The Possibility of Internalist Epistemology

Kurt L. Sylvan

1.   Introduction

Internalism holds that epistemic justification is determined by what is internal to the mind, not by facts about the mind-independent world. The implications of internalism depend upon what it is to be internal. Many internalists hold either that (1) to be internal is to be accessible to conscious awareness, or that (2) to be internal is to be a (non-factive) mental state.[1] I reject these views, but I think internalism is defensible if founded on a different conception of internality. My aim is to defend this thought.

The kind of internalism I prefer is rooted in rationalist ideas that have been neglected in recent epistemology, despite inspiring internalist projects in cognitive science.[2] Ignoring rationalist insights has, I think, damaged the prospects for internalism, by saddling internalists with empiricist burdens. Internalists can refuse these burdens by accepting a better philosophy of mind. Rather than looking to 17th century rationalists, I look to Kant, and especially his idea that normativity is grounded in the *constitution* of agency and mind.

I think a Kantian approach can avoid stock objections to internalism. It can also address new challenges from Littlejohn (2012), who prefers externalism because he thinks it better aligns with the philosophy of normativity. Kantian internalism embraces Littlejohn's methodology, since it fits well with Kantian ethics and philosophy of practical reason. Internalism hence needn't rest on ignorance of the philosophy of normativity.

With these ideas in mind, here is the plan. In §2, I get back to basics and consider more ways to satisfy the core internalist template than have become standard. I draw attention to rationalist insights about how normativity could be grounded in the constitution of mind, and hypothesize that if internalism respects these insights, some problems disappear. I develop this hypothesis in §3 with a Kantian view. I further defend this view in §4. I conclude in §5 by explaining why Kantian internalism answers Littlejohn's challenges.

2.   Groundwork for the Metaphysics of Epistemic Normativity

Consider the basic internalist template:

I: Epistemic justification is determined by what is internal to the mind.

This template could be read in two ways if we distinguish *ex ante* justification—i.e., there *being justification to φ*—and *ex post* justification—i.e., *φ-ing justifiedly*. Many contemporary

---

[1] For this division, see Littlejohn (2012: 12-16) and Feldman and Conee (2001).

[2] 17th century rationalism inspired Chomsky's appeal to innate knowledge of syntax to explain linguistic competence—hence he initially called his paradigm 'Cartesian Linguistics' (Chomsky (1966)). He frames this paradigm as an internalist one in Chomsky (2000). This paradigm has been highly influential in cognitive science. Contemporary epistemic internalists usually ignore it. But one exception is Pollock and Cruz (1999), who have a relative of the view below.

internalists treat *ex ante* justification as prior to *ex post* justification.[3] But internalists can prioritize *ex post* justification: just as Kantian ethics emphasizes the *good will*, so one neglected form of internalism emphasizes *quality of thought*, a property of one's thinking determined by whether it is *well-structured*. I defend a view of this kind.

Understood either way, our template invites four questions:

Q1: What is it to be internal to the mind?

Q2: What is internal to one's mind?

Q3: What is it for justification to be *determined* by the internal?

Q4: How could *justification* (not *excuse*) be determined by the internal?

Progress is possible for internalism if we consider more answers to Q1-Q3 than have become entrenched. Giving Q4 serious attention also helps: Littlejohn is right that internalists mostly either ignore Q4, by resting content with arguments from intuitions, or make assumptions undermined by recent philosophy of normativity (e.g., that justification = blamelessness).[4]

Here epistemic internalists can, I think, learn from internalists in practical philosophy, particularly in the Kantian tradition. To achieve harmony with the philosophy of normativity, epistemic internalists should consider developing a view analogous to Kantian meta-ethical views that ground moral normativity in *norms of reason*.[5] They should also, I think, learn from Kantian ethicists who view fundamental moral demands as demands of *respect*. Something close to Kant's own epistemology, I'll suggest, has the required features. It offers better answers to Q1-Q3 and mirrors defensible views in practical philosophy.

I'll now set the stage for this idea by revisiting Q1-Q4.

On Internality

What is it to be internal to the mind? A naïve analysis would treat internality *spatially*, so that being internal to the mind = being in the head. Epistemic internalists don't tend to understand internality this way, however. It is helpful to remember why.

One reason is that epistemic internalism shouldn't require that justification is grounded in *brain states*. For traditional internalists, the brain is just as external as the world. Even if minds must be physically realized, being relevantly internal doesn't consist in being *inside* physical realizers. Spatial properties shouldn't have any deep significance for internalists. This point matters, because it suggests epistemic internalists needn't spurn some 'externalist' ideas in philosophy of mind, which assume a spatial picture. They could allow justifiers to sometimes include parts of the world if they insist that these *count as justifiers* only when suitably related to minds. Meta-ethical internalists after Schroeder (2007) have stressed this point: internalists needn't hold that reasons *are* internal states, but only that something *counts as a reason* only if it is suitably related to internal states. Epistemic internalists can hold an analogous view, I'll suggest below.

While the spatial analysis of internality is wrong, it is a special case of a wider idea with better instances. The spatial view might be plausible if minds were material objects, since what's internal to a material object are its *material parts* and their features and

[3] See Fumerton (1995) and Conee and Feldman (2002). While this direction of analysis is typical, there are exceptions, like Pollock and Cruz (1999) and Wedgwood (2002).
[4] Some internalists (e.g., Fumerton (1995)) do better, but they also are skeptics.
[5] See especially Smith (1995), Korsgaard (2009), and Markovits (2015).

organization. This suggests a more general idea: what is internal to any X are X's *constituents*, their non-relational properties, and organization. This generalization doesn't entail a spatial picture, since constituency can be more abstract: consider syntactic constituents of sentences or constituents of works of music. Because mental properties are multiply realizable, mental constituents are plausibly constituents of a more abstract sort. Hence one might consider:

> **Constitutive Internality**: Being internal to S's mind =$_{df}$ being grounded in the constituents of S's mind, their non-relational properties, and their organization.[6]

Since constituents can plausibly be contingent and temporary—consider parts of cars or computers—Constitutive Internality doesn't commit internalists to thinking justification is grounded in *essential* features of minds, like the bare capacities for intentionality and consciousness. Other features of minds could qualify.

Granting Constitutive Internality, some common versions of internalism shouldn't be the default versions. Typical *mentalist* versions of internalism make justification determined by mental *states* (e.g., experiences) and their qualities. States and their qualities can be mental constituents. But it is unobvious why they are the only relevant ones. Similar points go for typical *accessibilist* views, which invoke *introspectable states* as justifiers: introspectable states can be mental constituents, but it is implausible that *only* introspectable states are. The mind's nature and organization are deeper features that are not straightforwardly accessible. Internalists could prioritize these features.

On the Internal

What is internal to minds? Most contemporary literature focuses on fleeting states like perceptual experiences. The most prominent forms of internalism seek states with intrinsic reason-giving force—*seemings* or states of *direct acquaintance*. Externalists and coherentists argue that this quest is futile. I agree, but one needn't go externalist or coherentist. One could go more deeply internalist.

Short-stay occupants of minds like experiences aren't all that matter. Essential and structural features of minds also matter. Some historical internalists—Kant and some rationalists—prioritized such features. *Rational capacities*, like the capacity of *understanding*, represent one important historical option. Structural features, like Kant's Forms of the Understanding (the Categories), could also play a role. To provide an up-to-date story, internalists can consider relevant cognitive science—especially recent work on innate structures in the mind, like the capacity for carving up the perceived world into *objects* and *features*.[7] These structures could play the epistemic role Kant assigned to the Categories: they could make empirical understanding—and justification for empirical beliefs—possible.

Treating experiences as the ultimate grounds of justification saddles internalists with *Humean* burdens. Hume was wrong: minds aren't mere bundles of experiential states. Minds have more internal structure. Subjects also actively structure their mental lives by exercising their capacity to reason. Internalists can incorporate these non-Humean points.

---

[6] It is worth stressing that internalists who accept Constitutive Internality can allow that the propositional content of some mental states (e.g., the belief that water is wet) is possessed extrinsically. Constitutive Internality only suggests that justification for believing such propositions *derives from* the internal. And the derivativeness here needn't be understood *inferentially*: we can have non-inferential justification for the belief that water is wet. The relevant non-derivative/derivative distinction crosscuts the non-inferential/inferential distinction; see Silins (2011: 355ff).

[7] See Spelke (2022), whose framework echoes Kant.

To bring out some possibilities, I'll focus on *cognitive* constitution. Your cognitive constitution includes (i) your innate and learned cognitive capacities, like your capacity for *reasoning* or your capacity for *object permanence*, (ii) any aspects of your mental life that manifest these capacities rather than *external interference* (say, a disruptive environment) or interference from *your physical condition* (tiredness or illness, for example), and (iii) the organization of your capacities—e.g., the *functional relationship* between your capacities for experience and thought.

Humeans take our mental life to be organized by laws of association and take our records as cognizers to include whatever happens in our minds. Internalists can reject this outlook in different ways. They may hold that cognitive life is more or less governed by rational principles. They may do more or less to explain non-rational mental phenomena by external factors, like how linguists treat everyday mistakes in the articulation of language ('performance errors') as not manifesting defects in our understanding of language. They may also have a more or less essentialist picture of cognitive identity. They may hold that our identity as cognizers is exhausted by innate rational principles, like traditional rationalists. Or they could mimic existentialists by allowing cognizers to actively determine some of the principles that govern their minds—for example, by allowing subjects to set different thresholds for what counts as *sufficient evidence*.[8]

Given these views, one could argue that whether a *state* should count as internal to a cognizer is partly determined by whether its presence in her mind genuinely reflects her constitution. Non-rational states like intrusive thoughts, hallucinations, or mere lapses due to exhaustion plausibly count as *external* to one's self for this reason: they are *arational* and not part of one's record as a cognizer. Here appealing to constitution is illuminating, since it is clear that not everything that is *in* something (e.g., a bug) is *part* of it (e.g., a computer). Application of this point to the mind is a key rationalist insight that contemporary internalists largely neglect. While the mind may not be prepackaged with true beliefs in the way that traditional rationalists believed, it does plausibly have a rational nature to which serious errors are not easily attributed, especially when external explanations are available.

On Internal Determination

Besides seeking a richer conception of internality, internalists can develop a more fine-grained account of internal determination. To bring this out, we can start with an approach from Conee and Feldman (2008). They suggest that the way internal factors determine justification is captured by *epistemic principles* like:

> (Experience) If S has an experience with the content that p, there is *prima facie* justification for S to believe p.

Conee and Feldman acknowledge four views about such principles: they are either

> broadly analytic truths underwritten by logic and *a priori* probability (Strawson (1952), Carnap (1962)),

> synthetic *a priori* truths discoverable through reflection in the way that ethical truths are (Chisholm (1957)),

> deeply held commitments or intuitions (Foley (1987), BonJour (1998)), or
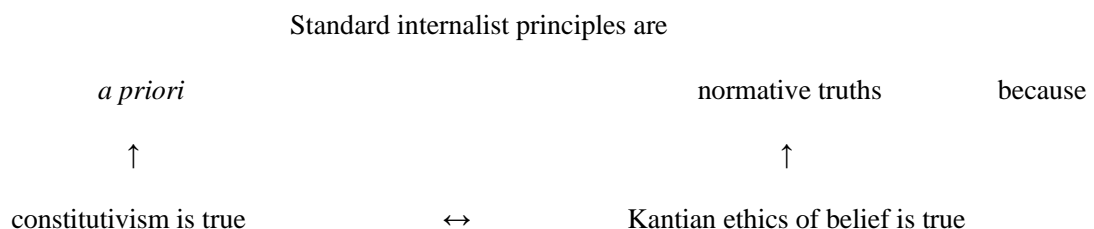
---

[8] See Callahan (2021).

principles of inference to the best explanation (Conee and Feldman (2008), McCain (2014)).

Further possibilities emerge if we take the analogy between epistemic principles and ethical principles seriously. Recall that there are two levels at which one can think about ethical principles. One can develop first-order views, like Kantian ethics and consequentialism. One can also develop meta-ethical views, like Korsgaard (2009)'s attempt to ground moral principles in requirements of reason. By analogy, we could give a first-order view about epistemic principles—an ethics of belief—or a meta-epistemological view—a meta-ethics of belief. One could also pursue projects at both levels, seeking mutually supportive views. These points help structure Conee and Feldman's list, which contains views of different genres: Chisholm had an 'ethics of belief', while Carnap and Strawson had meta-epistemological theories.

There are more views to explore given the analogy with ethics. Just as one might appeal to consequentialism to provide a systematic first-order justification of intuitive ethical principles, so one could defend epistemic consequentialism, as many have explored recently.[9] But there is also a more internalist option: one could seek an epistemic analogue of Kantian ethics.[10] Parallel remarks apply at the meta-normative level. Just as one might give a naturalistic account of ethical principles, so one might give a naturalistic account of epistemic principles. But again, there is a more internalist option: developing an epistemic analogue of views that ground moral normativity in principles of reason.

Although first-order and meta-epistemological theories are distinct, they interact. This point is central to internalism/externalism debates in meta-ethics. Consider how meta-ethical internalists feel that externalists like Parfit (2011), who reject internalism on first-order grounds, beg the question: they think meta-normative questions are prior. Skeptics like Fumerton (1995) share this methodology. But internalists can reject these strictures. Internalists could grant first-order intuitions (unlike skeptics), but still puzzle over why these intuitions are reliable. They could then solve this puzzle with a suitable meta-epistemology.

Kant's Critical philosophy solves a related puzzle. Kant thought the only way to secure both the objectivity and the *apriority* of epistemic principles is with transcendental idealism. Thankfully, we don't need transcendental idealism to solve our puzzle. Less radical Kantian ideas help. Part of what is needed is a Kantian ethics of belief: epistemic norms are norms of *respect for truth*, just as ethical norms are norms of respect for persons. But one shouldn't stop here, since one must explain how respect for truth could be internally determined. So I recommend a further Kantian view: epistemic norms are grounded in *constitutive principles* of cognition ('constitutivism'). Hence I suggest a Kantian approach depictable as follows, where the arrows represent support relations:

Standard internalist principles are

a priori   normative truths   because

↑   ↑

constitutivism is true   ↔   Kantian ethics of belief is true

This approach accepts internalist intuitions (e.g., brains-in-vats have justification) and principles (e.g., experience justifies belief), but sees them as requiring explanation. Kantian

---

[9] See Ahlstrom-Vij and Dunn (2018).
[10] See Sylvan (2020).

ethics of belief gives the first-order explanation. Kantian constitutivism gives the meta-normative explanation. These explanations are, in turn, mutually supportive.

I turn to develop this account.

3. Internalism as Constitutivism

The version of internalism I recommend is a form of *constitutivism*. Constitutivism seeks to ground the norms governing activities in the nature of those activities or the capacities they manifest. To understand what constitutive norms are, we can begin by considering chess. The rules governing chess pieces are constitutive of chess playing: you wouldn't be playing chess if you weren't trying to move the pieces correctly.

Now, activities like chess can be evaluated from an external perspective. You might think chess is a bad game: perhaps you think games should be more fun. But not all activities are like this, constitutivists stress. Some activities cannot be intelligibly questioned from the outside. Here we can consider activities that are constitutive of having a mind at all. Basic cognitive activities like *thinking* are examples. When Descartes claimed he was essentially a 'thinking thing', he had this idea in mind: he couldn't exist without the capacity for thought. If the capacity for thought is constitutive of having a mind, the norms of thought cannot be intelligibly questioned. To question, after all, is to produce reasons for doubt. Since doubting is a kind of thinking, giving reasons for doubt presupposes the authority of the norms of thinking. Hence rules of thought differ from rules of chess: you don't have to use the rules of chess to tell whether it is a good idea to play chess![11]

Constitutive norms of an activity have an interesting feature: anyone engaged in the activity automatically binds herself by these norms. Since these norms define the activity, to engage in it is to bind yourself by them. The norms of thinking are self-imposed in the deepest possible sense: they are constitutive of your mind and hence of your deepest self. Kantians thus deem these norms 'autonomous'—i.e., 'self-legislated', from Greek *autos* = self and *nomos* = law.

The first part of Kantian internalism draws upon these ideas. It treats epistemic norms as having authority in virtue of being constitutive of cognition as such, or of more determinate cognitive activities:

> **Kantian Epistemic Constitutivism (KEC)**: All epistemic norms that have authority for S have it in virtue of their compatibility with the constitutive principles of S's cognitive capacities and activities.

I will develop a more specific view below for illustration. The basic idea, though, is to find norms that are internal to cognition by examining the nature of cognitive activities. One place to look is to consider basic and generic cognitive activities like *recognizing objects* and *thinking*. One can also consider more specific activities like *recognizing pentagons* or *doing deductive reasoning*. The diversity of specific cognitive activities can generate normative diversity. For there might be cognitive activities that count as basic for you but not for others. So exactly which norms apply to you depends on the way you realize the general template for being a cognizer. Nevertheless, some norms apply to all cognizers—viz., the constitutive norms of generic and basic cognitive activities like thinking.

KEC is internalist partly because it honors the idea that legitimate authority must respect your capacity for self-governance—in this case, intellectual self-governance. Call

---

[11] Silverstein (2015) uses this point to address Enoch (2006)'s objection to constitutivism.

this the *Autonomy Precept* (AP). While recent forms of internalism don't invoke this percept, some older forms of internalism did. Hence Goldman (1980: 32) claimed that

> [o]n the [internalist] perspective […] a [doxastic principle] must not be 'heteronomous,' or dictated 'from without'. It must be 'autonomous,' a law we can give to ourselves….

Chisholm (1982: 63), for example, argues from AP to accessibilism. Unlike Chisholm, though, I don't think AP favors accessibilism. Constitutive principles can be hard to discover.

KEC further diverges from typical internalist views because it can allow *spatially* external factors—e.g., evidence in the world, like a bloody knife—to have *derivative* authority. KEC only says that the authority of epistemic norms over you is grounded in your mental constitution. Something spatially external could *acquire* such authority by being related in the right way to your mind. Again, this point is standard in meta-ethics after Schroeder (2007): meta-ethical internalists agree that reasons can include facts about the world. They just say such facts *have authority for S* because of S's constitution.

Justification and Constitutive Normativity

KEC could be developed in different ways. I will outline a specific Kantian framework for concreteness. One needn't accept all my baggage. But having a specific theory facilitates an existence proof that some internalist views escape Littlejohn's challenges.

To understand the framework, let's clarify the target. We can distinguish weaker and stronger concepts of justification. The weakest is a permissive concept: S has *minimal justification* to φ when S is permitted to φ. φ-ing might be permissible but only barely, by meeting minimum standards. There are more demanding concepts—φ-ing sufficiently above minimum standards ('adequate justification'), or φ-ing optimally ('full justification').

To give constitutivist analyses of these properties, I borrow ideas from Kant's ethics. Kant distinguished *perfect duties* like the duty not to use others as mere means, with which one *must* comply to meet minimum moral standards, from *imperfect duties* like the duty to help others, which set *ideals* that leave some playroom. These duties are illuminated by two tests. One is the *contradiction in conception* test: a maxim is flawed if it cannot be universally followed. There are perfect duties against acting on such maxims. Another is the *contradiction in will* test: a maxim is flawed if one cannot coherently will that it be universally followed. There are *imperfect duties* against acting on such maxims. Related tests distinguish two kinds of constitutive norms on φ-ing. Say that a norm N of φ-ing is

> *strongly constitutive* iff it is impossible to engage in φ-ing without being *disposed* to heed N

and

> *weakly constitutive* iff it is impossible to engage in φ-ing without *commitment* to N

Chess rules provide examples: rules about legal moves are strongly constitutive, whereas the ideal of checkmating is weakly constitutive. We can also distinguish weakly and strongly constitutive norms of a capacity C:

> *strongly constitutive norms* are ones one must be *disposed* to heed to have C,

and

> *weakly constitutive norms* are ones one must be *committed* to heeding to have C.

Linguistic rules are examples: L's basic syntactic rules are strongly constitutive of the capacity to understand L, whereas norms of usage for L's vocabulary are weakly constitutive.

These concepts facilitate generic constitutivist analyses of the three justification properties:

> (**Minimal**) S is minimally justified in φ-ing iff S's φ-ing is permitted by strongly constitutive norms of S's cognitive capacities and φ-ing.

> (**Adequate**) S is adequately justified in φ-ing iff S's φ-ing is minimally justified and S's φ-ing also sufficiently reflects commitment to weakly constitutive norms that non-conformity is significantly explained by performance error, luck, or interference.

> (**Full**) S is fully justified in φ-ing iff S's φ-ing is minimally justified and S's φ-ing also fully reflects commitment to weakly constitutive norms, so that non-conformity is wholly explained by performance error, luck, or interference.

To understand these claims, consider an example:

> EXAMPLE: S validly deduces a working hypothesis from premises that make better sense of S's evidence than the alternatives S considered. S reached these premises via induction. S's evidence is, however, so far slim.

EXAMPLE is mixed. S obeys perfect duties, like *Deduce validly!*, and partially obeys imperfect duties, like *Consider alternative hypotheses!* and *Seek good explanations!* Plausibly, S isn't fully justified. But I think S is a borderline case of adequate justification. It can be fine to accept working hypotheses on slim grounds if one plans to inquire further. So it is rash to deny adequate justification: we need more information.

What do we need to know? Here epistemology meets philosophy of psychology. I favor a principle of charity as an *a priori* constraint on interpretation.[12] As I understand charity, it doesn't preclude attributing non-conformity to norms to people: it just requires minimizing attributions of non-conformity *to people's mental capacities* rather than physical realization, circumstances, or interference. Hence, when people's beliefs are false, the errors are mostly not ascribable to *them*, understood as persons. In typical bad cases, error reflects one's physical realization or circumstances.

These claims may sound externalist, but they fit with internalism if charity is an *a priori* principle. Indeed, these claims make good internalist sense. For the same reason why matches between cognition and world are never guaranteed by internal states, *mismatches* are rarely explained by internal faculties. If so, error is no part of one's *record*. As Kant put it:

> All errors are, so to speak, crooked lines, which we determine while being driven on the one side by the understanding, from the other side by sensibility. (*Blomberg Logic* 87)

---

[12] See Lewis (1974) and Williams (2020).

Rationality isn't what leads us astray, in other words. This isn't to say we're angels: mismatches are common and can reflect human nature ('crooked timber').

So, I suggest that if φ-ing complies with strongly constitutive principles and manifests commitment to weakly constitutive principles, it is adequately justified.

The Constitution of Cognition

KEC needs an internalist-friendly metaphysics of cognitive capacities. Capacities are usually conjured by externalists like Schellenberg (2018), who understand them *teleologically*. But appealing to capacities alone doesn't make one externalist: only appealing to their extrinsic features (e.g., biological function) does. Internalists can welcome capacities and understand them differently. On one recent reading,[13] Kant has a capacity-based epistemology that invokes essential properties of cognitive capacities to ground epistemic normativity.

One Kantian idea is to understand capacities in terms of their '*forms*'—hence, Kant analyzes the Understanding by invoking the *Categories*. Another Kantian idea is to understand capacities in terms of *principles*: 'Experience therefore has principles of its form which ground it *a priori*, namely general rules of unity in the synthesis of appearances' (A156/B195). Cognitive science rooted in Chomsky relies on related ideas: capacities are grounded in *tacit knowledge of rules*, which is in turn grounded in mental architecture.[14]

What distinguishes Kantian epistemology from other views that summon tacit knowledge is that it treats foundational tacit knowledge as belonging to the mind's 'architecture' in Kant's sense—its '*a priori* form'. This knowledge is discoverable by philosophical reflection on the possibility of experience—call this *transcendental psychology* (Kitcher (1994)). Cognitive science explains how the transcendental story is *realized*. But the most general principles belong to transcendental psychology. This picture is surprisingly defensible: recent views about innate cognitive capacities strongly echo Kant.[15]

Kant focused on three capacities: Sensibility, Understanding, and Reason. I assume a related distinction between one's

> *presentational capacities*, which present one with apparent objects and features on the basis of sensory states (e.g., perceptions of pentagonal shapes),
>
> *recognitional capacities*, which enable one to recognize, of some presented object or feature, that it is that object or feature (e.g., recognizing something *as* a pentagon),
>
> *capacity of theoretical reason*, which enables one to reason from presentational and recognitional states to views about the world.

On KEC, what matters are constitutive principles of these capacities and their manifestations. To identify constitutive principles of a capacity C of S, we can ask:

> What principles of S's constitution make C possible as such?

To answer this question, we can apply two tests:

> **Strong Constitutiveness**: If S cannot conceivably have C unless S is *disposed to conform* to principle R, R is strongly constitutive.
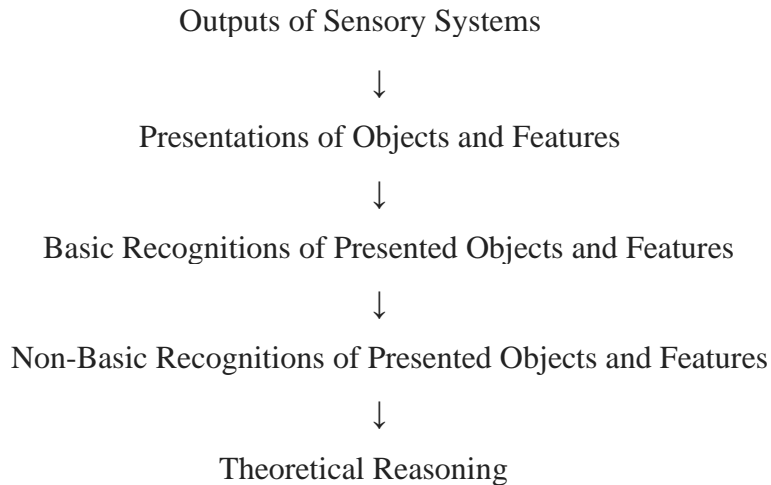
---

[13] See Schafer (2019).
[14] See Fodor (1968) and Chomsky (1986).
[15] See especially Spelke (2022).

Weak Constitutiveness: If S cannot conceivably have C unless S is *committed* to R, R is weakly constitutive.

To have a complete constitutivist view, we must understand the structure of cognition. I'll assume cognition properly so-called begins with the outputs of sensory systems and proceeds according to the following flowchart:

Outputs of Sensory Systems

↓

Presentations of Objects and Features

↓

Basic Recognitions of Presented Objects and Features

↓

Non-Basic Recognitions of Presented Objects and Features

↓

Theoretical Reasoning

We want to understand the constitutive principles of the capacities underlying the arrows. Kant and 20th century followers like Strawson (1959) had ideas. There are analogues in cognitive science that I'll mention.[16]

Presentation. How are presentations of whole, persisting objects possible? For Kant, they are possible thanks to the imagination and principles of the understanding.[17] In particular, one synthesizes the outputs of sensory systems, which represent parts of objects (e.g., the facing side of an apple), with imagined content (e.g., non-visible parts of the apple), where this synthesis is guided by innate understanding of objects. Only syntheses obeying these rules yield *objective* presentations rather than flawed ones, like hallucinations in psychosis.

       An up-to-date story is possible. Consider how infants can continue to track objects that have passed behind a screen. How is this possible? Spelke (1988) appeals to 'core knowledge' of objects—innate principles resembling Kant's First and Second Analogies. These principles allow us to 'connect the dots' and thereby apprehend whole, persisting ordinary objects. Such principles are strongly constitutive of our presentational capacities. Without being disposed to construct presentations of substantial, spatiotemporally continuous objects from sensory input, we couldn't be presented with ordinary things like apples. There are, however, different specific ways to satisfy these principles, which may constitute particular cognizers' specific capacities. Our capacities may also change with development.

       What can these norms explain? I think Kant is right that they explain how we can have evidence for an external world: it is by having experiences that satisfy these constitutive norms that we are able to gain evidence about the ordinary world.

Basic Recognition. How is recognition of objects and features *as* objects and features possible? For Kant, Categories like *substance* and *causality* do the work. There are modern

---

[16] See Spelke (1988, 1992) on presentational capacities, Carey (2009) on recognitional capacities, and 'mental logicians' (Rips (1994)) and theory theorists (Gopnik and Meltzoff (1997)) on rational capacities.
[17] See Sellars (1978) for lovely discussion.

versions of this idea. Spelke's project explains how we segment the perceived world into units subsumable under categorial concepts. Another project is needed to explain how we subsume these units under categorial concepts. Here an idea from Carey (2009) helps: the *concept* of an object makes recognition possible. Hence strongly constitutive norms of basic recognitional capacities might include being disposed to *conceive* of objects as such. More determinate ways of conceptualizing objects may constitute specific recognitional capacities. These may change with development.[18]

What can these norms explain? Again, I think Kant had the right idea: they ground an *a priori* entitlement to belief in the persistent objects encountered in everyday experience, by making basic recognitional capacities possible.

Non-Basic Recognition. How is recognition of secondary qualities (e.g., colors) or garden-variety natural kinds (e.g., trees) across fluctuating sensory input possible? Constitutive principles of non-basic recognition might include dispositions associated with *psychological essentialism* (Gelman (2003))—dispositions to treat observable patterns as explained by hidden properties. Kant (A652/B680) similarly invokes reason's need for *systematic unity* to explain acquisition of special science concepts and a taxonomic understanding of nature.

Such principles exemplify broader abductive principles. One plausibly couldn't glom onto worldly patterns, and hence acquire high-level concepts, without abduction. Still, there is room for variation. The abductive principle is only weakly constitutive. Strongly constitutive norms may include requirements on consistent taxonomizing and concept application, in line with whatever rules are constitutive of particular concepts. These make recognition of the ordinary world possible.

Theoretical Reasoning. Some principles of recognition are implicit principles of theoretical reason.[19] These principles can be made explicit and used in theoretical reasoning. We can usefully split theoretical reasoning into four stages:

> **Wondering**: Wondering what explains some phenomenon *E*.
>
> **Pondering**: Surveying explanations of *E* and getting input systems to deliver more presentations.
>
> **Evaluating**: Evaluating explanations and forming provisional judgments.
>
> **Concluding**: Drawing *conclusions*, which can be *close-minded* (e.g., hard agnosticism) or *open-minded* (e.g., suspension compatible with revisiting the question).[20]

Each step introduces epistemic duties. Pondering should obey imperfect duties on gathering explanations and evidence. Evaluating should obey the imperfect duty to favor good explanations and perfect duties to evaluate consistently. Concluding should obey perfect duties constitutive of the relevant mental action (e.g., deductions should be valid). These duties are constitutive principles: apparent thinking that shows no commitment to them is *pseudo*-thinking.

---

[18] See Carey (1985).

[19] For accounts of perceptual seemings as outputs of implicit reasoning, see Gregory (1980) and Rock (1983).

[20] See Lord and Sylvan (forthcoming) for discussion.

Summarizing and Comparing

The supreme principle of KEC is this:

> *Meta-Norm*: For any cognitive activity φ, S should φ in such a way that the maxim of S's φ-ing aligns with strongly and weakly constitutive principles of cognition.

With a more specific version of KEC in view, we can see how KEC improves upon standard internalist views and captures benefits of externalism without valorizing extrinsic properties of cognizers.

Note that, on KEC, whether your φ-ing is justified is determined by whether its structure aligns with constitutive principles. All cognitions except direct outputs of sensory systems must be *well-structured*. Since even seemings (e.g., the impression that someone is angry) reflect cognitive processing, they don't automatically justify. Hence KEC doesn't imply that you can be justified in believing that p just because p *seems true* to you: it might seem true to you only because of flawed processing.

Since perceptual presentations also reflect cognitive processing, there is a deeper epistemic story about empirical justification than naïve realist views suggest. But because the relevant principles are *strongly* constitutive, KEC doesn't demand more *work*. Hence KEC isn't as demanding as traditional versions of internalism that require hard labor to get behind the veil of perception.

KEC diverges from stereotypical internalist views in further ways.[21] Phenomenology and accessibility don't matter on KEC. What matters is the structure of your cognitive activity. Perhaps you necessarily have implicit *practical understanding* of it, like your understanding of how to move your arms. It is, after all, something you *do*. But this isn't knowledge via inner observation, and it needn't be *articulable*.

KEC also shares virtues of externalist views while avoiding their vices. Like process reliabilism, KEC treats justification as determined by the quality of one's processing. But unlike reliabilism, KEC denies that this quality is a contingent feature. Like virtue epistemology, KEC sees justification as determined by the quality of one's cognitive efforts. But it denies that this quality is determined by whether those efforts reliably hit upon truth. KEC is, however, compatible with a different kind of truth-connection: a cognitive activity that respects weakly and strongly constitutive principles cannot be *attributably* inaccurate.

So, KEC promises a plausible account of justification. While a complete evaluation requires a complete metaphysics of cognition, KEC avoids standard concerns. KEC avoids skepticism, unlike traditional foundationalism. It doesn't overextend justification, unlike views that invoke seemings as justifiers. It doesn't overintellectualize justification: while it has more structure than naïve realist views, it doesn't require more work. Furthermore, KEC isn't a 'time-slice' view and sidesteps worries about such views.[22] Although only the internal structure of cognitive activities matters on KEC, cognitive activities are extended in time.

4.   Ratifying the Constitution: Epistemic Non-Consequentialism

KEC may be plausible. But to show that internalist epistemology is possible, more is needed. One must explain why internalist intuitions track *justification* rather than weaker statuses. To address this challenge, I will respond by invoking an independently defensible ethics of belief and showing that KEC aligns with it.

---

[21] Wedgwood (2017) and Pollock and Cruz (1999) diverge in some similar ways.

[22] For full argument that this view avoids problems for time-slice views, see Sylvan (2017: 149-150).

Elsewhere I defend a non-consequentialist ethics of belief: justified beliefs are beliefs that manifest *respect for truth*.[23] This view is compatible with different analyses of justification, but its implications vary given different analyses. It would be good if there were an analysis that clarified why justified beliefs manifest respect for truth. Happily, KEC fits the bill. KEC combined with epistemic non-consequentialism predicts that whether a cognitive activity manifests respect for truth is determined by its constitution. This prediction is plausible: respect plausibly depends on intrinsic features of one's mental activity. If one thinker hits upon truth and another doesn't, both may equally manifest respect for truth if the latter was merely unlucky.

There are further reasons to accept this prediction. Respect requires the right kind of motive. Activities can be individuated finely to include motives (consider *murder*). Hence one might expect the (dis)respectfulness of cognitive activities to reflect their maxims. This seems right: whether an inference was fallacious depends on what kind of inference it was. Purported deductions from non-truth-preserving rules are fallacious, but inductions from non-truth-preserving rules aren't. So, the link between respect and motivation recommends KEC.

A further argument rests on the observation that whether you respect a value V depends on how you respond to relevant appearances, including misleading ones.[24] If all available information suggests X-ing is harmless to V and you take this into account, how could you manifest disrespect for V?

Here one might worry about negligence and biased perception. What if things seem some way for bad reasons, or because you have objectionably limited evidence? But such worries are rash. The problem with cases of objectionably limited evidence reflects the maxim of one's overall information-processing. If you have objectionably limited evidence, then plausibly you earlier manifested disrespect, by failing to expand your evidence. Persisting with maxims based on earlier disrespect transmits disrespect to the current time-slice, since the whole process manifests the implicit maxim: *I will draw conclusions after this [...too quick...!] survey of evidence.*

Biased perception may seem more concerning: one might doubt we can reiterate the foregoing story. But Siegel (2017) helps. Perceptual appearances derive from processing. The maxim of some processing may conflict with the processor's commitment to veridicality. If we can fault your attitudes for being based on appearances generated by bad maxims, we can attribute those maxims to your processing: you are proceeding *unconstitutionally*.

I suggest, then, that respect for truth could be internally determined. If so, and there are independent arguments for epistemic non-consequentialism, justification and not mere excuse can be internally determined. The constitutive norms of cognition matter because one must heed them if one is to heed the fundamental norm of respect for truth. Hence epistemic non-consequentialism supports KEC. The support is mutual, since KEC also supports epistemic non-consequentialism's predictions. The two cohere like Kantian ethics and philosophy of practical reason cohere.

5.   Ethics and Epistemology: Answering Littlejohn

The symmetry of Kantian epistemology and practical philosophy suggests a response to Littlejohn. Littlejohn (this volume) argues that the plausible view that rights and duties are external conflicts with global internalism. He suggests that since global internalism is false and epistemic and practical normativity are unified, epistemic internalism is false. I think this argument overlooks global Kantianism.

---

[23] See Sylvan (2020).
[24] See Sylvan (2021).

According to global Kantianism, all normativity derives from fundamental values to be respected, not promoted: accuracy (in epistemology) and personhood (in ethics). These values are not *goods*, but rather properties that figure in fundamental constitutive norms. Although these values are not denizens of individual minds, they have authority because they are objects of constitutive commitments of agency and cognition. This authority is hence internalist. Because a person's manifesting respect depends on the quality of her deliberation, the justification generated is also internalist.

Global Kantianism undermines Littlejohn (2012)'s argument for externalism. Littlejohn argued that justification must be objectively truth-connected, and that only externalism permits this. But epistemic Kantianism forges an objective connection between justification and fundamental epistemic value. It just understands this connection *non-instrumentally*. That is what one might reasonably expect by analogy with ethics.

Cognizers are constitutively committed to a truth norm. But one doesn't *violate* this norm merely by failing to conform to it. Violation requires that one's non-conformity is attributable to one's cognitive agency. Non-conformity cannot be so attributable if one cannot *comply* given one's epistemic position. Duties of respect can even require non-conformity: one may have a respect-based obligation to act on misleading appearances, if failing to respond would constitute disrespect. Hence acts and beliefs can be motivated by misleading appearances but remain justified because they obey a duty of respect.

One might wonder whether global Kantianism respects the idea that your duties and others' rights are not in your head. But it is compatible with that view. Here we should distinguish fundamental and derived rights, and consider how such rights and their correlated duties derive from fundamental values.

Fundamental rights on the Kantian view are rights against disrespectful intentional agency. These rights are correlates of a fundamental negative duty not to treat persons as mere means. Since problematic instrumental orientation can be manifested even in skeptical scenarios, this duty is internalist, though the value that grounds it is not in your head.[25]

Fundamental negative rights are rights against intentional agency, not mere causation. Falling rocks don't violate your rights. Sleepwalkers who step on your toes don't violate your rights. To be sure, the dignity of personhood doesn't only impose negative duties. Persons are ends-in-themselves. This status confers positive rights, like the right to conditions of dignity. But what positive rights imply for individuals is complicated. A patient's derived rights and an agent's derived duties depend on the agent's abilities and relationship to the patient. Not everyone has the same derived duties, since agents' abilities and position constrain what is required of them. This is a good feature of Kantianism.

Now, one may sometimes have duties to pay reparations when one merely causes bad states of affairs. If I sleepwalk in your house and break something, it could make sense to make me replace it. But this case involves no rights-violations. Perhaps you have a derived right to aid, but it may or may not be a sound convention to expect me to provide aid. More importantly, there are no general epistemic analogues of these phenomena. Facts lack rights. It is *incorrect* to believe falsehoods. But Kantians can agree, since the truth norm is a constitutive commitment of cognition.

There are more analogies than disanalogies between ethics and epistemology, since their fundamental norms are norms of respect. There are moral duties of respect that resemble the fundamental epistemic duty. But when objects of respect differ, we should expect some disanalogies. Global Kantianism explains the analogies and the disanalogies.

---

[25] For full defense, see Sylvan (2021).

References

Ahlmstrom-Vij, K. and Dunn, J. (eds.) 2018. *Epistemic Consequentialism*. OUP.
Bonjour, L. 1998. *In Defense of Pure Reason*. CUP.
Callahan, L. 2021. 'Epistemic Existentialism.' *Episteme* 18: 539-554.
Carey, S. 2009. *The Origin of Concepts*. OUP.
Carnap, R. 1962. *Logical Foundations of Probability*. U. Chicago Press.
Chisholm, R. 1957. *Perceiving*. Cornell UP.
Chisholm, R. 1982. *Foundations of Knowing*. U. Minnesota Press.
Chomsky, N. 1966. *Cartesian Linguistics*. Harper & Row.
Chomsky, N. 1986. *Knowledge of Language*. Praeger.
Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. CUP.
Conee, E. and Feldman, R. 2002. *Evidentialism*. OUP.
Conee, E. and Feldman, R. 2008. 'Evidence' in Smith, Q. (ed.) *Epistemology: New Essays*.
        OUP.
Enoch, D. 2006. 'Agency, Shmagency.' *Phil. Review* 115: 169-198.
Feldman, R. and Conee, E. 2001. 'Internalism Defended.' *American Phil. Quarterly* 38: 1-18.
Fodor, J. 1968. 'The Appeal to Tacit Knowledge in Psychological Explanation.' *J. Phil.* 65:
        627-40.
Foley, R. 1987. *The Theory of Epistemic Rationality*. Harvard UP.
Fumerton, R. 1995. *Metaepistemology and Skepticism*. Rowman & Littlefield.
Gelman, S. 2003. *The Essential Child*. OUP.
Goldman, A. 1980. 'The Internalist Conception of Justification.' *Midwest Studies in Philosophy*
        5: 27-51.
Gopnik, A. and Meltzoff, A. 1997. *Words, Thoughts, and Theories*. MIT Press.
Gregory, R. 1980. 'Perceptions as Hypotheses.' *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 290:
        181-97.
Kant, I. 1785/2012. *Groundwork for the Metaphysics of Morals*. CUP.
Kant, I. 1787/1998. *Critique of Pure Reason*. CUP.
Kant, I. 1992. *Lectures on Logic*. CUP.
Kitcher, P. 1994. *Kant's Transcendental Psychology*. OUP.
Korsgaard, C. 2009. *Self-Constitution*. OUP.
Lewis, D. 1974. 'Radical Interpretation.' *Synthese* 23: 331-344.
Littlejohn, C. 2012. *Justification and the Truth Connection*. CUP.
Lord, E. and Sylvan, K. Forthcoming. 'Beginning in Wonder' in Hughes, N. (ed.) *Epistemic
        Dilemmas*. OUP.
Markovits, J. 2015. *Moral Reason*. OUP.
McCain, K. 2014. *Evidentialism and Epistemic Justification*. Routledge.
Parfit, D. 2011. *On What Matters* (V.1). OUP.
Pollock, J. and Cruz, J. 1999. *Contemporary Theories of Knowledge*. Rowman & Littlefield.
Rips, L. 1994. *The Psychology of Proof*. MIT Press.
Rock, I. 1983. *The Logic of Perception*. MIT Press.
Schafer, K. 2019. 'Kant: Constitutivism as Capacities-First Philosophy.' *Phil. Explorations* 22:
        177-193.
Schellenberg, S. 2018. *The Unity of Perception*. OUP.
Schroeder, M. 2007. *Slaves of the Passions*. OUP.
Sellars, W. 1978. 'The Role of Imagination in Kant's Theory of Perception' in Johnstone, H.
        (ed.) *Categories: A Colloquium*. Penn State Press.
Siegel, S. 2017. *The Rationality of Perception*. OUP.
Silins, N. 2011. 'Seeing through the 'Veil of Perception'.' *Mind* 120: 329-367.
Silverstein, M. 2015. 'The Shmagency Question.' *Philosophical Studies* 172: 1127-1142.
Smith, M. 1995. 'Internal Reasons.' *PPR* 55: 109-131.

Spelke, E. et al. 1992. 'Origins of Knowledge.' *Psychological Review* 99: 605-632.

Spelke, E. 1988. 'Where Perceiving Ends and Thinking Begins' in Yonas, A. (ed.) *Perceptual Development in Infancy*. Laurence Erlenbaum Associates.

Spelke, E. 2022. *What Babies Know*. OUP.

Strawson, P. F. 1952. *An Introduction to Logical Theory*. Routledge.

Strawson, P. F. 1959. *Individuals*. Routledge.

Sylvan, K. 2017. 'Responsibilism out of Character' in Alfano, M. and Fairweather, A. (eds.) *Epistemic Situationism*. OUP.

Sylvan, K. 2020. 'An Epistemic Nonconsequentialism.' *Phil. Review* 129: 1-51.

Sylvan, K. 2021. 'Respect and the Reality of Apparent Reasons.' *Phil. Studies* 178: 3129-3156.

Wedgwood, R. 2002. 'Internalism Explained.' *PPR* 65: 349-369.

Wedgwood, R. 2017. *The Value of Rationality*. OUP.

Williams, R. 2020. *The Metaphysics of Representation.* OUP.