
+ • +

THE JOURNAL OF PHILOSOPHY

VOLUME 0, NO. 0, 2009

+ • +

INTERNALISM, EXTERNALISM, AND THE ARCHITECTURE OF JUSTIFICATION*

One of the liveliest and most persistent controversies in contemporary epistemology is the internalism/externalism dispute. Although first introduced in connection with knowledge,¹ the terms ‘internalism’ and ‘externalism’ rapidly came to be associated primarily with competing theories of justification.² A justification theory is classified as internalist or externalist as a function of how it answers the question: “What kinds of states of affairs determine, or make a difference to, the justificational status of a belief (or other doxastic attitude)?” Are these states of affairs internalist or externalist in character?

* Earlier versions of this paper were given at the University of Edinburgh (as the “Nature of Knowledge” lecture), State University of New York at Albany, Texas A&M University, University of Wisconsin (Madison), and Heinrich-Heine University, Duesseldorf (Germany). I received helpful comments from many members of these audiences. I am particularly indebted to Holly Smith, Dennis Whitcomb, Allan Sidelle, and Jim Pryor for extensive discussion, advice, and debate on many aspects of the paper.

¹ See D.M. Armstrong, *Belief, Truth and Knowledge* (New York: Cambridge, 1973).

² See William P. Alston, “Internalism and Externalism in Epistemology,” *Philosophical Topics*, ?? (1986): 179–221, and “An Internalist Externalism,” *Synthese*, LXXIV (1988): 265–83; Michael Bergmann, *Justification without Awareness: A Defense of Epistemic Externalism* (New York: Oxford, 2006); Laurence Bonjour, “Externalist Theories of Empirical Knowledge,” *Midwest Studies in Philosophy*, Volume 5, *Studies in Epistemology* (Minneapolis: Minnesota UP), pp. 53–73; Roderick M. Chisholm, *Theory of Knowledge* (Englewood-Cliffs, NJ: Prentice-Hall, 1966, 1st edition; 1977, 2nd edition; 1989, 3rd edition); Juan Comesaña, “We Are (Almost) All Externalists Now,” in John Hawthorne, ed., *Philosophical Perspectives*, xix, *Epistemology* (Malden, MA: Blackwell, 2005), pp. 59–76; Earl Conee and Richard Feldman, *Evidentialism: Essays in Epistemology* (New York: Oxford, 2004); Richard Fumerton, “The Internalism/Externalism Controversy,” in *Philosophical Perspectives*, Volume 2, *Epistemology* (Atascadero, CA: Ridgeview, 1988), pp. 443–61, and *Metaepistemology and Skepticism* (Lanham, MD: Rowman and Littlefield, 1995); Carl Ginet, *Knowledge, Perception and Memory* (Boston: Reidel, 1975); Alvin I. Goldman, “The Internalist Conception of Justification,” in Peter French, Thomas Uehling, Jr., and Howard Wettstein, eds., *Midwest Studies in Philosophy*, Volume 5, *Studies in Epistemology* (Minneapolis: Minnesota UP, 1980), pp. 27–52, and “Internalism Exposed,” this JOURNAL, xcvi, 6 (June 1999): 271–93; Hilary Kornblith, “How Internal Can You Get?” *Synthese*, LXXIV (1988): 313–27; and Alvin Plantinga, *Warrant: The Current Debate* (New York: Oxford, 1993).

Factors that (help to) fix justificational status are generally called *justifiers*, or *J-factors*. So the central question is whether justifiers, or J-factors, have an internalist or externalist character. Numerous types of factors, however, might influence justificational status. What if some are internalist and others externalist? Which side wins, in this case, internalism or externalism? Call this a question about the *terms of engagement*. One configuration of the terms of engagement is *existential*: externalism wins if there is at least one externalist type of J-factor. Internalism wins only if *all* J-factors are internalist. A second possible configuration is *majoritarian*. That side wins that has a majority of types of J-factors of the kind it promotes. The most common terms of engagement are the first, existential kind. This is because internalists have been quite explicit in claiming that all J-factors are internalist. For example, Roderick Chisholm writes:

A consequence of our “internalistic” theory of knowledge is that, if one is subject to an epistemic requirement at any time, then this requirement is imposed by the *conscious state* in which one happens to find oneself at that time.³

I shall argue, however, that it is relatively easy to show that at least one type of J-factor is externalist. So the existential configuration makes things too easy for externalism. To level the playing field a bit, the majoritarian configuration should also be considered. Applying the majoritarian approach hinges on the admittedly knotty question of how justifiers should be typed and thereby counted. Although this challenge might produce some bickering, I welcome the second configuration along with the first. Perhaps it takes an excess of confidence for an externalist like me to offer the majoritarian configuration as one of two reasonable ways to set the terms of engagement. But it is instructive to explore the internalist/externalist controversy from a variety of perspectives.

A further complication in setting up the debate is that epistemologists commonly admit two conceptions of justifiedness: *doxastic* and *propositional*. Under the first conception, justifiedness and unjustifiedness are properties of doxastic state tokens (beliefs, suspensions of judgment, degrees of credence, and so on). Under the second conception, justifiedness and unjustifiedness are properties of ordered triples, consisting of a subject, a proposition, and a time. Even if person *S* does not believe proposition *P* at time *t*, nonetheless, the person might be justified *in* believing *P* at *t* (given his/her evidence). It is unclear whether the J-factors for doxastic justifiedness coincide with those for propositional justifiedness. Most of my formulations are

³ Chisholm, *Theory of Knowledge*, 3rd edition, pp. 59–60.

intended to be largely neutral between the two, although that is difficult to achieve without repeatedly offering alternative verbal formulations. I will not inflict that on the reader.

Let us turn to the conception of a justifier. William Alston defines a justifier as “anything that affects the justification of a belief, positively or negatively.”⁴ Read strictly, the definition applies only to beliefs. For most of the paper, however, I will adopt an expanded definition:

- (J) A justifier of any belief or other doxastic attitude is any property, condition, or state of affairs (and so on) that is positively or negatively relevant to the justificational status of that attitude.

How should we understand the term “relevance” here? It certainly should not be understood purely causally. Although some J-factors may be causally relevant to the belief in question, they are not causally relevant to the *J-status* of that belief. We may say of a J-factor that it “helps make” the belief’s justificational status positive or negative, but the “making” in question is not the usual causal making. It is better to think of the relevant kind of relevance as *explanatory* relevance. *X* is a J-factor of a given belief’s justificational status if and only if *X* *helps explain* why the belief’s justificational status is what it is.

Next let us turn to the two leading conceptions of what it means for a J-factor to be internal as opposed to external. These conceptions are *accessibilism* and *mentalism*. There is considerable variation in how accessibilism is spelled out by different writers. The following characterization is intended to convey the core of the idea.

Accessibilist Internalism. A factor or condition *J* is an internal justifier for *S*’s having doxastic attitude *D* vis-à-vis *p* (at *t*) if and only if

- (i) *J* obtains (at time *t*),
- (ii) *J* is positively or negatively relevant to the justificational status of *S*’s holding *D* vis-à-vis *p* (at *t*), and
- (iii) *J* is “directly” accessible to *S* at *t*—that is, *S* is capable of truly believing or knowing, at *t*, that *J* obtains, and is capable of knowing this “directly.”

Mentalist Internalism. A factor or condition *J* is an internal justifier for *S*’s having doxastic attitude *D* vis-à-vis *p* (at *t*) if and only if

- (i) *J* is a mental state, event, or condition of *S*, obtaining or occurring at time *t*, and
- (ii) *J* is positively or negatively relevant to the justificational status of *S*’s holding *D* vis-à-vis *p* (at time *t*).

⁴Alston, “Internalism and Externalism in Epistemology,” reprinted in *Epistemic Justification* (Ithaca: Cornell, 1989), p. 192.

Accessibilism is culled from a variety of authors, including Chisholm, Carl Ginet, and Kent Bach.⁵ Mentalism is primarily derived from Earl Conee and Richard Feldman.⁶ For each interpretation of ‘internal’, there is a correlated sense of ‘external’ that simply rejects the distinctively internalist constraint (clauses (i) and (iii) of the accessibilist definition and clause (i) of the mentalist definition) on what it takes to be a J-factor. An externalist type of J-factor is one that is positively or negatively relevant to a belief’s justificational status without having to satisfy the distinctively internalist restriction(s).

A slight emendation of the foregoing characterization of mentalist internalism might be in order. Timothy Williamson argues that knowledge is a mental state.⁷ But knowledge, being a factive state, is not purely internalist. So if mentalists agreed that knowledge is a mental state and that knowledge is a potential justifier, they would have to conclude that not all justifiers are internal. This is something mentalists would certainly resist. The remedy is straightforward, however. In the definition of mentalist internalism, simply substitute the following clause for clause (i): “*J* is a *nonfactive* mental state, event, or condition of *S* (occurring at time *t*).” This is the interpretation of mentalist

⁵ Chisholm writes: “the things we know are justified for us in the following sense: *we* can know what it is, on any occasion, that constitutes our grounds, or reason, or evidence for thinking that we know”—*Theory of Knowledge*, 2nd Edition, p. 252. Ginet writes: “Every one of every set of facts about *S*’s position that minimally suffices to make *S*, at a given time, justified in being confident that *p* must be *directly recognizable* to *S* at that time. By ‘directly recognizable’, I mean this: if a certain fact obtains, then it is directly recognizable to *S* if and only if, provided that *S* at that time has the concept of that sort of fact, *S* needs at that time only to reflect clearly-headedly on the question of whether or not that fact obtains in order to know that it does obtain”—*Knowledge, Perception and Memory*, p. 34. Kent Bach writes: “Internalism ... treats justifiedness as a purely internal matter: if *p* is justified for *S*, then *S* must be aware (or at least be immediately capable of being aware) of what makes it justified and why”—“A Rationale for Reliabilism,” *The Monist*, LXVII (1985): 246–63, p. 250.

⁶ Conee and Feldman write that according to mentalist internalism, “internalism is the view that a person’s beliefs are justified only by things that are internal to the person’s mental life ... A mentalist theory may assert that justification is determined entirely by occurrent mental factors, or by dispositional ones as well. As long as the things that are said to contribute to justification are in the person’s mind, the view qualifies as a version of mentalism.” See Conee and Feldman, “Internalism Defended,” in *Evidentialism: Essays in Epistemology* (New York: Oxford, 2004), pp. ??–??, on p. 55. Elsewhere Feldman and Conee indicate that only the mental states and conditions that obtain *at the time of belief* (or other doxastic attitude) are relevant to its justificational status: “The epistemic justification of anyone’s doxastic attitude toward any proposition strongly supervenes on the evidence that the person has *at that time*.” See Feldman and Conee, “Afterword to ‘Evidentialism,’” in *Evidentialism: Essays in Epistemology*, pp. ??–??, on p. 101 (emphasis added). Note that this quotation occurs in a recent statement of their position.

⁷ Williamson, *Knowledge and Its Limits* (New York: Oxford, 2000).

internalism I shall henceforth adopt—although I will not bother to insert the nonfactivity qualification.⁸

I. THE RIGHT-RULE ARCHITECTURE OF EPISTEMIC JUSTIFICATION

Let me now explain the phrase ‘architecture of justification’ that appears in the title. The architecture I have in mind is inspired by the uncontroversial fact that ‘justification’ is a normative term. Calling a belief justified might be compared to calling an act right, permitted, or perhaps even obligatory (in the circumstances). When we enter a normative domain, it is often appealing to think in terms of rules or principles that govern the domain in question. In the moral domain, people talk about moral rules or principles. Arguably, actions attain the status of being morally right or wrong by a combination of what the agent does, what the circumstances are, and what the governing moral rules prescribe, permit, or prohibit. Similarly, actions are legal or illegal because of what those actions consist in, what the circumstances are, and what the governing legal rules prescribe, permit, or prohibit. Finally, by analogy, perhaps belief tokens (and other doxastic-state tokens) are made epistemically right or wrong—in the sense of justified or unjustified—in virtue of the conditions, circumstances, or causes of their occurrence plus the governing epistemic rules. Numerous epistemologists propound epistemic rules and principles in precisely this spirit. Rules of epistemic rationality (for example, Bayesian rules) are more commonly displayed in the literature than rules of justifiedness, but there is plenty of precedent for justificational rules as well.⁹

In all such normative domains, the power to confer (objective) rightness or wrongness on a piece of “conduct” (behavioral or doxastic) is not possessed by any random rule. Such power is possessed only by *right*, or *correct*, rules. Anyone can propound random moral or legal rules, but the acts that fall under them will not *be* (objectively) right or wrong unless some *correct*, *authorized*, or *legitimate* rule implies that they are wrong. Why should things be different in epistemology? So I shall assume that only *right* epistemic rules make a difference to

⁸ There is also the question of whether an externalist approach to mental *content* renders mental states noninternal. I will not try to tackle this problem systematically because it is not *my* problem; it is a problem for internalism. However, note 16? below offers a few additional comments on it.

⁹ Chisholm propounded principles of justifiedness (or epistemic statuses akin to justifiedness) in all three editions of his *Theory of Knowledge* (*op. cit.*). I have advanced a right-rule framework for thinking about justifiedness in two previous works: Goldman, *Epistemology and Cognition* (Cambridge: Harvard, 1986), and Goldman, “Epistemic Relativism and Peer Disagreement,” in R. Feldman and F. Warfield, eds., *Disagreement* (New York: Oxford, in press).

genuine justifiedness. This point should be equally acceptable to both internalism and externalism.

Once this theoretical fact is acknowledged, two important questions belong on the table. First, which rules of justifiedness (J-rules) are correct or right? Second, what is the *criterion, standard, ground, or basis* for the rightness of J-rules? Illustrative answers to such questions will be discussed as we proceed. Staying at the abstract level for now, however, we can formulate the following *linkage principle*, which articulates what I have been calling a right-rule architecture. This linkage principle is also intended to be neutral as between internalism and externalism.

- (L) *S* is justified in holding doxastic attitude *D* toward proposition *p* at time *t* if and only if there are some conditions *C* such that
- (i) *S* is in conditions *C* at (or before) *t*, and
 - (ii) one or more right J-rules jointly permit a subject who is (or was) in *C* to form or retain attitude *D* toward *p* at *t*.¹⁰

Given principle (L), the justificational status of a subject's doxastic attitude toward a proposition is a matter of whether the attitude *conforms* to what right J-rules authorize in light of the subject's epistemic situation or activity. Conforming to rules is here contrasted with *following* rules, where rule following consists in conforming to rules as a result of *mentally representing* them and "verifying" that the adopted act or attitude conforms to them. I shall not consider a rule-following linkage principle at all. Some theorists might prefer a slightly different type of linkage principle than (L) (such as a rule-following principle), but I shall stick with (L).¹¹ The remainder of the paper examines prospects for epistemic internalism versus externalism in the context of the right-rule architecture articulated by (L).

II. J-RULES AND TYPES OF JUSTIFIERS

To relate the right-rule architecture to the internalism/externalism dispute, I proceed as follows. First, I examine the *kinds* of J-rules

¹⁰ An alternative linkage principle might appeal to the cognitive *processes* a subject uses to arrive at the doxastic attitude in question. Right rules would authorize the use of selected cognitive processes. This kind of linkage principle is what I favored in *Epistemology and Cognition*. However, virtually all internalists would find linkage principle (L) more congenial than a cognitive-process linkage principle. Since I am trying to make the architecture as friendly as possible to internalism, so as not to prejudice my case against it, I shall proceed in terms of (L).

¹¹ Paul Boghossian argues that the rule-following conception of epistemic justification (or rational belief) is infected by a vicious regress. See his "Epistemic Rules," this JOURNAL, CV, 9 (September 2008): 472–500. This kind of vicious regress does not threaten (L), however. Some possible reservations about (L) emerge in my own paper, "Epistemic Relativism and Peer Disagreement." But I set those reservations aside for present purposes.

that hold initial promise of having the right form and general content, according to familiar epistemological theories. The antecedents of the rule schemata to be considered will specify types of conditions declared pertinent to the permissiveness of adopting certain doxastic attitudes toward certain propositions. If a rule schema is on target, material specified in the antecedent of one of its instances will be satisfiable by various types of *conditions, circumstances, states of affairs, facts, or processes*. In virtue of satisfying the antecedent of a right J-rule, these conditions, states of affairs, facts, and so on will qualify as justifiers. Ultimately, however, they qualify as justifiers because they are *relevant* to the justificational status of various doxastic attitudes, as our definition of ‘justifier’ stipulates. However, not *all* justifiers qualify as such in virtue of satisfying antecedents of right J-rules. In addition to justifiers that satisfy antecedents of (right) J-rules, I shall later identify justifiers that include, among other things, right J-rules themselves, as well as the contents of the criterion, standard, or ground of rightness. So J-factors can qualify as such in a variety of ways. For now, however, I focus on the case of J-factors that qualify as such by satisfying antecedents of right rules. A central thesis of the paper is that many internalism-friendly rule schemas would need to be replaced by externalism-friendly rule schemas in order to be right.

It might help the reader if I say a word about how I arrived at the present form of the argument. I originally planned to develop a hybrid position, one that would espouse a synthesis, or blend, of internalism and externalism. It would concede that internalism is right about those justifiers that satisfy the antecedents of right rules; they are indeed internal. But externalism is right about the contents of the *standard* of J-rule rightness. So internalism and externalism are each right about something important. On further reflection, however, I realized that the first thesis is incorrect. For the most part, right rules need antecedents that do not restrict the indicated conditions to internal factors. Thus, the position I adopt here is not very hybrid—in keeping with my past tendency to favor externalism over internalism.

I begin by distinguishing *inferential* and *noninferential* J-rules. Inferential rules license new beliefs (and other doxastic attitudes) based on logical, inductive, abductive, or probabilistic relations to pre-existing beliefs (and so on). Noninferential rules license new beliefs in virtue of the occurrence of nondoxastic states or events such as perceptual experiences and apparent memories. Of course, any such rule will specify conditions that permit the adoption or retention of doxastic attitudes in very general terms, so all of the illustrations below (and

the ones that follow) are really just rule schemas. Particular rules would have to supply more concrete content to replace (many of) the schematic letters.

An inferential rule would take approximately the following form:

(INF) If subject S , at time t , holds doxastic attitudes D_1, D_2, \dots, D_m toward propositions q, r, \dots, s , respectively, and if proposition p bears relation R to the conjunction of q, r, \dots, s , then S is permitted to adopt attitude D_n toward proposition p at t .

To specify a precise rule of this sort, ‘ R ’ would have to be replaced with a suitable relation that can hold between propositions, for example, the relation of deductive, inductive, or abductive support, or some sort of probabilistic relation.¹² The above schema could accommodate many different theories and approaches to good inference, for example, inference to the best explanation or Bayesian inference.

A noninferential rule would take roughly the following form:

(NON-INF) If subject S , at time t , is in nondoxastic mental conditions X, Y , and Z , then S is permitted to adopt doxastic attitude D toward proposition p at t .

Of course, these are not the only types of rule schemas that might have correct instances. In fact, these particular schemas might not subsume any appropriate rules at all. I start with these rule schemas because they seem to capture a good bit of (at least informal) epistemological writing, and also to be quite congenial to internalism. Whether they are ultimately acceptable as schemas for right rules, and whether they genuinely support internalism, remains to be seen. In fact, I shall argue that even as they stand, these rule schemas are not entirely congenial to internalism (at least in the case of the inferential rule schema). Furthermore, each needs to be revised in ways inimical to internalism.

Any right rule that might exemplify one of these schemas could well need a *ceteris paribus* clause, or a “prima facie” qualification. This is not so easy to express within a right-rule framework. We would need the notion of “prima facie permission,” which is not a very standard operator. In the interest of simplicity, however, I shall largely set aside problems involving defeaters and defeasibility that invite

¹²The sample inferential rule schema I have chosen is more congenial to propositional justifiedness than to doxastic justifiedness. Doxastic justifiedness would require something to be said about the basing relation in addition to a confirmational support relation. However, propositional justifiedness tends to be the favored type of justifiedness among internalists, so my choice again follows the plan of being as accommodating as possible to internalism.

incorporation of “prima facie” qualifications. There is enough to worry about without this wrinkle.¹³

III. INFERENTIAL J-RULES: THE SUPPORT RELATION

Let us begin with inferential rules. Our initial sample of an inference-rule schema has a conjunctive antecedent. The first conjunct refers to doxastic states that a subject is in, and the second conjunct refers to some relation R of entailment, inductive support, and so on, that holds between the propositional contents of those doxastic states and the target proposition p . In considering whether any rules exemplifying this rule schema could both be right and meet internalist criteria, I begin with the second conjunct.

The fact that a certain logical, inductive, or probabilistic relation holds between specified propositions is a fact independent of any individual’s mind.¹⁴ The relation of entailment between the conjunction of ‘All men are mortal’ and ‘Socrates is a man’, on the one hand, and ‘Socrates is mortal’, on the other, is not a mental fact or condition. It is independent of any and all minds. Thus, factors or states of affairs that satisfy an antecedent of an inferential rule, under our initial proposal, will include states of affairs that are not *mental states* of subject S . These factors clearly will not qualify as internal according to mentalism. Would they qualify as internal according to accessibilism? Much depends on what accessibilists mean by “direct” access. Surely there are many relations of the intended type that are not directly accessible to most epistemic agents. Extremely complex entailment relations, holding between a huge set of complex premises and a conclusion, will not be “directly recognizable,” or recognizable at all, for most epistemic agents. So even under the accessibilist notion of ‘internal’, many conditions or states of affairs that could satisfy the second conjunct of an antecedent of an inferential rule would not qualify as internalist J-factors.

Internalists sometimes respond to this problem by trying to “internalize” the ostensibly external facts that are candidates for justifiers. Conee and Feldman¹⁵ present an example about a logic teaching assistant (TA) and a beginning logic student who are looking over a homework assignment together. One question in the assignment features a sentence that the TA and the beginning student both know to express a truth and asks whether certain other sentences are true as

¹³ Our discussion of preservative memory in section v, however, will give some attention to defeasibility.

¹⁴ See Goldman, “Internalism Exposed,” and Comesana, “We Are (Almost) All Externalists Now.”

¹⁵ Conee and Feldman, “Internalism Defended.”

well. The student is clueless, whereas the TA can easily tell that some of the other sentences are logical consequences of the original one; so the TA is justified in believing that they are true as well. Conee and Feldman comment as follows:

Again there is an internal difference between the two [the TA and the beginning student]. The difference is that the TA has justification for her beliefs to the effect that certain propositions validly follow from the original ones. She is expert enough to “see” that the conclusions follow without performing any computations... [J]ust as in example 3, relevant internal differences make the difference (*ibid.*, p. 60).

The moral they draw from this and similar cases is the following:

It is reasonable to generalize from these examples to the conclusion that every variety of change that brings about or enhances justification either *internalizes an external fact* or makes a purely internal difference. It appears that there is no need to appeal to anything extramental to explain any justificatory difference. These considerations argue for the general internalist thesis that these epistemic differences have an entirely mental origin (*ibid.*, p. 61, emphasis added).

What do Conee and Feldman mean by “internalizes an external fact”? A natural interpretation is this: *S* internalizes an external fact that *p* just in case *S* believes the (true) proposition *p*. Pursuing this hint, Conee and Feldman might ask us to rewrite our schema for inferential rules as follows, where the new version adds a belief operator to the second conjunct of the antecedent:

(INF*) If subject *S*, at time *t*, holds doxastic attitudes D_1, D_2, \dots, D_m toward propositions q, r, \dots, s , respectively, and if *S* believes that proposition *p* bears relation *R* to the conjunction of q, r, \dots, s , then *S* is permitted to adopt attitude D_n toward proposition *p* at *t*.

This rewriting of the rule schema eliminates reference to any external fact that posed a problem for internalism. Unfortunately, (INF*) is both *too strong* and *too weak*.

It is too strong because it is implausible to require people to have *beliefs* about logical, inductive or probabilifying relationships as a precondition of having justified belief in conclusions that “rely” on such relationships. To get justified beliefs in those conclusions, it is sufficient to use psychological operations or computations that “realize” or “incorporate” corresponding principles. Ordinary cognizers rarely have a theoretical grip on these principles, or even have the concepts (for example, validity) in which the principles or relations could be expressed. Hence, unlike the logic TA described by Conee and

Feldman, they lack beliefs in the appropriate principles (for example, principles about validity). Nonetheless, their inferred beliefs can be justified. So (INF*) places too stringent a demand on inference-based justifiedness.

(INF*) is too weak in a different respect. By requiring *mere* belief in the holding of the relevant relation, rules that fall under (INF*) would allow a conclusion to be justifiedly inferred even though belief in the authorizing relationship is either *false* or entirely *unjustified*. Either of these scenarios is inappropriate. Conee and Feldman tacitly acknowledge the point about justifiedness. In their discussion of the TA example, they describe the TA (in part) as follows: “the TA has *justification* for her beliefs to the effect that certain propositions validly follow from the original one” (*ibid.*, p. 60, emphasis added). This sentence implies (or implicates) that *unjustifiedly* believing that certain propositions validly follow from the original one would not guarantee a justified belief in the conclusion (“such-and-such a sentence is true”).

(INF*) is also too lax because it does not require an objective, “external” relation of validity or strong support to hold between the contents of the antecedently held beliefs and a newly inferred conclusion. Suppose someone thinks that affirming the consequent is a valid form of inference, and derives a belief in some proposition based on affirming the consequent. Clearly, his newly formed belief is not justified. But it would be justified according to (INF*). If there is no (“external”) logical, inductive, or probabilistic *fact* of an appropriate sort, simply believing that there is one does not confer justifiedness on a resulting conclusion. Thus, the problem cannot be solved by formulating J-rules that “internalize” external facts, as Conee and Feldman express it. Right J-rules must feature objective, external (logical or confirmational) relations, a requirement that cuts against (strong) internalism of all stripes.¹⁶

In response, it might be argued that right inferential rules do not need to feature any conjunct of the second kind in their antecedents.¹⁷ The right sort of formal relationship between antecedent beliefs (premises) and conclusion beliefs can be built into the antecedent and consequent of the J-rule schema itself. For example, there could

¹⁶ Subjective Bayesianism is surely a form of internalism, so does it not internalize the *R*-relation? No, I would respond. Subjective Bayesianism retains the “external” requirement that all rational or justified changes in degrees of belief must conform to objective truths of the probability calculus, especially Bayes’s Theorem.

¹⁷ This was suggested by Alan Sidelle, in a conference commentary on the paper at the University of Wisconsin, Madison, and something similar was suggested by Thomas Grundmann, at a workshop in Duesseldorf, Germany.

be a J-rule schema of the following sort: “If S holds a pair of beliefs of the forms Q and (if Q then P), then S is permitted to adopt a belief in a proposition of the form P ” This would implicitly capture the pertinent formal relationship and eliminate the need to specify any relation R in the antecedent.

Suppose for the sake of argument that this is correct, and that the requisite content relationships can be adequately expressed in right inferential rules themselves, allowing us to dispense with the second conjunct in our original (or revised) rule schema. There is still another category of external facts that will play a justificatory role. This category of facts I shall call *instantiation facts*.

Melanie believes that Albert has just testified to P , that he has no motives for being disingenuous about P , and so forth. Does Melanie’s cognitive situation permit her to infer P ? Suppose there is a right inferential rule that authorizes one to believe a proposition if it is the best explanation of other things one already believes. And suppose that the truth of P is, as a matter of fact, the best explanation of Albert’s testifying to P and other things Melanie believes about Albert.¹⁸ Then Melanie is indeed justified in believing P . What facts or states of affairs are relevant to her being so justified? They include the antecedent beliefs from which Melanie infers P . Another such fact, however, is the fact that Melanie’s antecedent beliefs and prospective conclusion belief (namely, P) jointly *instantiate* the right inferential rule we have posited. This instantiation fact is highly relevant to Melanie’s J-status in believing P (where the J-status can be understood either in the doxastic or the propositional sense). If Melanie’s cognitive situation did not instantiate this right inferential rule, believing P would not be justified for Melanie, because (we may suppose) no other right rule would license this belief. So the instantiation fact helps justify her (that is, helps make her justified) in believing P . This instantiation fact, however, is an external fact, not a mental fact. So its status as a justifier cannot be accommodated by mentalism. The same problem holds for accessibilism, because for most ordinary cognitive agents this kind of fact is not readily available for belief. Among other things, most such agents lack the concept of inference to the best explanation. The same general point can be made for other

¹⁸ Peter Lipton defends an inference-to-the-best-explanation, or abductive, approach to the justification of testimonial belief, in “Alien Abduction: Inference to the Best Explanation and the Management of Testimony,” *Episteme: A Journal of Social Epistemology*, iv (2007): 238–51. This article does not focus much on the *justificatory* properties of such an inference, but merely on the descriptive claim that this is how we make testimonial inferences. It is the justificatory angle that interests me, of course.

types of right inferential rules. Even if there is no conjunct in their antecedent that features an “*R*-relation,” there is still another type of justifier that is externalist rather than internalist. These justifiers are not states of affairs that satisfy antecedents of right rules. But, as explained earlier, being a condition that satisfies the antecedent of a right rule is only one way, not the only way, in which a condition, state of affairs, or fact can qualify as a justifier.

To avoid possible misunderstanding, I remind the reader that according to our linkage principle (L), a belief (or other doxastic attitude) is justified just in case it *actually conforms* to a right rule. (L) does not require that epistemic agents believe or know that their inferential behavior conforms to a right rule. Upon seeing this threat to their thesis, internalists might seek to internalize instantiation facts. Such an attempt would run into analogous problems identified earlier with attempts to internalize support relations.

IV. INFERENTIAL RULES: THE JUSTIFICATION CONDITION ON PRIOR BELIEFS

I turn now from the second conjunct of the antecedent of (INF) to the first conjunct. This conjunct has an important inadequacy. Epistemologists generally agree that prior beliefs (or other doxastic states) can only *transmit* justifiedness to new beliefs; they cannot *generate* such justifiedness. If initial beliefs lack justificational “juice,” no such juice can be transmitted to a new belief. Thus, subjects starting with unjustified beliefs as premises cannot arrive at a new justified belief by *any* form of inference, no matter how impeccable. Hence, the first conjunct of the antecedent of (INF) must be revised to make it a model for genuinely correct J-rules. It must read roughly as follows:

(INF**) If subject *S*, at time *t*, *justifiedly* holds doxastic attitudes D_1, D_2, \dots, D_m toward propositions q, r, \dots, s , respectively, and if proposition p bears relation *R* to the conjunction of q, r, \dots, s , then *S* is permitted to adopt attitude D_n toward proposition p at *t*.

This amendment poses a new problem for the thesis that all justifiers are internal. Having a doxastic attitude toward a proposition is uncontroversially an internal matter. But I want to question whether *justifiedly* having a doxastic attitude is equally an internal matter. If it is an external matter, this will be a big problem for internalism. I contend that all such J-factors are external.

On what grounds do I base the suggestion that justifiedly believing a proposition is an external rather than an internal matter? Does this not beg the central question between internalism and externalism?

Internalists like Conee and Feldman claim that justifiedness supervenes on mental states that occur at the time of belief.¹⁹ Does this not imply that justifiedly believing a proposition is an internal matter? No, I reply. To say that something *supervenes* on states of kind *K* does not imply that it *is* a state of kind *K*. So even on the Conee-Feldman approach, justifiedly believing is not a mental state. However, if Conee and Feldman are right in holding that justifiedness supervenes on current mental states, does that not suffice for their purposes? If justifiedness supervenes on current mental states, then there will be a set of current mental states—call it *M*—on which the justifiedness of (token) attitudes D_1, D_2, \dots, D_m toward q, r, \dots, s (respectively) supervenes. If the justifiedness of these attitudes supervenes on *M*, does not the justifiedness of attitude D_n vis-à-vis p also supervene on *M*?

I shall not try to assess the truth of this last claim. Instead, I shall argue directly against the thesis that justifiedness supervenes on mental states obtaining at the time of belief (or states directly accessible at the time of belief). Specifically, I shall argue that the J-status of a belief held at time *t* partly depends, in the general case, on what transpired in the subject's cognitive history prior to *t*. Hence, some of the J-factors relevant to a given belief held at *t* will be states of affairs obtaining prior to *t*. The *historicity* of justifiedness is a major problem for internalism, because internalism maintains that all J-factors for a doxastic attitude held at time *t* are states of affairs that obtain at *t*.

V. PRESERVATIVE MEMORY AND THE HISTORICITY OF JUSTIFIEDNESS

The argument for the historicity thesis hinges crucially on the role of preservative memory in helping to determine a belief's J-status. To establish this point, we need an excursus into the epistemology of memory. One popular theory in the epistemology of memory is the "foundational" theory, which is very congenial to internalism. According to this view, conscious events of apparent recall provide foundational (prima facie) justification for believing the content of the recall. Just as a perceptual experience that *P* makes one prima facie justified in believing *P*, so an apparent memory experience that *Q* makes one prima facie justified in believing *Q*. Pollock and Cruz are among those who endorse this view.²⁰ On this theory, memory serves as a generative source of justifiedness, not merely as a mechanism that transmits justification (like inference). The foundational theory is congenial to internalism because a recall event at time *t* generates

¹⁹ Conee and Feldman, *op. cit.*

²⁰ John Pollock and Joseph Cruz, *Contemporary Theories of Knowledge* (Lanham, MD: Rowman and Littlefield, 1999, 2nd edition).

justifiedness for believing its content at the same time t . The J-factor (the recall event) is simultaneous with the (dated) J-status to which it is relevant, as internalism requires.

Michael Huemer identifies two problems with the foundational theory.²¹ Suppose I adopt a belief in P in an unjustified fashion, for example, by wishful thinking. The next day I seem to remember that P but have no recollection of how I originally formed it. According to the foundational theory, assuming there are no defeaters, I am now justified in believing P . But that seems counterintuitive: Why should the mere passage of time—plus some forgetfulness on my part—suddenly make me justified in believing P although initially I was unjustified? Second, suppose I initially learn that P by means of a proof. Given the possibility of error, the justification is less than conclusive, but still quite adequate. After a few moments, however, I am now able to recall that P , while still entertaining the proof. So I will now have *two* justifications for P , one inferential and one foundational, thereby increasing my level of justifiedness vis-à-vis P . This, says Huemer, seems wrong. Why should the mere passage of time increase my justifiedness?

An alternative approach to memory's role in justification is the "preservation" theory. This theory does not much concern *conscious* memory, that is, conscious events of apparent recall. It mainly concerns how memory *storage* (an unconscious process) preserves justifiedness (or unjustifiedness) over time. It transmits justification rather than generating or creating it *de novo*. Preservative memory is what enables me to continue having justification for—and often knowledge of—the innumerable facts that I once learned but for which I no longer recall the evidence. For example, I now know that the sun is approximately 93 million miles from the earth. Assuming that knowledge entails justified belief (which virtually all internalists accept), I must be justified in believing this. But I no longer recall the source from which I learned it. Nonetheless I am justified because, when I first learned the fact, I knew who the source was (a reliable one, presumably), and hence was justified. What keeps me justified in believing this, even now, is the preservation of the earlier justifiedness via memory. Thus, the following J-rule is plausibly a right J-rule:

- (PM) If S has a justified attitude D toward proposition P at t , and if S retains attitude D toward P until the later time t' , via memory, then, *ceteris paribus*, S is still permitted to have attitude D toward P at t' .

²¹ Huemer, "The Problem of Memory Knowledge," *Pacific Philosophical Quarterly*, LXXX (1999): 346–57.

The *ceteris paribus* phrase is intended to accommodate the possibility that new evidence is acquired between t and t' , thereby raising or lowering the appropriate level of credence for S vis-à-vis P . In light of (PM), a person's earlier J-status toward a proposition can clearly be relevant to its later J-status. Furthermore, as the *ceteris paribus* clause is intended to imply, any evidence S observes between t and t' inclusive can affect the justificational status of S 's attitude toward P at t' . That is, it is *relevant* to the justificational status of the attitude toward P at t' . Hence, under the definition of J-factors, all such evidence observations qualify as *J-factors* vis-à-vis S 's attitude toward P at t' . Thus, events or states of affairs occurring or obtaining prior to the time of justifiedness in question (here, t') can be genuine J-factors. (Remember: being a J-factor only requires being positively or negatively *relevant* to justificational status, not being decisively determinative of it.) This is what I mean by the "historicity" of justifiedness, and it is clearly implied by preservative memory.

A further example will be instructive. Ursula acquires by testimony an initially justified belief in a certain generalization, G . This occurs at t_1 . She subsequently encounters a mountain of counterexamples to G , each of which is very obvious. The joint effect of these counterexamples, at t_5 , is the total defeat of Ursula's original justification for G . She is now justified in *rejecting* G , not in believing it; hence she is unjustified in believing it. Nonetheless, Ursula continues to believe G . She either fails to consider how the observed counterexamples bear on G (an epistemically culpable omission), or she considers their bearing but fails to appreciate their falsification of G . Much later, at t_{20} , Ursula continues to maintain the same high confidence in G as she has held right along. On the first version of the story, she merely retains this belief in storage, along with her innumerable other beliefs. (She has no conscious apparent recall of its being the case that G .) This is a clear case in which Ursula's belief in G at t_{20} is unjustified. This assessment holds even if we also stipulate that at t_{20} Ursula has long since forgotten the defeating evidence she encountered at t_5 . Her belief in G became unjustified then and it has retained this J-status ever since.

Ursula's being unjustified in believing G at t_{20} is incompatible with mentalism because Ursula has no mental states at t_{20} that constitute defeaters of her original justification for G . All such defeating mental states—her observations of the counterexamples—lie in the forgotten past. Similarly, it is incompatible with accessibilism, because at t_{20} those past defeating states are not directly accessible to her. So the unjustifiedness of Ursula's belief in G (at t_{20}) cannot be accommodated by either brand of internalism. Preservative memory is a problem for

internalism because it shows that justifiedness at a given time does not supervene on states obtaining at that time.

Now the vast preponderance of things we know at a given time depend for their justifiedness on the rightness of a rule like (PM). This holds especially for items of knowledge stored in memory. Moreover, stored knowledge constitutes the great bulk of our knowledge, because there are narrow limits to the number of things held in consciousness at any given moment. The rest of what we know must be in states of storage. Although we cannot exclude the possibility that a stored belief can have positive J-status at a given moment without application of rule (PM), this possibility is problematic (see below). It is highly likely that almost all justified stored beliefs depend for their justifiedness on (PM). Thus, rejecting (PM) (and similar rules) would open the door to rampant skepticism about our common-sense knowledge. This would bestow a huge and unwarranted gift on the skeptic.

Consider the consequences of rejecting (PM) and its ilk. If (PM) is rejected, there will be an epistemic phenomenon that I shall call *continual clearing of the justificational slate* (CCJS). There will be no carry-over of a belief's J-status from one moment to the next. Assuming for simplicity discrete moments in time, after each such moment the epistemic slate is wiped clean. A person may mnemonically retain a belief from one moment to the next, but that retention makes no contribution to the justifiedness or unjustifiedness of the succeeding moment's belief. Its J-status is determined by other factors entirely. If CCJS holds—which is implied by the rejection of (PM)—it will have catastrophic results for the number of beliefs people are justified in holding at any time.

Perhaps this is not obvious. Could not many stored beliefs be supported evidentially (that is, inferentially) by other stored beliefs held at the same time? Each would receive its justification inferentially rather than mnemonically. (Another possibility is evidential support from conscious beliefs held at the time, but, as previously indicated, conscious beliefs are vastly fewer in number and not usually content-congruent with the wide assortment of stored beliefs we learned long ago, so they would not help very much.) Assistance from other stored beliefs, however, runs into problems if (PM) is out of the picture. Other stored beliefs are helpful only if they themselves are justified. How many of them will be justified if they no longer inherit justificational juice via memory preservation? CCJS again rears its ugly head. If the inferential regress is to be stopped along familiar foundationalist lines, it presumably must stop with beliefs based on current perceptions. But beliefs based directly on current perceptions are unlikely to be of much help for most of our stored beliefs. No belief based

on a current perception confers much if any justification on my belief that the sun is 93 million miles from the earth, that Versailles is in France, that my social security number is XYZ, and so forth. Thus, we cannot be rescued from serious skepticism if rule (PM) and similar rules are rejected.

However, the preservative theory of memorial justification (which inspires (PM)) has been criticized. Huemer criticizes it by invoking Russell's five-minute hypothesis embellished with a twin case (*ibid.*). Mike1 is Huemer himself who lives in a "normal" world and is justified in believing things about the past, for example, that he ate a bagel this morning. Mike2 is somebody God created just five minutes ago, endowed with exactly the same memories and environment as Mike1, except that in Mike2's case, the memories that purport to extend beyond the five-minute window are false. Intuitively, says Huemer, Mike2 should have the same memorially justified beliefs as Mike1 has, but the preservation theory disallows this. According to the preservation theory, memory is not an originating source of justification; it can only preserve justification initially conferred by some other means. There was no other originating source of justification in the case of Mike2. So the preservation theory must decree that Mike2's bagel belief is not justified.

Another case of interest is a variant of the original Ursula case (call it Ursula2) in which she has, at t_{20} , a conscious memory experience of G 's being true. She remembers nothing about how she originally formed her belief in G or her subsequent observation of counterexamples to G . She just seems to remember that G . Some people say that Ursula2 is justified in believing G at t_{20} ; or at least that the apparent memory muddies the water, making it less clear that her t_{20} belief is *un*justified.

What shall we say, then, about the preservation theory and rule (PM)? A solution Huemer offers is a "dualistic" one, which combines elements of the foundational theory and the preservation theory. Dualism says that a belief is justified full stop if and only if one had an adequate justification for adopting it at some point, and thenceforward was justified in retaining it. How does this handle the Mike2 case? By adding the "posit" that coming to believe something by seeming to remember it is an epistemically justified way of coming to believe it. This will allow Mike2 to be justified in believing that he ate a bagel this morning (though his mode of justification—coming to believe it by seeming to remember it—will be quite different from that of Mike1). In Ursula2, dualism's verdict seems to run in the other direction. She does not justifiably adopt G (for the first time) by seeming to remember it at t_{20} , because she already believed it before t_{20} . But nor does she have an adequate justification for adopting it and

thenceforward remain justified in retaining it; the defeaters at t_5 prevent this from being true.

Huemer's dualism has some definite virtues. It captures at least some of the attractions of both the foundational theory and the preservation theory. Internalists cannot accept it, of course, because the preservative aspect of dualism still implies the historicity of justifiedness. They will defend their rejection of it by insisting that dualism makes the wrong verdict in Ursula2. Some form of dualism must be right, however, something that balances the tension between preservationist intuitions and the skeptical consequences of CCJS, on one hand, and the pull of apparent-recall intuitions, on the other.²² For our purposes, the bottom line is this. Any general theory of the epistemology of memory must assign a substantial—indeed, very prominent—role to preservative memory. The details are not crucial here. As long as preservative memory receives its due, so that something like rule (PM) is right, the historicity of justifiedness follows. That is to say, for some cases of an attitude's justifiedness at a time t , certain events or states of affairs *relevant* to that attitude's justifiedness at t occur prior to t . Under our definition of a J-factor, it follows that those prior events or states of affairs are J-factors of the attitude's J-status at t . This implies historicity.

How can internalists cope with this problem? One option is to abandon the claim that justification supervenes on *current* mental states (or that J-factors must be directly accessible from the subject's current position). Simply get rid of the temporal restrictions embedded in our original articulations of internalism. This would amount, however, to an abandonment of traditional internalism. Traditional internalism, the kind inspired by Descartes, aspires to show how ordinary beliefs can be justified from the starting-point of one's current mental states and nothing more. "Solipsism of the moment" is the epistemic situation that launches classical skepticism. Many epistemologists (including internalists themselves) identify internalism with a kind of epistemology that takes classical skeptical challenges seriously.²³ To abandon the currency constraint, therefore, is to abandon a large part of the impetus

²² Huemer puts more flesh on his dualism by suggesting what amounts to a "vectorial" theory of justification (not his term). One dimension would involve a degree of justification associated with the adoption of a belief, a second dimension would involve a degree of justification associated with its retention, and the overall level of justification would be a product of those quantities. This is an interesting proposal (which probably does not handle the Ursula case, however), but the details are not relevant to us here.

²³ See Michael Bergmann, "Externalism and Skepticism," *Philosophical Review*, cix (2000): 159–94; Bonjour, *The Structure of Empirical Knowledge*, p. 37; Fumerton, *Metaphilosophy and Skepticism*, p. 180; Stroud, "Understanding Human Knowledge in General," in M. Clay and K. Lehrer, eds., *Knowledge and Skepticism* (Boulder: Westview, 1989), pp. ???–??.

for internalism. Moreover, a currency requirement is part of the characterization of their approach offered by many leading internalists. It is clearly enunciated in the passage quoted earlier from Chisholm:

A consequence of our “internalistic” theory of knowledge is that, if one is subject to an epistemic requirement at any time, then this requirement is imposed by the *conscious state* in which one happens to find oneself *at that time*.²⁴

To abandon the currency requirement—“at that time”—would be a major departure from traditional internalism.

VI. NONINFERENCEAL J-RULES

Where does the discussion to this point leave us? The discussion of inferential J-rules has yielded important consequences for three categories of J-factors. Two of these categories, but not the third, are types of J-factors that would satisfy antecedents of (plausible samples) of right rules. First, there are states of affairs consisting of the obtaining of support relations; these would be externalist J-factors. (If no reference to support relations appears in inferential rules, however, this case is of no consequence.) Second, there are states of affairs consisting of the subject’s premise beliefs being justified, states of affairs that would help satisfy the first conjunct of right rules. By the historicity of justifiedness, however, the justifiedness of these premise beliefs will be an externalist state of affairs. Thus, these J-factors will also be externalist. Third, *instantiation* facts—which will also be J-factors—are externalist facts. (Instantiation J-factors are not peculiar to inferential J-rules.) So, at least two and possibly three types of J-factors associated with inferential justifiedness will be externalist.

It should not be concluded that *all* J-factors associated with inferential beliefs are externalist. Although a J-factor that consists in a belief’s being *justified* is externalist, an arguably distinct J-factor that consists in the subject’s *having* that belief is not externalist.²⁵ Nonetheless, internalism is losing more battles over types of J-factors than its proponents would hope. Three out of four categories of J-factors for inferential J-rules (or two out of three, if support relations are omitted) are externalist.

What happens when we turn to noninferential J-rules? This is friendly territory for internalism, one would think. Are not nondoxastic experiences, for example, perceptual experiences, the paradigm case of internal events? And are not such experiences prime examples of

²⁴ Chisholm, *Theory of Knowledge*, 3rd Edition, pp. 59–60 (second emphasis added).

²⁵ This will depend on how finely we individuate states of affairs.

things that will figure in the antecedents of noninferential J-rules? Finally, since perceptual experiences do not have to *be* justified in order to *confer* justifiedness, J-factors for noninferential belief do not encounter the earlier problem of justifiedness appearing in the antecedent. In this section we shall see if the sailing for internalism here is as clear as these considerations might portend.²⁶

A sample noninferential J-rule might have the following content: “If *S* has a visual appearance as of seeing an airplane overhead, *S* is permitted to believe the proposition that an airplane is currently overhead.” This might be generalized into the following rule schema:

(PERC) If *S* has a perceptual experience as of *P*'s being the case, then *S* is permitted to believe that *P*.

No doubt, some kind of defeasibility clause will be needed, but we set aside that complication here.

This rule schema gets one thing right. It is a general schema for noninferential J-rules that uses a systematic connection between properties of a perceptual experience and a propositional content that would figure in a permitted attitude (for example, belief). It does this by having matched propositional content in the antecedent and consequent. Unless something like this is done, it is hard to see how to formulate a perception-invoking rule with suitable generality. At the same time, this strategy gives rise to difficulties.

A first difficulty is the Sellarsian dilemma.²⁷ The phrase “perceptual experience as of *P*'s being the case” seems to refer to a perceptual state with propositional content. This suggests that it is really some sort of disguised doxastic state. If so, it is arguable that such a state cannot confer justification on a belief that *P* except by inference. If

²⁶ Of course, there are many contemporary approaches to mental states and contents, including perceptual states and contents, that would automatically make perceptual experiences “external.” As far as perceptual contents are concerned, a Russellian approach to content would imply that a perceived object is part of the content of the state of perceiving it. Since perceptual objects are usually external, the contents of the perceptual state would also be external. There is still an additional step to be taken from the *contents* of the states being external to the *states themselves* being external. Many philosophers seem ready to take that additional step, perhaps on the assumption that a state's content is intrinsic to the state itself. In the epistemological literature, however, these two steps are not taken as commonly as they are in the philosophy of mind literature. It is either assumed that the contents are not Russellian (for example, they are Fregean), or that content-externalism does not imply state-externalism, or both. These matters often remain unclarified. It is simply assumed that mental states themselves are internal. Similarly, I shall assume that perceptual experiences do not qualify as external simply in virtue of their contents being external.

²⁷ Wilfrid Sellars, “Empiricism and the Philosophy of Mind,” in *Science, Perception, and Reality* (New York: Humanities Press, 1963), pp. ???–??.

so, the rule belongs in the category of inferential rather than noninferential rules. Furthermore, it is doubtful that this sort of state can confer justification on a belief unless it itself is justified. What could make it justified? We would need an additional class of perception-based rules, and it is far from transparent what they would consist in.

Suppose we try to circumvent the preceding difficulty by holding that a perceptual experience “as of P ’s being the case” is a nondoxastic state that is the categorical basis of a disposition to believe that P . The state’s content is not identical to P ; rather, it is some sort of perceptual, nonconceptual content. Nonetheless, the state tends to cause in the subject a belief that P . Although this circumvents the preceding problems, it is unclear whether the rule schema is right. Is an epistemic subject always justified in believing P if she is in a perceptual state that inclines her to believe that P ?²⁸

I argue to the contrary. Suppose Reginald is in a visual state that inclines him to believe an identity-ascribing proposition to someone he sees in a crowd. It might be the proposition, “That’s my long-lost love, Penelope, in the crowd” or the proposition, “That’s the guy who has been pursuing me with apparent murderous intent.” In each variant of the case, I imagine, Reginald has a visual experience that is partly caused by emotion. In the first variant it is an emotion of desire and longing. Reginald is so anxious to find Penelope that his visual experience is “as of Penelope being in the crowd,” although the visual cues leading to this visual experience are not at all indicative of its being Penelope. In the second variant, fear of the pursuer distorts Reginald’s perception, so that he has a visual experience “as of the murderous assailant being in the crowd,” despite the poverty of the visual stimulus. In both cases, I submit, Reginald is not justified in holding the indicated belief.²⁹ This cuts against the rightness of the proposed rule schema (under the proffered interpretation). Could it be amended so as to obtain a right rule schema?

What is needed to obtain an acceptable rule or rule schema is a characterization of the perceptual cues or operations that properly license belief in such a proposition. What might be examples of such cues or operations? Which perceptual cues would provide suitable perceptual *evidence* for what is believed? Here is an example, taken

²⁸ Of course, initial justifiedness could be defeated by concomitant beliefs. But that is not the kind of problem I have in mind here. I am not worried about mere defeasibility, but by the worry that even *prima facie* justifiedness is not guaranteed.

²⁹ Not long after incorporating these examples into this paper, I heard Susanna Siegel present a paper with similar examples, many of which she also considered to involve unjustified belief. See Siegel, “Cognitive Penetrability and Perceptual Justification” (unpublished, Harvard University).

from a simpler, well-studied domain, namely, depth perception. The visual system has multiple types of cues by which it estimates an object's distance from the perceiver. Two such cues are *binocular disparity* (the difference in images presented to the two eyes because of their different perspectives on the object) and *eye convergence* (the angle between the two eyes' lines of sight, that is, the extent to which the two eyes are turned toward each other to fixate an object, especially a nearby object). These conditions provide evidence to the visual system in its computation of an object's distance. Perhaps a subject whose visual system undergoes the right sorts of binocular disparity cues or eye convergence cues is thereby justified in believing that the observed stimulus is a certain distance from him. Borrowing pertinent details from vision science, we might be able to construct a J-rule that fits the circumstances in which subjects make accurate distance judgments based on binocular disparity or eye convergence.³⁰

Rules of this sort, however, would not be congenial to internalism—neither to accessibilism nor to mentalism. The problem for accessibilism is that the cues associated with binocular disparity and eye convergence are not conscious. Presumably they are not directly accessible to the subject. What about mentalism? It is not clear whether this is an equally severe a problem for mentalism because it is not clear what qualifies as “mental.” It is highly plausible, however, that being the mental state of a person requires more than being a representational state of some cognitive *module* in that person's brain. In most cases concerning cues or operations of the visual system that enable people to form accurate visual beliefs, these cues are only states or operations of one or more modules of the visual system, not ascribable to the person, and hence not mental states of the person.³¹ If this is correct, it seems difficult if not impossible to construct any plausible perception-based J-rules with antecedents that meet internalist strictures.

The upshot is that identifying a noninferential J-rule schema that meets internalist desiderata is not a slam-dunk. We have not definitively shown that such a J-rule cannot be had. But the prospects are murkier than is commonly thought. Thus far, then, we have identified several categories of J-factors that are clearly externalist, only one category that is clearly internalist, and another (perceptual experiences) where the classification is not wholly resolved. Under the “existential” terms of engagement, in which externalism wins if there

³⁰ Thanks to Jack Lyons for suggesting these examples. His *Perception and Basic Beliefs* (New York: Oxford, in press) provides an illuminating and subtle analysis of closely related issues.

³¹ For further discussion of this point, see Lyons (in press, chapter 3).

is even *one* externalist category of J-factors, externalism clearly wins. If we apply the majoritarian terms of engagement, under which externalism wins if there are more externalist than internalist categories of J-factors, externalism still wins at this stage. We are not finished, however. There are still additional categories of J-factors for which the analysis tilts toward externalism.

VII. OTHER EXTERNALIST J-FACTORS: RIGHT J-RULES AND THE
STANDARD OF J-RULE RIGHTNESS

I shall now argue that under the right-rule architecture and our definition of justifiers, right J-rules should qualify as justifiers. The argument is fairly straightforward. If a particular belief is justified, and it is justified in virtue of instantiating a certain right J-rule, then the content and rightness of that J-rule are positively relevant to its justifiedness. Given our definition of justifiers, this implies that the J-rule is a justifier of the belief's justifiedness. The argument can be amplified a bit by recalling that the species of relevance in question is explanatory relevance. It is explanatorily relevant to the belief's being justified that it instantiates the rule in question and that this rule is right. The fact that the rule has a certain content and is right helps explain *why*, or how it transpires, that the belief is justified.

Some people might respond as follows. "It is a mistake to regard either J-rules themselves or their rightness as J-factors. What is relevant to a belief's justifiedness are states of affairs that satisfy the operative rule's antecedents and the associated instantiation fact. But it is a *category mistake* to add the rules themselves, or their rightness, to the list of a belief's J-factors."³²

In reflecting on this matter, it may help to consider an analogue from another normative domain. When we ask whether an accused is guilty or innocent of a specified legal offense, what kinds of factors would (help) make it the case that he is guilty or innocent? Obviously, many material facts—what deeds were done, in what circumstances, in what state of mind, and so on—are relevant. In addition to these material facts, however, another group of relevant factors are legal facts. What is the content of the law concerning the offense in question? Is it a valid law (under the governing federal or state constitution)? What are the legal precedents? The material facts alone do not make it the case that the person is guilty or innocent. The material facts alone do not explain *why* the person is guilty (if he is). The contents of the law are equally relevant. In other words, legal rules are

³² Holly Smith has defended this position in conversation.

(partial) “determiners” of guilt and innocence as much as material facts are (which is not to deny the differences between them).

Given that right J-rules are justifiers, is their being right an internalist or externalist matter? That depends on what makes a right J-rule right. J-rule rightness is something conferred by the criterion, or ground, of rightness. If a J-rule meets the criterion, or ground, it is right; if it does not meet the criterion, it is not right. Is meeting such a criterion an internalist matter? Let us examine this first under the mentalist approach and then under the accessibilist approach.

Is meeting a rightness criterion a matter of an individual subject’s mental states? Presumably not. For any selected criterion, whatever its content, meeting or failing to meet the criterion is an objective matter of fact, not a matter of somebody’s belief or other mental state. Is meeting a rightness criterion directly accessible to any pertinent subject? Direct accessibility consists of being able to determine correctly, though directly, that the state of affairs in question obtains. Is the rightness of a J-rule determinable correctly and directly? The question should be clarified: determinable by whom? I assume that the answer is: by the epistemic subject, a person on the street. The person on the street cannot decide correctly which J-rules are right. Ordinary people have never considered such specialized issues (neither have many philosophers). The person on the street never dreams of these questions, much less their answers. He or she would have no clue as to what a standard or criterion of J-rule rightness is; nor would he or she be in a position to determine which J-rules meet such a standard. Hence, on neither approach to internality does the rightness of a J-rule qualify as an internal state of affairs. In this category, like most of those examined earlier, J-factors turn out to be external rather than internal.

Finally, what about the criterion, standard, or ground of J-rule rightness? Is this criterion itself a justifier of beliefs and other doxastic attitudes? Again, the answer seems to be affirmative. After all, the justificational status of each token attitude is a function, positively or negatively, of a specific right J-rule, or the total set of right J-rules. But the membership of the set of right J-rules is a function of the standard, ground, or criterion of J-rule rightness. So, by transitivity, the justificational status of each token attitude depends on the standard, ground, or criterion of J-rule rightness. Hence, this standard falls in the category of a justifier.

Is this justifier internalist or externalist? That is, does the *content* of the rightness-standard invoke internalist or externalist properties or states of affairs? An obvious externalist candidate for the rightness standard would invoke *truth-conduciveness*—for example, conduciveness

to a high truth-ratio (or veritistic value³³) among the beliefs that would result from conforming to a rule. The standard might say, for example, that a J-rule *V* is right if and only if, among the J-rules that compete for a certain category of J-rules, conformity with *V* would maximize the truth-ratio (or veritistic value) of the doxastic attitudes that would be generated if epistemic subjects conformed their doxastic “conduct” to *V* (as compared with the truth-ratio that would be generated by conforming their doxastic conduct to alternative rules).

Up to this point I am not trying to argue that the (correct) standard of J-rule rightness *is* externalist. I am just trying to clarify the issue and sustain the thesis that the standard itself (that is, its content) is among the justifiers of token beliefs. The content of the standard is one of the things that (directly or indirectly) help to fix the justificational status of doxastic attitudes.

No doubt other candidates for the standard of J-rule rightness would feature internalist contents. However, it might take more than a little delicacy to formulate some plausible ones. Some internalists—namely, mentalists—would probably resist the idea that a rightness criterion is needed. Conee and Feldman defend internalism primarily through the vehicle of a supervenience principle: their principle “S.”³⁴ They do not explicitly claim that this supervenience principle provides a *full* theory of justifiedness, but they strongly imply this. It is important to indicate, therefore, that this is incorrect. A supervenience principle like “S” does not provide everything a theorist wants to know about justifiedness.

To explain this important point, consider an analogy. Suppose a moral consequentialist holds that the moral rightness of actions supervenes on their consequences. Is this a theoretically complete story of moral rightness? Hardly. It leaves entirely open *which* (types of) consequences confer rightness. Is rightness fixed by the pleasure that actions cause, or by different types of consequences (for example, the surprise they occasion)? Even being told that rightness supervenes on hedonic consequences leaves crucial questions unanswered: Does *more* pleasure or *less* pleasure conduce to rightness? The answer might

³³ What is meant by “veritistic value” is explained in my *Knowledge in a Social World* (New York: Oxford, 1999). See note 35? below.

³⁴ Conee and Feldman’s supervenience principle *S* reads as follows: “The justificatory status of a person’s doxastic attitudes strongly supervenes on the person’s occurrent and dispositional mental states, events, and conditions” (“Internalism Defended,” p. 56. This principle does not explicitly restrict the supervenience base to states occurring at the same time as the doxastic attitudes in question. Another recent statement of theirs, however, quoted above in note 6, indicates that they do intend this restriction.

seem obvious. But supervenience per se only guarantees that if two items share the same base properties, they also share the same supervening properties. It does not necessarily tell us precisely *which* properties supervene on *which* base properties.

It is similarly under-informative to be told that epistemic justifiedness supervenes on a subject's mental states. Even if this were true, it would not tell us everything relevant. We also need to be told which combinations of mental states give rise to *positive* justifiedness and which combinations give rise to *negative* justifiedness. Knowing that mental states M_1, \dots, M_k jointly fix the J-status of an attitude toward q would not indicate whether the resulting J-status is positive or negative. This is why right J-rules are needed, and, at a higher level, why a criterion of J-rule rightness is theoretically desirable. Specifying a supervenience base (for example, the class of mental states) leaves a great deal indeterminate. This indeterminacy is what a truth-conduciveness criterion, or other such criterion, needs to resolve. I do not mean to imply that no internalist rightness criterion could possibly be up to the challenge. But it must be shown that an internalist rightness criterion *is* up to the challenge. One cannot dismiss the challenge on the ground that no rightness criterion is needed.

I next sketch an argument in support of an externalist, truth-conduciveness approach as the right approach to the job. The argument is simply that such a criterion would provide a satisfying explanation of why epistemologists find certain candidate J-rules right (or roughly right) and others wrong. I shall illustrate the point first with respect to noninferential J-rules and then with respect to inferential J-rules.

Why do so many examples of non-inferential J-principles center on perceptual experience, especially where the epistemic subject is in "good" perceptual circumstances? Because these are cases in which beliefs formed in accordance with these experiences are usually true. Or because these are cases in which epistemologists assume, and assume that others assume, that believing in accordance with the recommended principles leads predominantly to truths. What I say is intended to apply not only to reliabilists, who explicitly invoke a truth-conduciveness story. I also speak of epistemologists who offer an entirely different theory, but still adduce similar principles. I claim that the underlying appeal of these J-principles is a tacit recognition that they are truth-conducive, even when this is not the *official* doctrine being endorsed.

A similar example can be drawn from inferential J-rules. Which R-relations do epistemologists regularly choose to incorporate into inferential rules? Relations such as deducibility, strong inductive

support, and inference-to-the-best-explanation. We do not choose epistemic rules in which the support relation is weak probabilistic support, or inference-to-the-worst-explanation! Why not? Because these rules would yield a much lower rate of true-belief formation than the preferred principles. Or so it is assumed. What apparently underpins the choice of congenial J-rules is their tendency to produce a fairly high ratio of true beliefs. So that is the appropriate choice of a standard of J-rule rightness. Even if some of these rules are criticizable for one reason or another, the best unifying explanation of their appeal to epistemologists is that epistemologists expect good epistemic practice to be the kind of practice that is comparatively truth-conducive.³⁵

If all this is correct, the J-factor at the level of the rightness criterion is an externalist one. This adds yet another type of externalist J-factor to our list, covering many categories of justification determiners. So externalism appears to win the dispute with internalism even when the playing field is made more congenial to internalism than it normally is.

We have not examined any possible rightness criteria from the internalist camp. One such candidate comes from Richard Foley.³⁶ According to Foley a belief is rational if and only if it conforms to the believer's most deeply held principles about promoting the epistemic goal of now believing the truth. Translating this into a proposal for justification rather than rationality, this suggests that a J-principle is right just in case it is one of the believer's most deeply held principles about getting the truth. Being deeply held certainly looks like an internalist matter. It takes little reflection, however, to notice that Foley's criterion is highly relativistic. A principle deeply held by one epistemic subject is not necessarily deeply held by another. This makes Foley's approach ill suited to the objectivist, nonrelativistic spirit of our entire framework. As explained at the outset, our right-rule architecture links justifiedness to conformity with *right* rules, and it looks as if Foley's approach gives no sense to person-independent rightness of rules. A given rule may be right-for-Smith but not right-

³⁵ The explanation in the text cites only the propensity of the exhibited rules to produce true *categorical* beliefs. This cannot be the entire story, because many prominent epistemologies feature *graded* beliefs as appropriate outputs of good J-rules. I have sought to capture the value of graded beliefs in a truth-conduciveness framework by introducing the broader notion of *veritistic value*, in which even graded degrees of belief assigned to truths are good epistemic consequences (as compared with lower degrees of belief assigned to the same truths). For a systematic development of this approach, see my *Knowledge in a Social World*.

³⁶ Foley, *A Theory of Epistemic Rationality* (Cambridge: Harvard, 1987).

for-Jones, with no sense made of rightness *simpliciter*. It is not clear, then, that one can treat Foley's criterion as a *rightness* criterion. Instead it is a *relative-rightness* criterion, a different kind of beast altogether.

VIII. CAN INTERNALISM BE SAVED BY A NEW DEFINITION OF JUSTIFIERS?

After reviewing my case for externalism, readers might suspect that I have been unfair to internalism from the outset in my characterization of justifiers. I adopted Alston's characterization of justifiers as any states of affairs that are positively or negatively relevant to justification. This is too broad, it might be objected. In fact, Alston himself gives a rather different characterization of justifiers in another paper. In "An Internalist Externalism," Alston introduces the notion of a "ground" as a state on which a belief is "based" (*op. cit.*). Prime candidates for grounds are the subject's perceptual experiences and other beliefs. Suppose now that we restrict "justifiers" to *grounds* in Alston's sense. Would not this substantially change the terms of the debate, and would not internalism win the debate so defined? If the outcome of the debate is so easily changed, should not our original definition be deemed arbitrary and prejudicial?

There are several reasons not to choose grounds as the linchpin of justifiedness. First, a person can have a justified belief without having any ground at all for the belief. Someone can justifiably believe a certain factoid for which he once had evidence he has since forgotten entirely. Nothing qualifies as his (current) grounds for the proposition, but he still believes it justifiedly. True, some writers say that memory preserves his original justification, that is, his original grounds. Therefore, he must still "have" them. There is no pertinent sense of "possession," however, in which a person who has forgotten his original grounds still possesses them. This just shows that *being* justified cannot be equated with *having* a justification in the sense of having a ground. Notice, moreover, that if having had grounds in the past sufficed for justifiedness, this would provide no succor for internalism, because internalism denies epistemic force to the past.

Another reason not to privilege "grounds" as the only factor to consider in analyzing the contours of justifiedness is that, as we have seen, grounds are just one among many types of factors relevant to a belief's justifiedness. In addition to grounds, there are (1) support relations, (2) instantiation facts, (3) historical events that bear on premise-belief justifiedness, (4) right rules, and (5) the ground or standard of rightness. In addition to these, there is one we have not much discussed: cognitive processes that constitute or underlie the "basing" relation. The basing relation is critical for doxastic justifiedness in particular. If you justifiably believe both Q and (if Q then P), you have the "makings"

(grounds) for a justified belief in P . But even if you go ahead and believe P in these circumstances, it does not follow that your belief in P is doxastically justified. It would not be justified if you ignored Q and (if Q then P) and instead came to believe P by fallaciously inferring it from your prior beliefs in R and (if P then R). So, current grounds are not the be-all and end-all of justifiedness, and the dispute between internalism and externalism should not be equated with the dispute over the current-grounds factor in justifiedness.

IX. CONCLUSION

Two possible terms of engagement were delineated for addressing the internalism/externalism dispute: the existential and the majoritarian configurations. I have argued that there are numerous categories of J-factors that qualify as externalist, many more than J-factors that qualify as internalist. If at least one of those arguments is correct, externalism wins the dispute as judged by the existential terms of engagement. Even by the more stringent, majoritarian term of engagement, externalism comes out on the winning side of the ledger.

ALVIN I. GOLDMAN

Rutgers: The State University of New Jersey