which requires that condition C is realized, requires that she also stand in R to C itself. My successfully maneuvering a car through an obstacle course requires that the brake pedal be appropriately connected to the brakes. But it doesn't require that I connected them. Why should S's knowing that P, which requires that enabling condition C is satisfied, require also that she know that C is satisfied?

Some – but not all – of the enabling conditions for S's knowledge of P are implied by P itself. Suppose, for example, that P is “the gas tank is empty,” which S believes as a result of consulting the gas gauge whose needle points at “E.” An enabling condition of S's knowing that the tank is

¹³ See Chapter 10. I examine the skeptical closure argument itself in Chapter 3.

¹⁴ See Chapter 11 for discussion of bootstrapping and easy knowledge.

¹⁵ These include externalist accounts of both evidence and method, contextualism, pragmatic encroachment views, and safety accounts, among many others. It also includes brute-force reconciliations of closure with views that are not, on their face, closure-friendly by simply appending a closure principle; Sherrilyn Roush's 2005 tracking-with-closure account is an example.

empty this way is that the needle isn't stuck on "E." If it were stuck on "E," it would be so either while the tank isn't empty or while it is, coincidentally, empty. "The tank is empty" implies that the former possibility is not realized: if the tank is empty, then it's not the case that the tank isn't empty while the needle is stuck. Since it's an enabling condition of S's knowledge that the tank is empty that the needle isn't stuck (whether or not the tank is empty), it's also an enabling condition of that knowledge that the needle isn't stuck while the tank isn't empty.

This generalizes. For any enabling condition C, since knowledge of P is not compatible with the failure of C, it is also not compatible with the failure of C while some other fact is true, including $\neg P$. So $\neg(\neg C \ \& \ \neg P)$ is also an enabling condition, one that is implied by P. ($\neg C \ \& \ \neg P$) is incompatible with knowledge of P for two reasons: it is incompatible with P itself – and so with the facticity of knowledge – and it is incompatible with C, a condition of S's knowledge of P given how she acquires that knowledge.

Finding closure intuitive, one could insist that S need only know that those enabling conditions that do follow from P are satisfied. So she needs to know that it's not the case that the needle is stuck while the tank isn't empty (if, at least, she recognizes the inferential relation), but she doesn't need to know that it's not the case that the needle is stuck while the tank is empty (and so also doesn't need to know that the needle isn't stuck *simpliciter*).

But this is intuitively arbitrary, with respect to both what S needs to know and what she is in a position to know. It's *prima facie* unintuitive that she needs to know that it's not the case that the needle is stuck while the tank isn't empty but not that the needle isn't stuck. And it's similarly unintuitive that she could be in a position to know that it's not the case that the needle is stuck while the tank isn't empty but not in a position to know that the needle isn't stuck.

But if she does need to know that the needle isn't stuck then skepticism seems the inevitable result. The same considerations apply to any enabling condition.¹⁶ So S needs to know that each such condition is satisfied. That knowledge will, moreover, have its own enabling conditions. So S must know that those conditions are satisfied as well, and that each condition of

¹⁶ That is, for any such condition C, there is a condition $\neg(\neg C \ \& \ \neg P)$ that follows from P, and it seems correspondingly arbitrary that S needs to, and can, know that without also knowing $\neg(\neg C \ \& \ P)$ and so, simply, C.

that knowledge is satisfied, and so on. The imposition of such requirements seems destined for skepticism.¹⁷

1.3 Strategy

It's hard to argue against a principle that is widely treated as an epistemological axiom grounded in intuition. Even if a closure denier were to develop a theoretical account that rejects closure and that successfully answers every other conceivable objection (unlike those of Nozick and Dretske), she remains susceptible to the critique that her view doesn't preserve closure. She would likely be accused of acquiring the fruits of her theory by theft over honest toil: *of course* she has an answer to the skeptic, for example, but only because she hasn't done the necessary hard work of providing such an answer that is compatible with closure.

In order to proceed under such dialectical circumstances, the arguments I will offer against closure are, I suggest, not just intuitively compelling, but are so from a variety of epistemological standpoints. I will also show that the intuition behind closure is not as forceful as it seems at first glance, and that ultimately it does not support closure. Finally, I will show that the abominable conjunction and spreading problems directed against closure denial can be answered.

The closure denier does owe an account of when, and why, closure fails. Such an account might attempt to isolate all and only closure failures; one would certainly expect this of a full-blooded, closure-denying theory of knowledge.¹⁸ A more modest aim is to identify conditions under which closure fails, without claiming that these are the only such conditions. The result might be less satisfying than a "closure fails if and only if X" account. But, for the closure denier's purposes, no more than the "if" direction is required: it would suffice, and is more secure, to present a

¹⁷ Indeed, as we'll see in Chapter 6, skepticism results from the demand that S need only know that those conditions hold that follow from P. (At this point I'm only describing a motivation for resisting closure, not an argument for doing so. The argument comes in the rest of this book.)

¹⁸ This implies that some restricted version of closure is true. Since closure is surely not a property randomly distributed over inferences, there is some general characterization of those cases in which it does hold. Such a principle would, however, apply in a more restricted class of cases than would closure principles that are typically endorsed by those who identify themselves as closure advocates, and so will still count as closure denial in the relevant sense. (However, some philosophers who so identify themselves offer versions of closure that are in fact more restricted than those typically endorsed by mainstream closure advocates; Baumann 2012 and Roush 2012 are examples. It is, as a result, disputable whether they really should count as closure advocates.)

sufficient-but-perhaps-not-necessary account with broad theoretical and intuitive appeal. Developing a view satisfying the “only if” direction as well is likely to require a full-blooded theory of knowledge, in which case it will then run into the theft-over-honest-toil objection.

I will not, therefore, attempt to derive closure failure from some particular theory of knowledge (or class of such theories). This might disappoint some closure advocates since many of the objections against closure have in fact been directed against particular theories – especially those of Dretske and Nozick – that imply closure failure.¹⁹ But, on the face of it, such an argumentative strategy is inadequate; that T, which entails $\neg C$, is false does not imply that C is true.

So I will claim that there are conditions such that, when they are realized, closure fails, although there may well be other conditions with the same effect. Moreover, the cost of endorsing closure under those conditions will be very high indeed. The overall result will be that, far from closure denial’s being a theoretical disadvantage, it is incumbent on any defensible theory of knowledge that it accommodates closure failure.

In the remainder of this chapter I take up the challenging task of formulating a defensible closure principle. The next chapter presents a version of Dretske’s argument by counterexample, which appeals to putative counterexamples to closure. That argument will then structure the discussion for Chapters 3–9, in which I examine the different strategies the closure advocate might adopt in way of responding to Dretske’s argument.²⁰ The conclusion of Chapter 9 is that each such strategy fails, and so Dretske’s argument succeeds.

In Chapter 10 I will respond to two popular arguments against closure denial: the abominable conjunction problem and the spreading problem. In the course of doing so I’ll also examine closure-preserving contextualism and the non-skeptical invariantist closure denier’s response to skepticism.²¹ In Chapter 11 I’ll examine the bootstrapping problem, epistemic circularity, and the relationship between knowledge and justification closure.

¹⁹ Much of Hawthorne’s 2005 defense of closure, for example, is less an attack on closure denial *per se* than the presentation of counterexamples to Dretske’s conclusive-reasons account of knowledge (naturally enough, since he was responding to Dretske).

²⁰ I provide a more detailed description of those chapters at the end of Chapter 2, after Dretske’s argument is in place.

²¹ Contextualists claim that the semantic value of “know” varies across contexts of knowledge attribution. Invariantists deny that there is such variation.

1.4 Defining Closure

Notwithstanding broad agreement that some sort of closure principle is true, it turns out to be very difficult to formulate a principle that is immune to counterexamples that are recognized as such by both friends and foes of closure. I will hereafter limit attention to single-premise closure, which concerns only inferences with one premise. As is well known, there are objections to multi-premise closure that don't apply to single-premise closure (but not vice versa). So, if single-premise closure is undermined, then so is closure overall.²²

The simplest version of closure – appearing primarily in studies of epistemic logic – is that, if S knows that P and P implies Q, then S knows that Q. But this is an obviously inadequate description of actual epistemic agents. If such an agent has no grip whatsoever on the fact that P implies Q, it is highly implausible that she nevertheless must know that Q.

A common formulation declares that if S knows *both* that P and that P implies Q, then S knows that Q (call this the *Classical Formulation*). But it doesn't follow from the antecedent that S even believes Q; knowledge of Q, however, requires belief that Q. And, even if S believes Q, it is compatible with this formulation that she doesn't do so because it follows from P. She could believe Q solely on the basis of wishful thinking and so, presumably, would not know it.

A more recent, and widely adopted, formulation is offered by John Hawthorne, inspired by Timothy Williamson's suggestion that closure is an expression of the capacity of deductive inference to increase what one knows.²³ "Williamson has an insightful take on the root of epistemic closure intuitions," says Hawthorne, "namely the idea that 'deduction is a way of extending one's knowledge'."²⁴ Call this *Williamson's insight*. Here is Hawthorne's formulation, with its scope and necessity made explicit and the clauses labeled for convenience:

²² For similar reasons, I will not consider closure over inductive inferences. The problem for multi-premise closure is that small probabilities of error for each premise can add up so that, while each premise is probable, the conclusion is not. On probabilist conceptions of knowledge, according to which knowing P requires that P is probable on one's evidence, closure can fail as a result. This does not apply to single-premise inference; if P implies Q, the probability of Q is at least as high as P. (Nevertheless, Lasonen-Aarnio 2008 argues that the same problem can be extended to single-premise inferences.)

²³ Hawthorne 2004 and 2005. ²⁴ Hawthorne 2005, 41, fn. 6, quoting Williamson 2000b, 117.

Hawthorne's Formulation

Necessarily, for all agents *S* and propositions *P* and *Q*: if (a) *S* knows that *P* and (b) competently deduces *Q* from *P*, thereby (c) coming to believe *Q*, while (d) retaining her knowledge that *P* throughout, then (e) she knows that *Q*.²⁵

Clause (d) is designed to exclude cases wherein, during the course of performing the inference, the agent somehow loses her knowledge of *P* (because, perhaps, the performance somehow brings misleading evidence to light). Note that closure, so formulated, is diachronic: clause (a) refers to *S*'s knowledge at one time and clause (e) refers to her knowledge at a subsequent time. *S*'s performance of the inference takes up the intervening time.²⁶

1.5 KC

Clause (b) of Hawthorne's Formulation replaces "knows that *P* implies *Q*" in the Classical Formulation. The extent of its departure from that formulation depends on how "competent deduction" is to be interpreted.

A competent deduction might consist in only a single inferential step from *P* to *Q* or a sequence of such steps from *P* to *Q*. It is, however, better to characterize the latter as involving successive instantiations of closure rather than a single instantiation. After all, an intermediate conclusion in the sequence is the next inference's premise. If *S* doesn't know that intermediate conclusion, then she doesn't know the premise of the next inference. But, if so, it is unintuitive that she, nevertheless, must know the subsequent conclusion inferred from that premise. So knowledge of the ultimate conclusion requires that closure succeeds for each inferential step. We might as well, then, construe closure as applying to single-step inferences from the outset.

But a single-step inference seems to involve no more than the recognition that *P* implies *Q*, which is synchronic: one recognizes that *P* implies *Q* at a time, rather than across an interval of time. One might think that some span

²⁵ Hawthorne 2004, 34. In Hawthorne 2005, 29 he substituted "comes to know that *Q*" for (e). However, and as Hawthorne recognized in the earlier work, an agent could know *Q* already, before performing the inference, and so satisfy the antecedent without satisfying the consequent (Hawthorne 2004, 34, fn. 86). But such a case should obviously not count as a counterexample to closure. So I cite his earlier formulation here.

²⁶ For this reason, Hawthorne's Formulation is not, strictly speaking, a closure principle, since closure principles specify conditions on set membership (at a time). For reasons that will soon be apparent, I won't attempt to revise it further.

of time is involved, since S initially believes P and then acquires her belief in Q by performance of the deduction. But “S competently deduces Q from P” should not be understood to imply that S believes either P or Q. The former is the purpose, in part, of clause (a), and the latter of clause (c). Appeal is also made to closure in way of explaining “retraction” phenomena, wherein S, taking herself to not know Q and realizing that Q follows from P, proceeds to deny that she knows P, despite having previously claimed knowledge of P. S’s competent, single-step deduction, then, just consists in her recognition that P implies Q, without commitment to either.

To recognize that P implies Q is not merely to know that it does (and so this does not amount to a reversion to the Classical Formulation). S might know that P implies Q by testimony from a logician, without having any grip on the inferential relation herself. Some might think that this would suffice. But suppose S knows P, and knows that P implies Q by testimony. How do these pieces of knowledge fit together in order to deliver her knowledge of Q? Presumably by a *modus ponens* inference: she recognizes that, since P is true and P implies Q, Q is true as well. But perhaps she only knows that (P and (P implies Q)) implies Q by testimony as well. Then how do *these* pieces of knowledge fit together in order to deliver her knowledge of Q? Presumably by another MP inference from (P and (P implies Q)) and (if (P and (P implies Q)) then Q) to Q. But perhaps she only knows this by testimony as well . . .

If we model S’s relation to the fact that P implies Q as merely something S knows, and so just another proposition that she believes and can wield as a premise then, like Lewis Carroll’s tortoise, she will never be in a position to detach the conclusion.²⁷ S’s recognition that P implies Q cannot be construed as merely something that S knows and so believes, so that an independent disposition to infer Q from her beliefs that P and that P implies Q must be postulated. Rather, to recognize that P implies Q is to be *inherently* disposed to believe Q if one believes P, and not-P if one believes not-Q; if S does not have those dispositions then S does not recognize that P implies Q.²⁸ That disposition is then manifested by S’s believing Q when she believes P.²⁹

²⁷ Carroll 1895. Closure is sometimes represented as involving a *modus ponens* inference from S’s beliefs that P and that P implies Q to Q. As per Carroll’s story, however, that’s a mistake. S infers from P to Q, not from P and (P implies Q) to Q.

²⁸ Lasonen-Aarnio 2008 makes essentially the same point against knowledge-of-inference formulations of closure, albeit in the course of arguing for a competent-deduction formulation of closure.

²⁹ Does recognition imply knowledge? Perhaps not. Suppose that S receives excellent, though misleading, testimonial evidence to the effect that P doesn’t imply Q. On some views, excellent