

Knowledge-first functionalism

Mona Simion

<http://eprints.gla.ac.uk/196576/9/196576.pdf>

DEPTH	EXPERTISE	VOICE	GENERATED
Overview (10-15 min)	Balanced	Adam	August 28, 2026

What does it mean for a belief to be justified? Not morally justified, not socially acceptable — epistemically justified, in the sense that gives you the right to call something knowledge. Let that question sit for a moment. For most of the history of epistemology, the answer has started with evidence, reason, or probability. Mona Simion's paper argues that all of those approaches get the order of explanation exactly backwards. Knowledge comes first. Justification is downstream of it.

To understand why that inversion matters, you need to understand what Simion thinks justification has to do. She opens with a general Deontic Thesis — justification, in any domain, tracks permissibility. Applied to belief, this becomes the Deontic Thesis for Belief, or DTB: a belief is epistemically justified if and only if it is epistemically permissible to hold it. That might sound like a tautology, but it has real teeth. If justification is about permissibility, then any theory of justification is also, necessarily, a piece of normative theory. And normative theories have to meet normative standards.

Simion identifies three minimal desiderata. First, normative prior plausibility: your account of justification should fit coherently within mainstream normative theory or give a good reason why epistemic normativity is a special case. Second, normative generalizability: the account should extend beyond the narrow epistemic setting where it was developed, applying across practical, moral, and social domains. Third, normative evidential adequacy: the account should actually capture the distinctions our intuitions make between different cases. These aren't exotic philosophical demands. Simion calls them fairly minimal, and thus fairly uncontroversially constraining. The challenge is that existing knowledge-first views struggle to meet even this minimal bar.

The strongest knowledge-first claim is also the simplest: justified belief just is knowledge. Call it JB equals K. If you have knowledge, your belief is justified. If you don't, it isn't. Prominent defenders include Brian Ball, Clayton Littlejohn, Adrian

Haddock, Jonathan Sutton, and Tim Williamson. The appeal is real — it's clean, it's unified, and it makes epistemic normativity match epistemic success. But the problems arrive fast.

Knowledge is factive. You can only know something that's true. Which means JB equals K has no room for justified false belief — and intuitively, we think justified false belief is possible all the time. A victim of sophisticated deception forms a belief based on everything available to them. A Gettier case — where someone has a justified true belief that still isn't knowledge because of some lucky coincidence — seems to involve a genuine epistemic achievement, not a failure. JB equals K would have to say both people are simply unjustified. That's a hard bullet to bite.

The standard response from JB equals K proponents is to invoke normative pluralism: our sympathy for the deceived victim isn't tracking justification, it's tracking blamelessness. The victim complied with derivative norms — they lacked control, and they were ignorant of the breach — even though they failed the primary knowledge norm. Williamson's version treats blamelessness-conferring conditions as themselves derivative from the knowledge norm.

But Simion presses on this. The blamelessness strategy is too coarse-grained. Consider Goldman's benighted cognizer, a character Simion calls Ben — someone who forms beliefs by astrology because of cultural determination. Ben seems blameless in roughly the same way a Gettier victim is blameless. Yet intuitively, we'd advise Ben to change his epistemic practices in ways we simply wouldn't advise the Gettier victim. Something normatively important separates them. The blamelessness account treats both as mere norm violators and can't explain that difference. That's a failure on normative evidential adequacy.

There's also a generalizability problem. In moral and practical domains, a justified action can fail to achieve its end and still be justified — the doctor who prescribes the right treatment and loses the patient anyway hasn't acted unjustifiably. JB equals K collapses that gap in the epistemic case, making epistemic justification uniquely identical to success. Without a positive argument for why epistemic normativity should work differently from every other normative domain, that's a suspicious move.

Knowledge-first theorists have developed milder variants to relieve the pressure. Alexander Bird distinguishes praiseworthy belief from blamelessness and from knowledge itself. Jonathan Ichikawa ties justified belief to intrinsic mental states consistent with having knowledge. Steven Reynolds construes justification as the appearance of knowledge. These views allow for justified false belief and separate

justification from success. But each one inherits problems. Reynolds's appeal to assertion raises issues because assertion is socially embedded and may require sophisticated reflective capacities. Ichikawa's mapping to blamelessness runs into the same coarse-grained difficulties. Bird's tripartite taxonomy preserves distinctions without explaining why the praiseworthiness standard, rather than some other internal standard, should carry the justificatory weight. The tension remains.

This is where Simion's own proposal enters. She calls it Knowledge First Functionalism, or KFF, and states it precisely: a belief is epistemically justified if and only if it is generated by a properly functioning cognitive process that has the etiological function of generating knowledge. Two concepts do the work — proper functioning and etiological function — and both need unpacking.

An etiological function is a historical fact about why a trait exists. Simion formulates the E-Function condition: a token of type T has the etiological function of producing effect E in system S if, first, tokens of T produced E in the past; second, producing E benefited S or S's ancestors; and third, that benefit contributes to explaining why T still exists in S. This is a consequence etiology — the trait persists because its effects fed back into the system's survival or continuation. The heart's function is biological because it pumped blood, and that pumping contributed to survival. An art object's function is aesthetic because the relevant benefit is aesthetic. What individuates function is the type of benefit.

Now apply that to cognition. Simion argues that generating knowledge satisfies the E-Function condition for our cognitive processes: they have produced knowledge in the past, producing knowledge benefited our ancestors — she appeals to the widely accepted thesis that knowledge has distinctive value — and that benefit helps explain why our cognitive systems persist. The main epistemic function of our belief-forming faculties, on this account, is to generate knowledge.

The crucial distinction is between functioning properly and fulfilling the function. Simion uses the heart in a vat: a heart plugged into a nutrient tank and pumping orange juice is still functioning normally in the etiological sense — it's doing what its history equipped it to do — but it is failing to fulfill its function of pumping blood in an organism. Proper functioning is a process-level norm. Function fulfillment is about outcomes in the current environment. And these can come apart.

That gap is exactly what KFF needs. A Gettierized agent, a deceived victim, a conscientious false believer — all of them are running cognitive processes that are functioning properly. Their belief-forming systems are doing what those systems were historically constituted to do. The fact that the environment conspired against

knowledge on this particular occasion doesn't change the process-level verdict. As Simion puts it, Gettierized and deceived victims employ properly functioning cognitive processes that have the function of generating knowledge, and thus come out justified. Ben, the astrology-driven believer, does not — his processes don't have the right etiological history to count as aimed at knowledge in the relevant sense. KFF draws exactly the distinction the blamelessness strategy couldn't.

Now run KFF against the three normative desiderata. On generalizability: KFF allows justified belief to come apart from successful belief, just as justified action comes apart from successful action in moral and practical domains. That's the structural symmetry JB equals K lacked. On evidential adequacy: KFF captures the intuitive gradient between Gettier victims, deceived subjects, and Ben — grounding the distinction in process history rather than in the blunt instrument of blamelessness. On prior plausibility: the etiological story provides naturalistic grounding. Justification is anchored in evolutionary and developmental history, not in purely internalist intuition. The account earns its normative credentials from outside epistemology as well as within it.

Simion also notes a resonance with Williamson's broader picture: her knowledge-centered etiological story receives further support if one accepts Williamson's claim that knowledge is a state of mind, since knowledge construed as a state makes intuitive sense as the target that cognitive processes are historically tuned to generate.

Here is what KFF actually changes. The traditional picture treats knowledge as something built up from justified beliefs — justification is the foundation, knowledge is what you get when everything goes right. KFF inverts that entirely. Knowledge is the target. Justification is what it looks like when your cognitive machinery is properly aimed at that target and working as it should. Epistemic failure, on this view, isn't a shortfall of evidence or reasoning — it's a failure of function. Either your processes aren't working properly, or the environment has intervened between proper functioning and its intended product. That's a different diagnosis, and it changes what questions epistemology should be asking.

This lecture was created by ennepō.

Go to <https://ennepo.ai> to Discover, Create and Follow the latest research in your field.

Read when you can. Listen when you want to.

CREDITS

AI-generated lecture by ennepō.ai, adapted from:

- “Knowledge-first functionalism”, Mona Simion (Philosophical Issues, 2019). doi:10.1111/phis.12152. CC-BY-SA (<https://creativecommons.org/share-your-work/cclicenses/>)

Restructured, summarized and explained. Not endorsed by the authors.

This lecture (script and audio) is licensed under CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0/>).

<https://ennepo.ai/lecture/c279231a-bb87-4916-bebb-faf24f5f0bbb>

Generated by ennepō