

ORIGINAL ARTICLE

A priori perceptual entitlement, knowledge-first

Mona Simion

COGITO Epistemology Research Centre,
University of Glasgow

Correspondence

Mona Simion, COGITO Epistemology
Research Centre, University of Glasgow.
Email: mona.simion@glasgow.ac.uk

Abstract

Tyler Burge notably offers a truth-first account of perceptual entitlement in terms of *a priori* necessary representational functions and norms: on his account, epistemic normativity turns on natural norms, which turn on representational functions. This paper has two aims: first, it criticises Tyler Burge's truth-first *a priori* derivation on functionalist and value-theoretic grounds. Second, it develops a novel, knowledge-first *a priori* derivation of perceptual entitlement. According to the view developed here, it is *a priori* that we are entitled to believe the deliverances of our perceptual belief formation system, in virtue of the latter's constitutive function of generating knowledge.

1 | INTRODUCTION

Tyler Burge¹ once famously argued that perceptual epistemic entitlement can be derived on *a priori* grounds (2013). The claim, scandalous and exceptionally thoroughly defended, got a lot of attention and made a very nice career for itself. At the same time, it's fair to say, not many philosophers rushed to embrace it. Burge himself gave up some of its initial ambition as of late (e.g. 2003, 2020).

This paper attempts a novel *a priori* derivation of perceptual epistemic entitlement. Like Burge's, it does so on functionalist grounds: it proceeds from the representational function of our perceptual capacities. *Contra* Burge, it argues that the function at stake is generating *knowledge*, and that this function gives rise to a constitutive epistemic norm of belief which, when met, delivers epistemic entitlement.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Philosophical Issues* published by Wiley Periodicals LLC

To this aim, I first (Section #2) unpack Burge's argument in its most recent form, and argue that it fails at two junctures: first, we miss support for Burge's claim that a norm of reliability in function fulfilment—which is supposed to ground epistemic entitlement—drops out of the function in question (Section #3). Second, we don't have support for the function of our belief formation capacities being to generate *mere* true beliefs (Section #4). Last (Section #5), I propose my own derivation that, in avoiding the second juncture by making a knowledge function claim, does not require the reliability claim downstream either. In the last section I conclude.

2 | BURGE'S DERIVATION

Tyler Burge (e.g. 2003, 2010, 2013, 2020) gives a truth-centric account of perceptual entitlement in terms of *a priori* necessary, representational functions and norms: on his account, epistemic normativity turns on natural norms, which turn on representational functions.

Burge takes it to be *a priori* that if a system is a representational system, it has a representational function. He also takes it that, insofar as a system is a perceptual system, it is a representational system. In turn, he argues, veridicality is the chief representational good. For propositional representations, such as beliefs, the measure of veridicality is truth. According to Burge, every false belief is a failure *qua* belief. Since, according to Burge, failure is evidence of function, it follows that the function of every token belief is to be true (2003, p. 509; 2010; 2020). Attributively (Geach, 1956), then, Burge takes it that good belief is true belief: that is, in order for a belief to be a good token of its type, it needs to be true. As such, he argues, it is knowable on *a priori* grounds that the representational function of systems of belief formation is to generate true beliefs.

According to Burge, wherever there are functions or, more generally, wherever there is purposiveness, there are standards for how to go about realizing the function or the end state of the purposiveness (2010, p. 339). This applies, he says, "to all biological organisms and their subsystems, to artefacts, to animal agency, to perception and belief, to inference, to knowledge" (2010, p. 339). The relevant standards are natural norms,² in that they drop right out of the function in question.

Importantly, according to Burge, and *contra* views notably defended by Peter Graham (e.g. 2012) and Ruth Millikan (e.g. 1984), these representational functions and norms do *not* reduce to biological etiological functions. This does not mean that belief-forming structures cannot have forming true beliefs as a biological etiological function too; it just follows that the latter is a contingent matter.

According to Burge, entitlement is the fulfilment of norms associated with the representational good, i.e. truth. As such, epistemic entitlement requires (a) that the belief was formed or sustained through a competence whose outputs have a very high ratio of true beliefs over false beliefs in normal, content-fixing³ conditions, and (b) that the belief was well-formed (not through malfunction). In short, warranted beliefs will be beliefs formed via a properly functioning and normally reliable representational competence.

Here is my reconstruction of Burge's argument:

1. Representational systems have representational functions.
2. The representational function of a system is to successfully represent.
3. The perceptual belief formation system is a representational system.
4. The representational function of the perceptual belief formation system is to successfully represent (from 1–3)

5. Successful representation implies veridicality.
6. The representational function of the perceptual belief formation system is to veridically represent (from 4 and 5).
7. Veridical propositional representations are true propositional representations.
8. Beliefs are propositional representations.
9. The representational function of the perceptual belief formation system is to form true beliefs (from 6–8).
10. If a system S has the function of type T to *phi*, then it (T)-should do so reliably.
11. The perceptual belief formation system constitutively (representationally)-should⁴ reliably form true beliefs (from 9 and 10).

Recall, also, that, according to Burge, entitlement is the fulfilment of norms associated with the representational good, i.e. truth. Epistemic entitlement, then, will be the fulfilment of (11).

The above is an attempted *a priori* derivation of Burge's view on perceptual entitlement from the nature of our perceptual systems. Let's call this the 'Modest Derivation' henceforth: it is modest in that, crucially, it only claims to derive a normative claim: perception *should* be reliable.

In contrast, earlier Burge stood by an even more ambitious claim: in former work, Burge used to think that it is *a priori* that our perceptual systems are reliable in normal, content-determining conditions, and, as a result, that it is *a priori* that we are entitled to believe the products of our perceptual systems (Graham, 2020). Here are some relevant passages from his early work from the 1990s (collected in 2013) to this effect:

I think that being reliable in normal circumstances can be shown to be necessary to a capacity [...] to register perceptual appearances [...] (2013: 294) Being a perceiver necessarily involves certain reliabilities in perceiving normal perceptual objects in normal circumstances (Burge, 2013, p. 347).

So, to put it differently, while late Burge merely believes the norm (11) follows on *a priori* grounds (Modest Derivation), early Burge also believes it follows on *a priori* grounds that (11) is, *as a matter of fact*, met in normal conditions (henceforth, Ambitious Derivation). The Ambitious Derivation was backed by the following generalized principle about functions:

12. A system S has a function F only if it reliably Fs in normal conditions.

This latter claim, in turn, together with (9), implies:

13. The perceptual belief formation system reliably forms true beliefs in normal conditions.

Unfortunately, (12) is false. Several systems in nature have a function that they only very rarely fulfil in normal conditions: perhaps *the* paradigmatic examples are sperm, and danger detection mechanisms in rabbits (Graham, 2019). Since (12) is false, Burge does not have an *a priori* derivation for (13), and we don't get an *a priori* entitlement to believe based on perception.

Burge abandoned the Ambitious Derivation in "Perceptual Entitlement" (2003); here is Burge: "The nature of perceptual states does not require that they be reliably veridical in their normal content-determining environment [...]—the circumstances in which their contents were formed." Importantly though, Burge still stands by the derivation of (11) on *a priori* grounds.

In a nutshell: even though it does not follow *a priori* that perception is reliable in normal conditions, it may well still follow that it *should* be, in which case Burge is left with an *a priori* derivation of his view of perceptual entitlement.

To put my cards right on the table, in what follows, I will take issue with two junctures in Burge's Modest Derivation. First, I will argue that (10) is problematic on functionalist grounds. Second, I will argue that (6) requires further defence.

3 | AGAINST THE A PRIORI RELIABILITY NORM

One important problem for Burge is how to account for the fact that many unreliable (in normal conditions) representational systems can be found in nature. Furthermore, these systems seem, intuitively, to be properly functioning in spite of their unreliability. Here is Peter Graham (2019):

The rabbit is a skittish creature. At the slightest sign of danger, it runs away. But it is nearly always running away from nothing. Are its beliefs (suppose it has them) reliable? In a sense, the rabbit's danger belief is highly reliable: if danger is present, it believes danger is present. Harty Field (1990) called this "world-to-head" reliability: if the world is a certain way, the odds are very high you believe it is that way. Given what the rabbit needs to produce—a danger belief in the presence of danger—the rabbit's beliefs are world-to-head reliable. But the rabbit's beliefs are not reliable in the "head-to-world" direction: if it believes the world is a certain way, it probably isn't. It has a dependable but not a ratio-reliable "design." If truth over error is your goal, don't ask a rabbit (Graham, 2019, p. 21).

In short, here is the problem: on Burge's account, systems that have generating truth as their function should do so reliably in normal conditions. The case of the rabbit seems to prove otherwise: indeed, plausibly enough, this is how the rabbit's representational systems are supposed to work to begin with: world-to-head reliability keeps the rabbit alive. Head-to-world reliability does not.

Burge discusses this problem in many places (e.g. 2003, pp. 517–518, 537). He acknowledges that reliable representation need not be biologically good representation, nor the other way around. Even so, he argues that the danger representation in rabbits, *should* be head-to-world reliable *qua representational kind*. The relevant should is the should of representation. That is because, according to Burge, given a representational function, the corresponding reliability norm applies: even if head-to-world reliability would be biologically bad for the rabbit, it would be representationally good. Being reliable is always a representational norm, given a representational function.

Now, recall that according to Burge, the representational reliability norm drops right out of the representational function. The thought is that if a trait has the function to generate truths, it should do so reliably. Recall, also, that Burge takes this principle to be perfectly general. Let's spell this out:

FUNCTIONALIST RELIABILITY CLAIM: For any function F of system S , system S should^F generate F reliably (where should^F is a should typed by the function in question, and it is, crucially, normative, not descriptive).

Why think, in line with Burge, that the FUNCTIONALIST RELIABILITY CLAIM holds? At first glance, one might believe it to be quite a simple and plausible idea: if the function of a system S is to produce apples rather than pears, it'd better be the case that S does so more often than not;

that is, that more often than not, S successfully produce apples rather than pears. Similarly, if S's function is to produce true beliefs rather than false beliefs, it'd better be the case that S does so more often than not.

Note, though, that this plausible thought is not enough for Burge's purposes: after all, not just any reliability rate is enough for epistemic purposes: epistemic normativity requires a very high level of reliability (50% +1 just won't do!). If so, the general principle that we actually need is much more stringent: it should be the case that, when something T has the function F in S, T Fs with a very high degree of reliability:

FUNCTIONALIST RELIABILITY CLAIM*: For any function F of system S, system S should^F generate F with a high degree of reliability.

Is FUNCTIONALIST RELIABILITY CLAIM* plausible? For some systems, it seems to hold: Think of coffeemakers; their function is to produce coffee. If they fail to do so very reliably, they're bad coffee makers: they don't meet the reliability norm for coffee makers.

Note, however, that what is enough, reliability-wise, for norm conformity seems to be domain specific, and to vary from high thresholds to, indeed, quite low requirements. The example of the coffee machine is one where high reliability seems required for norm conformity. In other domains, however, the degree of reliability in reaching the aims internal to the domain required for norm conformity seems quite low. Consider, for instance, baseball: good players rarely even reach 30% reliability, but that's fine. They are not thereby falling short in any way. Similarly, consider, again, sperm: the biological function at stake—reproduction—is extremely unreliably fulfilled, and that's *biologically* perfectly fine: indeed, too much reliability would likely be disastrous for the species. The FUNCTIONALIST RELIABILITY CLAIM* is false.

One way out for Burge at this point would be to make the claim less ambitious, by restricting the FUNCTIONALIST RELIABILITY CLAIM* in the following way:

FUNCTIONALIST RELIABILITY CLAIM:** For any function F of system S, system S should^F generate F with a D-appropriate degree of reliability (where D stands for the normative domain at stake).

At least on the face of it, the FUNCTIONALIST RELIABILITY CLAIM** seems correct. The problem with it, though, is that it fails to serve Burge's purposes in two ways. Recall that Burge wanted functions to deliver norms. It is plain to see, though, that the content of the norm outlined by the FUNCTIONALIST RELIABILITY CLAIM** is not fully delivered by the function at stake anymore, but rather it hinges on independent assumptions about the domain in question. More reliability in function fulfilment is required for some functions than for others, depending on how the relevant domain D is regulated. Nothing in the description of the function of coffee machines—making coffees—determines how reliable coffee machines should be: rather, further details pertaining to the domain of coffee making are what sets the threshold for tolerance for error: how hard or easy it is to produce a coffee is likely relevant, for instance. Similarly, think of baseball again: there's nothing about the function of scoring that delivers the relevant reliability threshold. Rather, other details about the domain step in: likely, e.g., how hard or easy it is to produce a base hit.

If that is the case, however, the representational function will not deliver the high reliability threshold for Burge; rather, further assumptions about the domain are needed in order to set the reliability norm in place. If this is so, however, we are back to square one: warrant is not (fully) grounded in our cognitive capacities' representational functions. Also, if the function does not fully determine the content of the norm, the claim that the norm in question is an *a priori* derivable *natural* norm remains undefended at two junctures: since it does not solely drop out of the function at stake, it is not clear the norm in question is a natural norm in Burge's sense.

Second, there is reason to think that domain specificity spoils a priority. Consider baseball again: had our species been better at hitting, plausibly, the content of the norm would be different, in that the domain-specific reliability threshold would be higher. Similarly, one could think that, had we been better at perceiving (better eyesight, more fine-grained sensory motor skills etc.), the domain-specific reliability threshold would be higher.

4 | AGAINST THE TRUTH FUNCTION

We have seen that Burge's premise (10), according to which the fact that a system has a particular function implies there is a norm that requires the system in question to fulfil it reliably fails on functionalist grounds. While its failure is enough to render the *a priori* derivation unsuccessful, in what follows, I would like to point to another, earlier place in the argument where the derivation breaks down: premise (6). The reason for this is that identifying this particular weakness will be instructive in developing my own derivation in the section to follow.

For the reader's convenience, here is my unpacking of Burge's derivation one more time:

1. Representational systems have representational functions.
2. The representational function of a system is to successfully represent.
3. Our perceptual belief formation system is a representational system.
4. The representational function of our perceptual belief formation system is to successfully represent (from 1–3)
5. Successful representation implies veridicality.
6. The representational function of our perceptual belief formation system is to veridically represent (from 4 and 5).
7. Veridical propositional representations are true propositional representations.
8. Beliefs are propositional representations.
9. The representational function of our belief formation system is to form true beliefs (from 6–8).
10. If a system S has the function of type T to phi, then it (T)-should do so reliably.
11. Our belief formation system constitutively (representationally)-should reliably form true beliefs. (from 9 and 10).

Note that (6)—the truth function claim—is supposed to follow from the conjunction of two claims: The claim that the representational function of our perceptual belief formation system is to successfully represent, and the claim that successful representation implies veridicality. Note, though, that the latter claim is a mere necessity claim: veridicality is necessary for successful representation. However, (6) takes veridicality to be *the* function of our perceptual belief formation system; it takes it, then, to be the success condition for function fulfilment. What would be needed, however, in support of this claim would be a stronger claim than (5), one that stipulates necessary and sufficient conditions for successful representation rather than just necessary conditions. (6) is too strong: what (5) lends support to, in conjunction with (4), is the weaker:

(6') The representational function X of our perceptual belief formation *implies* veridical representation (from 4 and 5).

Veridicality is necessary for successful representation, therefore, since the function of perception is to successfully represent, it follows that the function of perception, whatever it may be,

implies veridicality. (6)' is, of course, compatible with the function of perception being to generate beliefs that are true and that, on top of this, also feature a bunch of other properties. For all the derivation shows, then, it may well be that the function of our perceptual mechanisms is, e.g., to generate pleasant true beliefs.

In turn, if we replace (6) with (6)', the argument for Burge's favourite truth-first reliabilist view of perceptual entitlement also breaks down (independently of the failure of (10)).

1. Representational systems have representational functions.
2. The representational function of a system is to successfully represent.
3. Our perceptual belief formation system is a representational system.
4. The representational function of our perceptual belief formation system is to successfully represent (from 1–3)
5. Successful representation implies veridicality.
6. (6)' The representational function of our perceptual belief formation *implies* veridical representation (from 4 and 5).
7. Veridical propositional representations are true propositional representations.
8. Beliefs are propositional representations.
9. (9)' The representational function of our belief formation system is formation of beliefs with property X, where X implies truth (from 6–8).
10. If a system S has the function of type T to phi, then it (T)-should do so reliably.
11. (11)' Our belief formation system constitutively (representationally)-should reliably form beliefs with property X that implies truth (from 9 and 10).

The argument for Burge's preferred view of perceptual entitlement breaks down: for validity, (9) needs to be replaced by (9)', which, in turn, delivers a novel (11)' instead of (11). Since (11), but not (11)' is Burge's view of entitlement, Burge does not have an *a priori* derived view.

The alternative is, of course, to replace (5) with a biconditional in the original formulation of the argument, which would deliver the goods. (5) would then become:

(5)' A representation is successful if and only if it is veridical.

Which would, in turn, offer the needed support to (6) and, in turn, to (11). The issue is: why think (5)' holds? Recall that Burge's support for 5' relied on two main claims: one value theoretic—that truth is the chief good of representation—, and one normative—that belief that is not true is defective qua belief.

It should be clear that the normative claim will not offer support for the biconditional (5)'. After all, it is compatible with false belief being defective qua belief that true belief is also defective qua belief, in virtue of missing other crucial properties. Again, the normative claim is but a necessity claim: good belief requires truth. As such, it is too weak to lend support to the more ambitious (5)'. *Contra* Burge, failure is not *enough* evidence for *complete* function.

How about the value theoretic claim? Unfortunately, although less straightforwardly, this claim is also too weak to do the job. Even if beliefs are representational devices, and truth is the chief good involved in representation, it does not follow that truth is enough for attributively good belief—i.e., belief that is a good token of its type. And here is why: normative requirements on types need not transmit unchanged to species: in virtue of x being a species of y, norms governing y will transmit to x in that x will be governed by norms that are at least as demanding as the ones governing y. Norms governing x can, however, be more stringent than norms governing y, which

is what makes x into a particular species of y to begin with. My favourite example is that of dancing and waltzing. In virtue of being a species of dancing, waltzing will be governed by norms that are at least as demanding as those of dancing. Two-year-old children can dance. Twenty-two-year old college graduates can learn how to waltz, if they practice enough. Take the following plausible constitutive norm for dancing, N : ‘One should perform a selected sequence of movements following a rhythmic pattern.’ This requirement should be transmitted from dancing in general to waltzing on pain of waltzing not being a species of dancing to begin with. However, type membership merely requires that the corresponding constitutive norm governing waltzing is not weaker than N . The norm governing the species may be stronger, in which case norm correspondence fails. Here is the corresponding constitutive norm for waltzing: N' : ‘One should perform a triple-time sequence of movements in close position.’ Every instance of waltzing will be governed by N' . In fact, a particular sequence of movements will not qualify as waltzing to begin with if it does not comply with these norms.

To go back to Burge’s value-theoretic claim: recall that the thought was that truth is the chief good involved in representation, beliefs are representational devices, from which it would seem to follow that good beliefs are true beliefs. Just because beliefs are species of the type ‘representational device,’ though, it does not follow that the condition for ‘goodness’ is the same for beliefs and representational devices generally. In other words, the evaluative norm for belief need not be the same as that of representational devices in general: it can be stronger. Just like with dancing and waltzing, it can be that representational devices in general are governed by a truth norm, while the norm governing beliefs is stronger. If so, again, just like with the normative claim, the value theoretic claim will only support the necessity claim involved in (5)’’.

5 | KNOWLEDGE AND A *PRIORI* PERCEPTUAL ENTITLEMENT

In my view (Simion, 2016, 2019a, 2019b), generating *knowledge* is the function of our perceptual belief formation systems,⁵ and belief formation processes more generally. Crucially, in what follows, I *will not* be engaging in deriving my particular view of entitlement on *a priori* grounds, i.e. I will not take on what I dubbed ‘the Modest Derivation’ project. Rather, I would like to put forth a knowledge-first Ambitious Derivation of *a priori* perceptual entitlement: i.e., if correct, the argument below delivers the result that we are, as a matter of fact, entitled to believe based on perception, insofar as it’s properly functioning.

Let’s start with a rough sketch of the argument: I take the constitutive function of representational belief formation systems to be generating knowledge. In virtue of it being such a system, my perceptual belief formation system—if I have one—will follow suit. Constitutive functions, in turn, generate constitutive norms: as such, constitutively, my perceptual system should generate knowledge. The next step is to notice, via an analogy with games and languages, that norms constituting an entity cannot be broken with maximal systematicity, while still counting as engaging in/being the same entity: I cannot only utter ‘kakakakaka’ for the rest of my life and still count as speaking English. If so, it seems to follow that, if I have a perceptual system, it has generated knowledge at least once. Since systems that have generated knowledge at least once deliver epistemic entitlement when properly functioning, I conclude that, if I have a properly functioning perceptual belief formation system, I am entitled to believe its deliverances.

With the rough sketch in play, let’s move on to the more detailed argument. Here it goes:⁶

1. Representational systems have uniquely constitutive representational functions.

2. The representational function of a system is to successfully represent.
3. If a subject S has a perceptual belief formation system, then it is a representational system.
4. If S has a perceptual belief formation system, then its uniquely constitutive representational function is to successfully represent (from 1–3).
5. Perceptual belief formation systems successfully represent if and only if they generate knowledge.
6. If S has a perceptual belief formation system, its uniquely constitutive representational function is to generate knowledge (from 4 and 5).
7. Uniquely constitutive functions generate uniquely constitutive norms: when x 's constitutive function of type T is to *phi*, x constitutively T-should *phi*.
8. If S has a perceptual belief formation system, then it is governed by a uniquely constitutive knowledge norm (from 6 and 7).
9. If X is in maximally systematic breach of the uniquely constitutive norm for Y, then X is not identical to Y.
10. If S has a perceptual belief formation system, it is not in maximally systematic breach of the knowledge norm (from 8 and 9).
11. If a system M is not in maximally systematic breach of a norm that requires M to R, then M has R-ed.
12. If S has a perceptual belief formation system, then it has (in the past or present) generated knowledge (from 10 and 11).
13. One is *prima facie* entitled to one's belief if generated by a properly functioning system that has generated knowledge.
14. S is *prima facie* entitled to S's beliefs generated by a properly functioning perceptual system (from 12 and 13).

Note, crucially, that (14) is, *de facto*, a conditional claim: *if* we have a properly functioning perceptual system, then we're entitled to believe its outputs. Compatibly, of course, we might not have one. Most importantly for present purposes, (14) is compatible with the idea that the systems that we take to be perceptual belief forming systems are, as a matter of fact, not perceptual belief forming systems.

I will now run through all the premises in turn and, whenever needed, offer reason to believe they hold.

I take (1) and (2) to not be in need of a lot of defence. It is eminently plausible that some systems wear their constitutive functions up their sleeves, as it were, in that their dubbing is function-driven. Plausibly, toasters have toasting functions, and the toasting function of any system is to toast successfully. Washing machines have washing functions, and the washing function of any system is to wash successfully.

(3) assumes a view on the nature of perception whereby its being a representational system is an essential feature of perceptual systems *qua* perceptual systems. I take this not to be in need of much defence either. That is not to say, of course, that one can know on *a priori* grounds that my eyes etc. generate representations: surely, there could be a world populated by creatures that have the exact same mechanisms that we use for representing, only they use them, say, for digesting. Burge too acknowledges that it is an empirical question, to be answered by psychological methods, whether some given system is a perceptual system (2013). But Burge and I agree that, if a psychological system is a perceptual system, it is *a priori* that its fundamental and constitutive function is to perceive, hence to represent. This is why the argument above establishes a

conditional claim: if we have a perceptual system, then we're entitled to believe its outputs. Compatibly, we might not have one.

(4) follows from 1, 2 and 3.

(5) will likely be the source of much disagreement. For what is worth, I think disagreement is *prima facie* justified: representation seems to be about truth, not knowledge. Here is, though, a simple argument that follows Burge's quite closely to see why that's wrong. Recall Burge's plausible normative claim discussed above: a belief that is not true is defective *qua* belief. Now, note that a similar claim holds for knowledge: a belief that is true by luck is defective *qua* belief; a belief that is not justified is defective *qua* belief; a belief that is Gettierized is defective *qua* belief. Importantly, we know all this on *a priori* grounds: beliefs that fall short of knowledge in the various ways described above are criticisable (which, importantly, is not to say that the believers in question are: to the contrary, very plausibly, justified believers are not criticisable). In turn, criticisability implies evaluative norm violation: it cannot be that something X is criticisable although there is nothing attributively defective with X. If all this is so, (5) follows. As Williamson puts it, mere belief is botched knowledge (2000).⁷ Crucially, note also that knowledge is enough for non-defective belief: there's nothing wrong with a knowledgeable belief that falls short of certainty, for instance.

(6) follows from (4) and (5).

I take (7) to be in no need of defence; functions in general are widely taken to generate corresponding norms. Constitutive functions will generate constitutive norms: since the function in question will be sourced in the nature of the relevant item, the norms thereby generated will follow suit.

(8) follows from (6) and (7).

To see why (9) holds, take a paradigmatic example of a constitutive norm: norms of games. Uncontroversially, if you break too many norms of chess too flagrantly and too often, you don't count as playing chess anymore. To see this, note that if I just throw the chess pieces in the air all the time, I don't count as playing chess. When it comes to games governed by one and only one constitutive norm, this principle translates, at a minimum, in (9): if I break the only constitutive norm governing game G all the time, I'm not playing G (Kelp & Simion, 2020). For instance, if I'm breaking the only constitutive norm governing the game 'Ball in the Air,' which stipulates that I should throw the ball in the air, then I am not playing 'Ball in the Air.'

(10) follows from (8) and (9). (11), I take it, is in no need of defence, and (12) follows from (10) and (11). The perhaps hard to swallow result implied by (12) is the contrapositive (12)': If a system S has not generated knowledge, then it is not a perceptual belief formation system. Note also that, in fact, the claim implied by my argument is even stronger: given that (12)' about perception in particular follows from its being a representational belief formation system, together with the knowledge norm of belief, the assumption underlying the argument is the stronger (12)''': If a system S hasn't generated knowledge, then it is not a representational belief formation system. I think (12)''' is correct: in failing to ever—i.e., in the past or present—generate knowledgeable beliefs, the system in question has systematically failed to put us in cognitive contact with the world. In that, it is not a representational belief formation system. Compatibly, of course, it might still be a representational system *simpliciter*. The rabbit's danger detection mechanism, for instance, might still be a representational system, while it is not a representational *belief formation* system.

To see the plausibility of this further, recall Putnam's famous ant case:

An ant is crawling on a patch of sand. As it crawls, it traces a line in the sand. By pure chance the line that it traces curves and recrosses itself in such a way that it ends

up looking like a recognizable caricature of Winston Churchill. Has the ant traced a picture of Winston Churchill, a picture that depicts Churchill? Most people would say, on a little reflection, that it has not. The ant, after all, has never seen Churchill, or even a picture of Churchill, and it had no intention of depicting Churchill. It simply traced a line [...] (Putnam, 1981, p. 1).

This case lends support to content externalism: the ant does not have the concept of Churchill, since she's never seen or heard of Churchill, and, as a result, the line it draws in the sand does not stand in the right relation to Churchill to count as representing Churchill.

Now, consider the following variation on this case:

Mary has been introduced to Winston Churchill and knows him well. Now, Mary is dancing salsa at a beach party and, unbeknownst to her, by pure chance, as she dances, she traces a line in the sand that ends up looking like a recognizable caricature of Winston Churchill.

Has Mary drawn a picture that represents Winston Churchill? Again, most people would say no: the line in the sand does not stand in the right relation to Churchill to count as representing Churchill. There is no difference between the ant's line in the sand and Mary's. Mary has the concept of Churchill, but the right relation still fails to obtain: the causal chain, as it were, is interrupted on the segment from Mary's concept to the line in the sand: what explains the correlation is intervening luck (Pritchard, 2005). What this suggests is that representation is incompatible, in an important sense, with intervening luck. This, in turn, lends support to (12)" above: systems exclusively generating luckily true beliefs are not representational systems, in virtue of not putting us in cognitive contact with the world. Similarly, Mary's dancing does not count as a representational process, no matter how many times she luckily ends up drawing Churchill-looking lines in the sand. Representational belief formation systems are knowledgeable belief formation systems: they are systems that do not fail to deliver knowledge with maximal systematicity.⁸

Last, I take (13) to rest on *a priori* reflection, which reveals it to be true, in that it does not suffer exceptions. We are, indeed, intuitively *prima facie* entitled to believe whatever our properly functioning knowledge-generating processes deliver, even if they have only generated knowledge in us *once* (in the past or present). The stress is on the *prima facie* here: defeat will make all the difference. If Mary's testimony has generated knowledge in me when we first met (say, about her name being Mary), but she has been compulsively lying to me ever since, and I know it, I am not entitled to believe her. My *prima facie* entitlement gets defeated. However, absent defeat, my *prima facie* entitlement stands.⁹ For instance, if my perception delivered just of piece of knowledge and then an evil scientist started generating false perceptual beliefs in me (e.g. by exposing me to extraordinarily veridical holograms), I am entitled to believe based on perception. The reason why this is so is because, in order for a system to generate knowledge at all, it needs to be a fairly sophisticated epistemic machine, with an excellent track-record to begin with.¹⁰

Finally, the entitlement claim follows from (12) and (13): we are *prima facie* entitled to believe based on perception; again, that is, on the assumption that we have it.

6 | CONCLUSION

Tyler Burge argued that the norm according to which perception should reliably deliver truths is derivable on *a priori* grounds. I have argued that Burge's attempted derivation encounters two main difficulties, and that both are sourced in his truth-first picture, i.e., a picture on which truth is the chief representational good. In turn, I have argued, if we favour a knowledge-first picture—which, I claimed, Burge's own arguments for truth-first support—we can deliver the much more ambitious result that it is *a priori* that, insofar as we have a properly functioning perceptual system, we are *prima facie* entitled to believe its deliverances.

ENDNOTES

- ¹ I am deeply indebted to Peter Graham, Chris Kelp, Jack Lyons and Nicholas Silins for extensive comments on this paper. I owe whatever understanding I might have of Burge's views to Peter Graham (2019, 2020, forthcoming and p.c.).
- ² For an excellent overview and critical discussion of functionalist normativity, see Hazlett (2013).
- ³ Burge (e.g. 1979, 1982) is one of the main defenders of content externalism.
- ⁴ Normatively, not descriptively.
- ⁵ For knowledge-first views in the epistemology of perception, see e.g. Millar (2019), Miracchi (2017), Schellenberg (2018). For criticism, see e.g. McGrath (2016).
- ⁶ Many thanks to Chris Kelp who helped spot a number of problems with several previous versions of this derivation.
- ⁷ See also Kelp (2014) for arguments that knowledge is the goal of inquiry.
- ⁸ See also Dretske (1983).
- ⁹ In line with most literature of recent years, I bracket generality problems that may and likely will arise at this stage. See Lyons (2019) for an excellent proposal.
- ¹⁰ Is having generated knowledge only once really enough? What if the system in question did so once twenty years ago, and I've forgotten all about it? The answer is 'yes:' even so. Again, the reason why this is so is because, in order for a system to generate knowledge at all, it needs to be a fairly sophisticated epistemic machine, with an excellent track record to begin with. To see this, all that one needs to do is fully spell out the case: what is the process in question supposed to be? As soon as this is done, I predict, one of two things will happen: either the process will not have the credentials to have generated knowledge in the past, or it will intuitively deliver entitlement. Thanks to Adam Carter for pressing me on this one.

REFERENCES

- Burge, T. (1979). Individualism and the mental. In P. French, T. Uehling, & H. Wettstein (Eds.), *Midwest Studies in Philosophy 4* (pp. 73–121). Minneapolis: University of Minnesota Press.
- Burge, T. (1982). Two thought experiments reviewed. *Notre Dame Journal of Formal Logic*, 23(3), 284–293. <https://doi.org/10.1305/ndjfl/1093870087>
- Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research*, 67(3), 503–548. <https://doi.org/10.1111/j.1933-1592.2003.tb00307.x>
- Burge, T. (2010). *Origins of objectivity*. Oxford: Oxford University Press.
- Burge, T. (2013). *Cognition through understanding: Self-knowledge, interlocution, reasoning, reflection, philosophical essays* (Vol. 3). Oxford: Oxford University Press.
- Burge, T. (2020). Entitlement: The basis of empirical warrant. In P. Graham & N. J. L. L. Pedersen (Eds.), *Epistemic Entitlement*. Oxford University Press.
- Dretske, F. (1983). The epistemology of belief. *Synthese*, 55, 3–19. <https://doi.org/10.1007/bf00485371>
- Geach, P. (1956). Good and evil. *Analysis*, 17, 33–42. <https://doi.org/10.1093/analys/17.2.33>
- Graham, P. (2012). Epistemic entitlement. *Noûs*, 46, 449–482. <https://doi.org/10.1111/j.1468-0068.2010.00815.x>
- Graham, P. (2019). Why is warrant normative? *Philosophical Issues: A Supplement to Noûs*, 29, 110–128. <https://doi.org/10.1111/phis.12142>

- Graham, P. (2020). Why should warrant persist in demon worlds? In P. Graham & N. Pedersen (Eds.), *Epistemic Entitlement* (pp. 179–202). Oxford University Press. <https://doi.org/10.1093/oso/9780198713524.003.0005>
- Graham, P. (Forthcoming). What is epistemic entitlement? In C. Kelp & J. Greco (Eds.), *Virtue theoretic epistemology*. Cambridge University Press.
- Hazlett, A. (2013). *A luxury of the understanding: On the value of true belief*. Oxford: Oxford University Press.
- Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51(3), 227–232.
- Kelp, C., & Simion, M. (2020). The C account of assertion: A negative result. *Synthese*, 197, 125–137. <https://doi.org/10.1007/s11229-018-1760-5>
- Lyons, J. C. (2019). Algorithm and parameters: Solving the generality problem for reliabilism. *Philosophical Review*, 128(4), 463–509. <https://doi.org/10.1215/00318108-7697876>
- McGrath, M. (2016). Schellenberg on the epistemic force of experience. *Philosophical Studies*, 173(4), 897–905. <https://doi.org/10.1007/s11098-015-0529-7>
- Millar, A. (2019). *Knowing by perceiving*. Oxford University Press.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Miracchi, L. (2017). Perception first. *Journal of Philosophy*, 114(12), 629–677. <https://doi.org/10.5840/jphil20171141244>
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Putnam, H. (1981). *Reason, truth and history*. Cambridge University Press.
- Schellenberg, S. (2018). *The unity of perception: Content, consciousness, evidence*. Oxford University Press.
- Simion, M. (2016). Perception, history and benefit. *Episteme*, 13(1), 61–76. <https://doi.org/10.1017/epi.2015.56>
- Simion, M. (2019a). Knowledge-first functionalism. *Philosophical Issues: A supplement to Nous*, 29, 254–267. <https://doi.org/10.1111/phis.12152>
- Simion, M. (2019b). *Epistemic norms and epistemic functions*. Manuscript.
- Simion, M., Kelp, C., & Ghijsen, H. (2016). Norms of belief. *Philosophical Issues*, 26(1), 375–392. <https://doi.org/10.1111/phis.12077>
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

How to cite this article: Simion M. *A priori* perceptual entitlement, knowledge-first. *Philosophical Issues*. 2020;30:311–323. <https://doi.org/10.1111/phis.12187>