THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Two Notions of Circularity

Citation for published version:

Kallestrup, J 2015, 'Two Notions of Circularity', *Inquiry: An Interdisciplinary Journal of Philosophy*, vol. 58, no. 5, pp. 486-510. <https://doi.org/10.1080/0020174X.2015.1026391>

Digital Object Identifier (DOI):

[10.1080/0020174X.2015.1026391](https://doi.org/10.1080/0020174X.2015.1026391)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Inquiry: An Interdisciplinary Journal of Philosophy

Publisher Rights Statement:

This is an Accepted Manuscript of an article published by Taylor & Francis in *Inquiry: An Interdisciplinary Journal of Philosophy* in July 2015, available online: <http://www.tandfonline.com/10.1080/0020174X.2015.1026391>

© Kallestrup, J. (2015). Two Notions of Circularity. *Inquiry: An Interdisciplinary Journal of Philosophy*, 58(5), 486-510. 10.1080/0020174X.2015.1026391

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Two Notions of Circularity

Abstract. Crispin Wright's epistemic response to McKinsey's paradox is to argue that introspective knowledge of the first premise fails to transmit across the semantic externalist entailment in the second premise to the conclusion that one has such untoward knowledge of the external world. This paper argues first that Stewart Cohen and Jonathan Vogel's bootstrapping arguments suffer from a novel kind of epistemic circularity, which triggers failure of transmission but safeguards the possibility of basic perceptual knowledge. It is then argued that McKinsey's paradox falls out as a special case of this template for transmission failure. The circularity in play is semantic: the paradox illicitly imports semantically relevant properties of knowledge-individuating sources into the contents of the knowledge states that those sources individuate by instantiating those properties. Importantly, this diagnosis permits the possibility of basic introspective knowledge as propounded by Tyler Burge and other semantic externalists.

Keywords: epistemic circularity, reliabilism, semantic externalism, transmission of knowledge, epistemic bootstrapping, McKinsey's paradox.

I. The Incompatibilist Recipe

On the face of it, privileged access—the thesis that a subject (*S*) is always able to have a priori knowledge via introspection of the contents of thoughts—seem incompatible with semantic externalism—the thesis that those contents are known a priori to fail to supervene on her internal features.¹ Incompatibilists have deployed various strategies, but we shall only probe into what Brown (2004: 194) calls the *consequence problem*: the joint assumptions of semantic externalism and privileged access pave the way for a priori access to those external environmental features which partly determine the thought contents in question. But it is supposedly absurd that armchair philosophical theorizing should lead to a priori knowledge of the external world given a priori knowledge of thought contents. Consider what Davies (2003b) dubs the (MC) form:

- (1) *S* has mental property *M*
- (2) *S* meets non-mental condition *C* if *S* has mental property *M*

¹ Four caveats: first, some prefer to call introspective knowledge 'non-empirical' rather than 'a priori', but as nothing hangs on this terminology we shall henceforth use the latter label. Second, 'content' is throughout understood as conceptual content. Third, by 'thoughts' we shall from now on mean occurrent thoughts. Fourth, semantic externalists have typically taken the failure of supervenience of thought contents on internal features to be justified by Putnam- or Burge-style thought experiments which is why the view is defined as making an a priori claim. Those semantic externalists who take the view to be only empirically supported do not face our incompatibilist challenge. See Author (2011) for more details.

(3) So, *S* meets non-mental condition *C*

In order to use this schema to mount an incompatibilist argument, *C* must not only be an external condition the obtaining of which makes *M* a wide content mental property, *S* must also be in a position to know a priori that *C* obtains if she has *M*. For if *S* then combines that knowledge with her a priori knowledge that she has *M*, then it looks as if she can know a priori that she meets *C*. For instance, if *M* is the property of thinking thoughts involving the concept **water**, then *C* might be the condition that *S* has causally interacted with water. Thus McGinn (1989: 30-48) and Davies (2003b), argued that semantic externalists should embrace a fairly strong constraint on concept-possession:

(CC) If the concept of *X* is an atomic natural kind concept, then *S* cannot possess that concept unless *S* has causally interacted with instances of *X*

Here an *atomic concept* is one that lacks conceptual constituents, and a *natural kind concept* is one that both purports to and actually succeeds in picking out (instances of) a natural kind. To require that *S* herself has encountered water may be too demanding. Perhaps, as Wright (2000: 156) suggests, all that can be inferred from *S* possessing atomic natural kind concept *X* is that either *S* has causally interacted with instances of *X* or else *S* is a member of a linguistic community in which others have causally interacted with such instances. But it would be no less surprising if all that one could know a priori were that others have causally interacted with instances of *X*. So, in the following we shall assume a *natural-kind-dependent thought* view according to which it is a priori knowable that if *S* possesses the atomic natural kind concept **water**, then *S* has a history of causal encounters with water.² Consequently, it is a priori knowable that water exists in *S*'s external environment if *S* possesses such a concept. Consider now the following argument for incompatibilism³:

² Note three features of this view: (i) *S*'s grasp of **water** may be incomplete yet sufficient to think thoughts involving that concept. (ii) *S* can express contentful thoughts when uttering sentences containing 'water' despite the absence of water in the immediate vicinity. (iii) The apriority stems from armchair reflection on possible scenarios. On Boghossian's Dry Earth (1998), for instance, where despite all appearances water has never been present, there is no atomic natural kind concept of water to exercise in thought for those who have always inhabited this waterless planet. Note that such dry-earthlings may still succeed in expressing a different concept of water such as the compositional concept **the watery stuff**.

³ Here I am grateful to an anonymous referee.

- (P1) If privileged access and semantic externalism are both true, then *S* can know the premises (1) and (2) of an MC-argument a priori
- (P2) If so, then *S* can know the conclusion (3) of an MC-argument a priori
- (P3) But no conclusion (3) of an MC-argument can be known a priori
- (C) So, privileged access and semantic externalism are not both true

Most semantic externalists reject the natural-kind-dependent thought view on which premise (P1) relies. Thus, Korman (2006) argues that since the compositionality of concepts is an a posteriori matter, *S* cannot know a priori that **water** is an atomic concept. On his view, 'water' expresses an atomic concept only if such tokens successfully refer to instances of water. If all tokens of 'water' are empty, as on Boghossian's Dry Earth (1998), then a compositional concept such as **the watery stuff** is expressed. Since *S* cannot know a priori whether her tokens of 'water' are empty, she cannot know a priori which concept these tokens express. Relatedly, McLaughlin and Tye (1998) argue that although *S* can know a priori that she has the concept **water**, she cannot know a priori that **water** is a natural kind concept. On their view, **water** is a natural kind concept only if water is a natural kind. And since *S* cannot know a priori that water is a natural kind, she cannot know a priori that **water** is a natural kind concept. After all, for all *S* knows a priori, water might be a motley kind. On this view, there is no external condition the obtaining of which follows a priori from the conjunction of privileged access and semantic externalism.

Rejecting (P1) on grounds of specific views about atomic or natural kind concepts may get the semantic externalist off the hook vis-à-vis our incompatibilist argument. Nevertheless, here are two reasons why a different tack is preferable. The first is that there are related views according to which *S* can know the premises (1) and (2) of an MC-argument a priori. Thus, Pryor (2007) observes that Evans (1982) and McDowell's (1984) view of object-dependent thoughts implies that exemplifications of (1) and (2) are a priori knowable. They define so-called *de re* thoughts to be related to specific objects of acquaintance in such a way that if, despite appearances, those objects were absent, there would be no thoughts at all. *S* would undergo an illusion of content. But they advocated a Fregean view of thoughts

according to which their contents are composed of a priori accessible modes of presentation or ways of thinking of those objects. If that's right, an influential stripe of semantic externalism cannot reject (P1).

The second reason is that since all (2) requires is that condition *C* be non-mental, the (MC) form besets unrelated views too. Any putative response that relies on specific semantic externalist commitments is thus at best lacking in generality. Take a view on which *C* in (2) pertains to internal features of embodiment such as information-processing mechanisms inside the head. Consider what Davies (2003a) calls *architecturalist entailments* of the form:

(2*) *S* has internal cognitive architecture *A* if *S* has mental property *M*

As an example of such machinery being a priori connected to the mental Davies (op. cit.) offers an a priori two-stage argument for the language of thought hypothesis. If *M* is the property of thinking certain contents, the (MC) form seems to underwrite a priori knowledge that *S* has *A*. Gaining a priori access to such architecturalist features is obviously equally problematic. Or think of *analytical functionalism*, which McLaughlin (2003: 93-94) cites as subject to the (MC) form. This view holds that for any mental property *M* there is some causal role *c* such that it is a priori knowable that *S* has *M* if and only if *S* has an internal (physical) property that plays role *c*. For a range of mental properties *M*, *S* can also know a priori that she has *M*. Yet by any reckoning *S* cannot know a priori that she has an internal (physical) property that plays role *c*.⁴

The upshot is that we should ideally opt for a response to the incompatibilist argument which is available to *all* semantic externalists, and which can be adapted by architecturalists and analytical functionalists to avoid similar arguments leading to a priori knowledge of internal cognitive architectures or causal role-playing physical properties. What is needed is a diagnosis of the incompatibilist argument which permits *S* a priori knowledge of the premises (1) and (2) of an MC-argument, no matter what *M* and *C* are, but denies *S* a priori knowledge of the conclusion (3) of such an argument. That is to say, the aim is to find a principled way for the semantic externalist to reject (P2).

⁴ Davies (2003a) and McKinsey (2003) provide additional examples of views that are threatened by the (MC) form. These concern, for instance, indexical thoughts, color concepts and singular thoughts as expressed by proper names and demonstratives.

One last preliminary remark is crucial. If both (1) and (2) are a priori knowable, the question is which epistemic principle would allow S to attain a priori knowledge of (3). At first blush, what is required here is not merely that knowledge, or specifically a priori knowledge, be closed under known entailment:

(Closure) If S knows that p , and S knows that p entails q , then S knows (or at least is in a position to know) that q

This principle imposes a consistency requirement on knowledge ascriptions. Thus, (Closure) bans the possibility of consistently ascribing to S knowledge that p , and that p entails q , without also ascribing to S knowledge that q . Likewise, if (Closure) holds, then S cannot consistently be ascribed knowledge that p in conjunction with ascriptions of knowledge of the entailment from p to q and lack of knowledge that q . The point is that (Closure) is silent on the source of S 's knowledge that q , and so allows for the possibility that S 's knowledge that q is independent of the kind of knowledge S has that p and that p entails q . But the distinctive incompatibilist claim is that if privileged access and semantic externalism are true, then S could come to know (3) *in virtue of* knowing (1) and (2).⁵ The (MC) form is designed as a recipe for knowledge-transmission and not merely knowledge-preservation. What is called for is rather Wright's principle (2002, 2003), roughly that⁶:

(Transmission) If S knows that p , and S comes to believe that q on the basis of competently deducing q from p , then S can thereby come to know that q

⁵ The in-virtue-of relation is intended as logically stronger than the epistemic basing relation. The latter is a relation between S believing that p and a reason R , e.g. an experience or belief, such that R is the reason for which S holds that belief. We shall take the former to be a dependency relation between S knowing that p and those epistemic features which make S know that p given S 's true and properly based belief that p . In this case, S 's belief that q is based on her belief that p in that she deduces the former belief from the latter belief. What must be added in order for S to know that q in virtue of knowing that p is something about what turns S 's true belief that q into knowledge given the way that belief is based. See also Section II and Tucker (2010) who uses the in-virtue-of locution as a placeholder for whatever non-deviant causal relation is required for the transmission of epistemic properties.

⁶ Wright (op. cit.) prefers the notion of 'warrant' and occasionally 'justification', but to keep things simple we shall use 'knowledge' throughout.

Here ‘deduction’ applies only to valid inferences. If p does not entail q , S can infer but not deduce q from p . Competence might consist in recognizing the validity of the inference; or so Wright (op. cit.) proposes. Maybe additional fine-tuning is called for, e.g. to ensure that S retains knowledge that p while deducing q from p . What matters for present purposes is the idea that S bears an epistemic relation to q at least in part because S bears that relation to p . In other words, knowledge that q is partially dependent in kind on knowledge that p . The claim is typically not that the very same kind or piece of knowledge is transmitted across entailment. While the incompatibilist could reasonably couch a version of (Transmission) in terms of a priori knowledge, other kinds of knowledge clearly fail to transmit. Thus Pryor (2004) contends that (Transmission) holds in Moore’s proof, but while S has merely visual, non-inferential knowledge that she has hands, believing that the external world exists on the basis of competent deduction would at most yield visual-cum-inferential knowledge.⁷ Let’s say one piece of knowledge is *epistemically prior* to another piece of knowledge if and only if the latter asymmetrically depends in kind on the former. In that case (Transmission) sanctions epistemic priority of knowledge that p over knowledge that q . That is to say, if S comes to know that q by forming the belief that q on the basis of competent deduction from p of which S already has knowledge, then S ’s knowledge that p is epistemically prior to her knowledge that q . Such epistemic priority is inessential to (Closure).

Here’s the plan for the remaining parts of this paper. Our contention was that a satisfactory response to the incompatibilist argument to the effect that S can know the conclusion of an (MC)-argument a priori should not hang on specific views about the externalist nature of concepts, our cognitive architecture or the functionalist character of mental properties. Only an epistemic solution will afford sufficient uniformity as it identifies a fault with the cogency of the reasoning behind such an argument. Such a solution would be independent of particular exemplifications of M and C , and so would be available regardless of any such specific views. In the next two sections a novel epistemic solution to the (MC)-argument will be expounded. Section II argues that while Wright correctly blames the deployment of (Transmission) in so-called

⁷ See also Wright (2007: 44-47).

bootstrapping arguments, his influential diagnosis implies a claim which is incompatible with the type of epistemological externalism that is troubled by these arguments. Instead a new template is developed for diagnosing a generalized notion of epistemic circularity as in bootstrapping arguments. Section III argues that while Wright also correctly takes the (MC)-argument to trigger failure of (Transmission), his diagnosis implies a claim which is incompatible with the type of semantic externalism that allegedly succumbs to this argument. Instead it is shown how the incompatibilist argument falls out as a special case of the improved template—one that exhibits a distinctive kind of semantic circularity. Finally, Section IV briefly sums up the foregoing arguments.

II. Epistemic Circularity

Epistemic bootstrapping is typically taken to consist in *S* coming to know that an object or process is reliable by relying (at least in part) on the reliability of that object or process. For example, suppose *S* forms the belief that the NYT is reliable by reading in the NYT that this newspaper is reliable. Even if *S*'s belief is true, it falls short of knowledge because *S* cannot intuitively come to know that such a knowledge source is reliable by relying (at least partially) on the reliability of that very source. After all, had the NYT been unreliable, *S* could not come to know any of its reports by reading them. Or suppose *S* comes to know by visual perception that a table is red, and suppose *S* also comes to know by introspection that visual perception produced the belief that the table is red. *S* can then come to know by reasoning that visual perception produced a true belief in this occasion. If *S* repeats that process on numerous occasions, *S* can come to know by induction that visual perception is reliable. But *S* cannot intuitively come to know that visual perception is reliable by relying (at least in part) on the reliability of visual perception. What exactly explains that intuition is a vexed question to which we shall attempt an answer in what follows.

Vogel (2000; 2008) has argued that process reliabilism is susceptible to epistemic bootstrapping, and Cohen (2002; 2005) has shown that this problem can be generalized to evidentialist internalism, roughly the view that justification consists in evidence that is internal to one's mind. As we are primarily interested in whether (Transmission) fails in certain arguments against—epistemological and semantic—

externalist views, we shall from now on restrict attention to process reliabilism. Thus suppose a fairly unadorned version of process reliabilism (PR) is true: *S* knows that *p* if and only if (a) *S*'s true belief that *p* is produced by a reliable process *r*, and (b) the epistemic status of *S*'s belief that *p* is not undermined. As regards (a), belief-producing processes are henceforth to be understood as cognitive processes such as those involved in visual perception, introspection, etc. As regards (b), Goldman (1986: 62-3, 111-2) invoked a specific non-undermining condition in terms of neither having misleading evidence nor believing that *r* is unreliable, so as to handle reliable clairvoyance and Truetemp cases, but (b) could well include other types of undermining defeat. For instance, Vogel (2008: 534-535) takes bootstrapping arguments to suffer from rule-circularity, which he in turn views as a kind of undercutting defeater of justification.⁸ And Brueckner (2013: 596) argues that the process reliabilist can simply mimic Vogel's evidentialist response by tacking on to the reliabilist view such a 'no-rule-circularity', non-undermining condition. According to Brueckner, evidentialist internalism and process reliabilism can thus both avail themselves of the notion of undercutting defeat, as triggered by the kind of rule-circularity that Vogel detects in bootstrapping arguments. The key point for now is that nothing stops process reliabilists from helping themselves to additional non-defeater conditions as part of a response to the challenge posed by bootstrapping. Consider again the following *bootstrapping argument*:

- (4) The table is red
- (5) *S*'s visual perception produced the belief that the table is red
- (6) *S*'s visual perception produced a true belief that the table is red
- (7) *S*'s visual perception produced true beliefs on numerous occasions
- (8) *S*'s visual perception is reliable

⁸ Roughly, an undercutting defeater is one that undermines the connection between the justification for believing *p* and the belief that *p*. In contrast, a rebutting defeater is one that outweighs the original justification by providing some other justification for believing *not-p*. We are only concerned with types of undercutting defeat.

S 's beliefs in (4) and (5) are produced by S 's visual perception and introspective ability, respectively—both of which are reliable. (6) follows deductively from (4) and (5), and competent deduction is reliable. (7) repeats (6), thus amassing track-record evidence for the inductive conclusion (8). Memory and induction are both reliable. Since each step seems to be underwritten by (PR), it looks as if S is in a position to know (8) on the basis of knowing (1) – (7). But intuitively S should not be counting on the reliability of visual perception in an argument with the conclusion that visual perception is reliable. This raises Vogel's (2008) *rollback problem* of where to cut the straps. We submit that S 's knowledge of (6) is problematic. Given that visual perception produced S 's belief in (4), S knows (4) only if visual perception produced a true belief on this occasion, but S should not be counting on the accuracy of visual perception in an argument with the conclusion that visual perception is accurate. Surely, to determine whether S 's visual perception was correct on this occasion, S would rather need to scrutinize the conditions for observation, ensure her perceptual apparatus is not malfunctioning, or ask another reliable perceiver to take a good look at the table. Apart from the intuition that knowledge of (6) requires some independent inquiry into the credentials of S 's visual perception, there are at least three reasons why this is the right place to undo S 's knowledge. (i) Cohen (2002: 317) points out that even if only the track-record evidence in (7) can facilitate knowledge of (8), (6) itself constitutes some evidential support for (8). The problem is not so much whether S can use the bootstrapping argument to acquire knowledge of (8), but that S should be able to bootstrap any evidence for (8) via that argument, whether sufficient for knowledge or not. (ii) White (2006: 546-547) argues that (7) provides the best explanation of (8), indeed it is hard to see why (7) should be true unless (8) is also true. So, if the argument is allowed to proceed to (7), (8) ought to follow. (iii) Titelbaum (2010: 120-121) observes (using a different example) that if S knows independently that S 's visual perception is either anti-reliable (always wrong) or else

reliable, then she can deductively infer (8) from (6). *S* can thus bootstrap knowledge of (8) while bypassing (7).

The foregoing suggests that (Transmission) fails in the argument (4) – (6).⁹ Given the way *S* knows (4), *S* cannot come to know (6) by coming to believe (6) on the basis of competent deduction from (4) and (5) even if such deduction is reliable. How exactly such failure is best explained is a vexed issue. Wright (2000: 155) argues that certain arguments display failure of (Transmission) just in case they are such that (i) *p* entails *q*, (ii) *p* is incompatible with *p*^{*}, (iii) *S*'s state of knowing that *p* is subjectively indistinguishable from the belief state *S* would be in if instead *p*^{*} were true, and (iv) if *q* were false then *p*^{*} would be true. Wright (op. cit.) thinks any argument that fits this *disjunctive template* triggers failure of (Transmission). The key feature is that *S*'s knowledge that *p* is conditional on *S* having antecedent knowledge that *q*. The reason is that *S* can know *p* only if *S* has additional knowledge that *not-p*^{*}, which in turn requires *S* to already know that *q*. Consequently, *S*'s knowledge that *p* fails to transmit across the entailment to *q*. True, Wright prefers to cash this template out in terms of warrant and acceptance rather than knowledge, but since Vogel, Cohen and others formulate the bootstrapping problem in terms of knowledge, we shall henceforth follow their lead to avoid further complexities.¹⁰

Now, in the short bootstrapping argument, *p* = (4), *q* = (6) and *p*^{*} = the table is white with red lights shining on it. (5) is an extra premise needed to deduce *q* from *p*. The key contention is that knowing that *q* is a precondition on knowing that *p*. More precisely, the idea is that in arguments meeting Wright's disjunctive template, *S*'s knowledge that *p* depends on *S* having *default* a priori knowledge that *q*. Since default

⁹ Strictly speaking, we need here a two-premise version of (Transmission), but in order to ease exposition we shall continually refer to the single-premise version.

¹⁰ Wright's claim (2003: 57-58) is then that warrant fails to transmit from *p* to *q* in cases where "...there is warrant for [*p*] in the first place only because [*q*] is antecedently warranted". See also his (2000: 157). We shall return to this point. Compare also with Davies (2008: 365): "For one of the premises, *P*, the warrant, *W*, to believe *P* depends on an antecedent warrant to believe the conclusion, *Q*." The same thought is expressed by Pryor (2004: 359) when he talks about type 5 epistemic dependence: "having justification to believe the conclusion *is among the conditions that make you have the justification you purport to have for the premise*."

knowledge is non-acquired knowledge and all evidential knowledge is acquired knowledge, Wright's proposal is that *S*'s knowledge that *p* depends on *S* having *non-evidential, a priori knowledge* that *q*. Again, bearing in mind that Wright isn't primarily concerned with knowledge, we can approximate this default, non-evidential epistemic standing to his notion (2003: 67) of an *entitlement*. Obviously, if knowing that *q* is part of what makes *S* know that *p*, then *S* cannot acquire knowledge that *q* on the basis of coming to believe that *q* by competently deducing *q* from *p* of which *S* already has knowledge. Such reasoning would be viciously circular. If the argument (4) – (6) *did* display that epistemic structure, it would beg the question when propounded with the purpose of convincing someone who had antecedent doubts about the conclusion (6).

The point is now that Wright's generic explanation of why (Transmission) fails cannot be extended to the bootstrapping argument as levelled against (PR), even if, as we assume, that explanation could be extended to include knowledge.¹¹ The main reason is that process reliabilists are committed to a claim which is incompatible with Wright's diagnosis. Let me explain. What, according to (PR), matters for knowledge is primarily that *S*'s true belief in actual fact be produced by a reliable process. As Cohen (2002: 310) remarks, process reliabilists permit the possibility of:

(Basic Perceptual Knowledge) *S* knows that *p* even though *S* has no antecedent knowledge that the process *r* that produced *S*'s belief is reliable

where 'antecedent' denotes epistemic priority.¹² Reflect that this principle allows for *S* to know that *p* even though *S* lacks antecedent knowledge of any kind that *r* is

¹¹ Wright (2011: 34-35) argues that "...the easy warrant inferences involve a failure of *transmission* of warrant..."

¹² To be precise, what process reliabilists permits is the *possibility* of someone coming to know that *p* without prior knowledge that the process that produced the belief that *p* is reliable. That possibility is compatible with the existence of someone else who knows that *p* yet cannot help but also know (perhaps tacitly) that the process by which the belief that *p* was formed is reliable. Here I thank an anonymous referee.

reliable. Wright's claim (2003: 70) that *S* knows that *p* only if *S* has antecedent, non-evidential, a priori knowledge that *r* is reliable is not one process reliabilists would take on board. Whether that antecedent knowledge is evidential or a priori is simply irrelevant from their perspective. Of course reliability is just one property of *r* of which *S* need have no knowledge antecedent to her knowledge that *p*. Producing a true belief that *p* on this occasion will be another such property. Thus friends of (PR) will be equally happy to accept the possibility of:

(Basic Perceptual Knowledge*) *S* knows that *p* even though *S* has no antecedent knowledge that the process *r* that produced *S*'s belief produced a true belief on this occasion

To wit, (PR) has it that *S* knows that *p* as long as (a) *S*'s true belief is produced by some reliable process *r*, and (b) the epistemic status of that belief is not undermined, regardless of whether *S* has antecedent knowledge that *r* produced a true belief on this occasion. Again, this principle imposes no qualifications on the character of the prior knowledge in question. In particular, process reliabilists would balk at Wright's claim (2003: 67) that *S* knows that *p* only if *S* is entitled "...to assume the proper functioning of [her] perceptual apparatus on a particular occasion." On their view, having such entitlement is consistent with, but not a necessary condition on, knowing *p*. That is to say, *S* can know *p* even though *S* lacks antecedent, non-evidential, a priori knowledge that *r* produced a true belief on this occasion. As far as they are concerned whether the latter knowledge is evidential or a priori is neither here nor there. One might suggest that failure of (Transmission) by the disjunctive template is compatible with (PR) if such failure is classified as a way for the epistemic status of *S*'s belief to be undermined. True, but in that case (PR) would be incompatible with (Basic Perceptual Knowledge*). For condition (b) would then include a type of undermining defeat which would rule out the possibility of (Basic Perceptual Knowledge*).

The upshot is that Wright's disjunctive template implies the falsity of a claim, which is part and parcel of process reliabilism, namely the possibility of (Basic

Perceptual Knowledge*). Consequently, process reliabilists cannot consistently help themselves to that template as part of a response to the bootstrapping argument. Bear in mind that such disagreement over basic knowledge is consistent with commitment to the possibility of knowledge that is not acquired on the basis of inference from some other known proposition. In fact, Wright (2003: 62, 70) explicitly applies his template to *non-inferential* epistemic standings: perception and other cognitive faculties provide non-inferential knowledge (or warrant) conditional on having antecedent entitlements to rely on those faculties. And process reliabilists are also keen to embrace the acquisition of knowledge directly via the reliable and accurate operation of some such faculties. The key difference is that only process reliabilists can endorse the combination of non-inferential and basic perceptual knowledge. On their view, there simply is no need for antecedent knowledge of any kind that the belief-producing process instantiates any of the epistemic properties in virtue of which the belief counts as knowledge.¹³

As mentioned above, the foregoing appears to ignore an important feature of Wright's entitlements. In (2004: 176-177) Wright proposes that an entitlement renders a proposition worthy of *acceptance* (or trust) but not belief. If knowledge requires belief and not *mere* acceptance, it thus seems as if, after all, his view is consistent with the possibility of (Basic Perceptual Knowledge*). In response, note that Wright (op. cit.) elsewhere suggests that acceptance be construed as a more general attitude that includes belief as a sub-case. So, if both entitlement and knowledge involve acceptance, thus understood, then Wright's view would still be inconsistent with the possibility of (Basic Perceptual Knowledge*). In any case, we could avoid assuming that *S* actually forms belief in the target proposition *p* by simply reformulating both

¹³ For a related but different objection to Wright's template see Brown (2003). Her line is that epistemological externalists would reject Wright's first claim that in order for *S* to know that *p*, *S* must have antecedent knowledge that any incompatible and subjectively indistinguishable alternative *p** does not obtain. *S* need only know that nearby or relevant alternatives fail to obtain. See also McLaughlin (2003: 87-89) who points out that not every argument that fits this disjunctive template fails to transmit knowledge (or warrant). The template incorrectly predicts that no valid argument from *p*, when known in some non-inferential, non-logically conclusive way, to *q* would transmit knowledge. Wright's reply (2003, 62, fn. 4) is to change condition (ii) such that *p** is incompatible, not with *p* itself, but with the reliable operation of the cognitive capacities involved in coming to know *p*. However, this amended version of (ii) has no material impact on the current objection to the template.

the bootstrapping argument and the transmission principle in terms of propositional justification, i.e. justification that S has to believe p whether or not S forms that belief. In that case, corresponding versions of basic perceptual propositional justification seem incompatible with Wright's take on transmission-failure of justification.

Here is an alternative proposal. Suppose S is in knowledge state K with content p in virtue of belief-producing source r (or process if you like) instantiating epistemically relevant property F . What counts as epistemically relevant properties of belief-producing sources depends to some extent on one's preferred epistemology. So, epistemological externalists and their opponents might disagree whether conscious awareness or reliability of visual perception are epistemically relevant properties. We need not resolve that disagreement here. Note instead that in order for S to know that p in virtue of r instantiating F , S 's belief that p must be based on deliverances of r .

Assuming a causal view of epistemic basing, this means roughly that S 's belief that p must be non-deviantly caused by such deliverances. For instance, if S is to know that the table is red in virtue of visual perception instantiating the property of being reliable, S 's belief that the table is red must be suitably causally connected to an experience of seeing the red table. But as adumbrated in fn. 3, such proper basing is merely a necessary condition. In order for the in-virtue-of relation to obtain it must additionally be the case that r instantiating F is part of what converts S 's properly based belief that p into knowledge. To use our example, what converts S 's properly based belief that the table is red into knowledge is at least in part that visual perception is a reliable belief-producing source, or perhaps that visual perception produces a justifying experience of which S is consciously aware. We can say that belief-producing source r constitutes a *knowledge-producing source* just in case (i) deliverances of r non-deviantly cause the belief that p , and (ii) r instantiates enough epistemically relevant properties to convert S 's belief that p into knowledge. Fortunately, we can elucidate the notion of an epistemically relevant property without undertaking the highly controversial task of specifying exactly which properties are

required for the conversion in (ii). The proposal is to understand *epistemic relevance* in counterfactual terms:

(ER) If belief-producing source r generates S 's belief that p , then F is an epistemically relevant property of r if and only if: had r not instantiated F , then S 's belief state, as produced by r ¹⁴, could not qualify as a knowledge state K with content p

For instance, reliability is arguably an epistemically relevant property of visual perception in that the following counterfactual is true: had visual perception not been a reliable source of belief-production, then S 's belief that the table is red, as produced by that source, could not amount to knowledge that the table is red. Thus, many epistemological externalists insist that only reliable belief-producing sources also produce knowledge. But note that super-reliability, i.e. being more reliable than needed for knowledge, is epistemically irrelevant by (ER). Assume S 's visual perception would have been reliable (enough for knowledge), had it not been super-reliable. In that case, the following counterfactual is false: had visual perception not been a super-reliable source of belief-production, S 's belief, as produced by that source, could not amount to knowledge. The same is true of infallibility in cases where S 's visual perception would have been reliable (enough for knowledge), had it not been infallible.¹⁵ That some such properties turn out epistemically irrelevant is a slightly odd consequence of (ER) counting properties as epistemically relevant only if they *make the difference* between belief and knowledge, as captured by the truth of such counterfactuals. Reliability is a difference-maker in that technical sense, but super-reliability and infallibility are not.

So far, we have assumed that S is in knowledge state K with content p in virtue of belief-producing source r instantiating epistemically relevant property F . Now assume also that S is in knowledge state K with some different content q , where some distinct source r^* produces the belief that q . With those two assumptions in place, we

¹⁴ This clause is needed to avoid cases where the nearest possible world in which r fails to instantiate F is one where S 's belief state is produced by distinct source r^* instantiating enough epistemically relevant properties for that state to count as a knowledge state.

¹⁵ I owe this example to an anonymous referee.

can now identify epistemically circular arguments as those in which the conjunction of p and q , or in a special case just p itself¹⁶, entails Fr . Given that S knows each of the premises p and q , (Transmission) would then permit S to come to know Fr . Yet, S has knowledge of premise p in virtue of r instantiating F . In short, epistemically circular arguments fit the following template:

(9) $K(p)$ in virtue of Fr

(10) $K(q)$

(11) $p \ \& \ q \Rightarrow Fr$

(12) $K(Fr)$

Consider again our short bootstrapping argument:

(4) The table is red

(5) S 's visual perception produced the belief that the table is red

(6) S 's visual perception produced a true belief that the table is red

This argument illustrates our template in the following way: (9) S knows that (4) the table is red in virtue of (6) S 's visual perception producing a true belief that the table is red; (10) S knows that (5) S 's visual perception produced the belief that the table is red; (11) if (4) the table is red and (5) S 's visual perception produced the belief that the table is red, then (6) S 's visual perception produced a true belief that the table is red; (12) S knows that (6) S 's visual perception produced a true belief that the table is red. So, the contention is that the bootstrapping argument (4) – (6) exhibits a kind of epistemic circularity pertaining to a distinct epistemically relevant property. Such circularity is typically characterized in terms of knowing a premise in virtue of a certain belief source being reliable or trustworthy in an argument with the conclusion that that source is reliable or trustworthy.¹⁷ Given that (6) pertains to the property of producing a true belief rather than these epistemic properties, a more encompassing

¹⁶ To use our earlier example, you cannot come to know that the NYT is reliable by reading in the NYT that it is reliable, assuming reliability is an epistemically relevant property of that newspaper qua belief-producing source.

¹⁷ Alston (1986: 9-10) and Bergman (2004: 710) are cases in point.

characterization is called for. Importantly, our template (9) – (12) makes no special assumptions about which properties count as epistemically relevant. So, against the backdrop of that template, *generalized epistemic circularity* can be defined as follows:

(GEC) An argument is epistemically circular if and only if (i) *S* knows at least one of the premises *p*, *q* in virtue of belief-producing source *r* instantiating epistemically relevant property *F*, and (ii) the conclusion that *r* has *F* can be deduced from the conjunction of the premises *p* & *q*, or in a special case just *p* itself.

In our bootstrapping argument, *p* = (4) the table is red, *F* = producing a true belief that the table is red, *r* = *S*'s visual perception, *q* = (5) *S*'s visual perception produced the belief that the table is red, and *Fr* = (6) *S*'s visual perception produced a true belief that the table is red. The property of producing a true belief that the table is red is epistemically relevant to *S*'s knowledge that the table is red given that her belief that the table is red is based on what the table looks like. If *S*'s visual perception had failed to instantiate that property, *S* could not know that the table is red on the basis of looking at the table. According to (ER), the relevant counterfactual that tests for epistemic relevance is thus true. The problem is that when *S* puts her knowledge that the table is red together with her knowledge that her visual perception produced her belief that the table is red, she can deduce that her visual perception produced a true belief that the table is red—in fact given (Transmission) she is positioned to know exactly that. Knowledge of the first premise relies on a belief source *r* instantiating an epistemically relevant property *F* in an argument with the conclusion that *r* instantiates *F*. The argument is therefore afflicted by (GEC).¹⁸

¹⁸ This bootstrapping argument also exhibits a knowledge-version of Pryor's (2004: 358-359) type 2 epistemic dependence: in order for *S* to know one of the premises, it is necessary that the conclusion be true, and his type 4 epistemic dependence: evidence against the conclusion would undermine the knowledge *S* purports to have of one of the premises. Importantly, the epistemically benign arguments that Pryor cites as exhibiting type 2 and type 4 epistemic dependencies are not properly classified as epistemically circular by (GEC). So, these arguments provide no reason to think that (GEC) is not always an epistemic vice.

The thought is then that banning such epistemically circular knowledge blocks our bootstrapping argument. More precisely, the generalized epistemically circular nature of the putative knowledge of (6) undercuts the epistemic support that (4) and (5) would otherwise provide for (6). Consequently, such circularity triggers failure of (Transmission): *S* cannot acquire knowledge of (6) by competently deducing (6) from (4) (and (5)) of which *S* already has visual-perception-dependent knowledge. The dependency claim is important: the competent deduction of (6) from (4) and (5) would transmit knowledge had *S*'s knowledge of (4) been independent of whether her visual perception instantiated the epistemically relevant property that features in (6). Note that generalized epistemic circularity triggers failure of (Transmission) even if competent deduction is reliable. Still, our solution is consonant with (PR) provided the non-undermining condition in (b) is plausibly interpreted so as to include generalized epistemic circularity as a kind of undercutting defeat.¹⁹ That may of course be challenged. After all, bootstrapping inferences might well be reliable. We could stipulate that *S* does not indiscriminately engage in what Kornblith (2009: 265) dubs “promiscuous” bootstrapping, but only bootstraps in cases where the pertinent processes are reliable. And indeed if we opt for a safety version of (PR), such as Vogel's (2000: 605) neighborhood reliabilism, according to which *S*'s belief is reliably produced if and only if in most nearby worlds, *p* is true if *S* believes that *p* as a result of *r*, then *S*'s belief in (6) is safe: in most nearby worlds, *S*'s visual perception produced a true belief that the table is red if *S* believes just that as a result of visual perception, introspection and deductive reasoning.²⁰ To be sure, the claim is not that

¹⁹ Pryor (2004) distinguishes between the ineffectiveness of an argument and its epistemic structure, Burge (2003) rejects the entailment from being dialectically useless to being lacking in warrant, and Bergman (2004) maintains that epistemic circularity is benign if only the argument is propounded in an unquestioned context. This raises the question of whether epistemic circularity renders an argument unfit for resolving doubt while still being subject to (Transmission). To provide a thorough answer to this intriguing question is beyond the scope of this paper, but see Wright (2007) and Author (2011) for further details.

²⁰ Reflect that *S*'s belief in (6) is insensitive in the sense that if *S*'s visual perception produced a false belief that the table is red, then *S* would still believe that *S*'s visual perception produced a true belief that the table is red as a result of visual perception, introspection and deductive reasoning. The bootstrapping argument is easily blocked if Vogel's (op. cit.) counterfactual (or tracking) reliabilism is adopted.

(Transmission) fails because drawing the inference from (4) and (5) to (6) takes *S* from safe to unsafe beliefs. On this view, safe belief is merely necessary for knowledge. As illustrated by the non-undermining condition, process reliabilists have typically invoked additional conditions on knowledge, and there is no reason why the absence of generalized epistemic circularity should not constitute one such condition. An inference may thus transmit safe belief without knowledge due to the affliction of such circularity. To repeat, Vogel (2008: 534) takes rule circularity as an undercutting defeater of the epistemic structure of the bootstrapping argument, and Brueckner (2013: 596) shows how the process reliabilist can mimic Vogel's response on behalf of evidentialism. Hence, when faced with the bootstrapping problem, process reliabilists can equally well embrace our proposal without inconsistency.

III. Semantic Circularity

Now, let's revisit the incompatibilist charge posed by the (MC) form. Recall that we are taking for granted a natural-kind-dependent thought view on which it is a priori knowable that *S* cannot have the concept **water** unless water exists in her global external environment. We allow for *S*'s grasp of **water** to be potentially incomplete yet sufficient to think thoughts involving that concept. Assuming *S* has privileged access to the fact that she possesses **water**, we arrive at the following exemplification of the (MC) form in which both premises are a priori knowable:

- (13) *S* has the concept **water**
- (14) Water exists if *S* has the concept **water**
- (15) Water exists

On the assumption that (13) and (14) are true, this argument is sound. Is it also cogent such that *S* could come to know (15) a priori in virtue of knowing (13) and (14) a priori? That seems incredible by any reckoning. Unless one is willing to embrace a priori knowledge of deeply contingent propositions of this kind, (Transmission) better fail in our incompatibilist argument.²¹ That way neither privileged access nor the existence of natural-kind-dependent thoughts is jeopardized. Moreover, as this epistemic response hangs on no semantic externalist commitments, it applies equally

²¹ In contrast to superficially contingent propositions, Evans (1979) took the truth of a deeply contingent proposition to depend upon some contingent feature of reality.

well to the other exemplifications of the (MC) form sketched in Section I. Thus Wright (2000: 155-156; 2003: 65-67) claims that (Transmission) fails, and the explanation he offers is reminiscent of the one rehearsed in Section II. Note again that Wright's account is couched in terms of warrant, but since McKinsey (1991) and other incompatibilists speak of knowledge, we shall in the interest of simplicity continue to follow their lead. Now consider this alternative to (13):

- (16) *S* is on Dry Earth where she unawares fails to express the concept **water** when she utters sentences containing 'water'

Wright's claim (op. cit.) is then that since (13) and (16) are incompatible yet *S*'s state of knowing (13) a priori is subjectively indistinguishable from the belief state *S* would be in if instead (16) were true, *S* knows (13) a priori only if *S* has additional knowledge that (16) is false.²² Assuming that (16) would be true if (15) were false, *S* knows (13) a priori only if *S* knows (15) antecedently.²³ In other words, part of what makes *S* know (13) a priori is that *S* already knows (15). It would therefore be viciously circular to think that *S* could come to know (15) a priori on the basis of knowing (13) and (14) a priori and then competently reason in the way she does. This means that *S*'s a priori knowledge fails to transmit across the entailment in (14).

But Burge (1988: 654) and other semantic externalists are keen to stress that not only can *S* think thoughts involving wide concepts without knowing that the external conditions on thinking those thoughts obtain, *S* can also know a priori that she thinks those thoughts without knowing that such enabling conditions obtain. More precisely, *S* can know a priori that she thinks wide content thoughts without knowing that certain external conditions obtain, or indeed that they are conditions on thinking those thoughts. For instance, *S* can know a priori that she thinks **water is wet** involving the natural-kind-dependent concept **water**, and yet not know that water exists or even that thinking that wide content thought is conditional on the existence

²² In Section II Wright's template required that the state of knowing that *p* be subjectively indistinguishable from the belief state *S* would be in if instead *p** were true. Obviously, if (16) is true then *S* cannot even falsely believe (13). Instead the corresponding state in which (16) is true must be characterized as resembling the state *S* is in when she knows (13) a priori without assuming that *S* has the concept **water**.

²³ Note for the record that this assumption ignores the possibility that the closest worlds in which there is no water are worlds in which *S* is on Twin Earth.

of water. The conditions that must obtain if *S* is to think thoughts involving wide concepts must also obtain if *S* is to know a priori that she is thinking those thoughts. For *S* is thinking those thoughts in the very acts of thinking knowledgeably that she is thinking them. So, just as advocates of (PR) embrace the possibility of basic perceptual knowledge, semantic externalists endorse the possibility of:

(Basic Introspective Knowledge) For a range of content-bearing, mental properties *M* and external, content-individuating conditions *C* which are such that *C* must obtain if *S* is to instantiate *M*, *S* knows a priori that she has *M* even though *S* has no antecedent knowledge that *C* obtains

Now, Wright (2000: 172-173) is clear that the knowledge (or warrant) *S* has that (15) water exists antecedent to her knowledge that (13) she has the concept **water** must be a priori. For if *S*'s knowledge of (13) were conditional on her having empirical knowledge of (15), then her knowledge of (13) would be at best a posteriori. Wright (2003: 68) also insists that the a priori knowledge (or warrant) *S* has of (15) has a non-evidential character. *S* has an entitlement for (15) "...conferred not by positive evidence for [(15)] but by the operational necessity, so to speak, of proceeding on the basis of such so far untested assumptions if [*S*] is to proceed at all." In this case, *S* proceeds by deploying the concept **water** in thought and reasoning on the so far untested assumption that water exists in her external environment.

Reflect in response that (Basic Introspective Knowledge) places no special constraints on the character of the antecedent knowledge in question. In particular, *S* is permitted to have a priori knowledge that (13) she has the concept **water** despite lacking non-evidential a priori knowledge that (15) water exists. As far as the semantic externalist is concerned, whether the supposed antecedent knowledge is a priori or non-evidential is irrelevant. Burge (1988: 653-654; 2003: 264-265) does stress that in the act of thinking a thought involving a wide concept, *S* must presuppose the relevant external enabling conditions, but *S* need have no positive epistemic attitude towards them. The act of presupposing is not to be assimilated to a propositional attitude, let alone one that plays an epistemic role in justifying *S*'s claim to be thinking a wide-content thought. Burge (1993; 1996) argues that absent reason for doubt *S* enjoys an a priori, defeasible entitlement to rely on introspective

deliverances for belief-formation. The idea is that *S* has an epistemic entitlement to her beliefs about her own wide content thoughts that derives from their function in critical reasoning. These beliefs are an integral part of the overall procedures of critical reasoning that *S* engages in. Burge (1993: 459) thus concurs with Wright that *S* has immediate entitlements even when she cannot provide reasons or evidence for them: having an entitlement “does not require being able to justify reliance on these sources [e.g. introspection], or even to conceive such a justification.”²⁴ Indeed Burge would agree with Wright’s (2003: 60) claim that in a range of basic cases *S* has non-inferential knowledge of her own—wide content—thoughts. But the vital point is that Burge would eschew Wright’s claim that *S* can know a priori what she is thinking only if she has a positive epistemic entitlement to the external enabling conditions. Burge, but not Wright, can accommodate the possibility of (Basic Introspective Knowledge). This suggests that although Wright correctly blames the use of (Transmission) in the incompatibilist argument, the explanation he offers is not one of which Burge-style semantic externalists will want to avail themselves.

Here is a different tack. Recall our bootstrapping template from Section II:

(9) $K(p)$ in virtue of Fr

(10) $K(q)$

(11) $p \ \& \ q \Rightarrow Fr$

(12) $K(Fr)$

The claim is now that the (MC) form is just a *special case* where $q = (p \Rightarrow Fr)$, so that:

(17) $K(p \Rightarrow Fr)$

together with (9) and (12) give rise to viciously circular knowledge. There are, however, three important differences. (i) While *F* is an epistemically relevant property in the bootstrapping argument, *F* is a *semantically relevant property* in the incompatibilist argument. (ii) While *r* is a belief-producing source in the

²⁴ Note that while Burge conceives of entitlements as externalist warrants that need not be accessible or even understood, justification is a kind of internalist warrant involving accessible reasons. Both Burge and Wright use ‘warrant’ disjunctively to pick out either justification or entitlements.

bootstrapping argument, r is a *belief-individuating source* in the incompatibilist argument. (iii) While the in-virtue-of relation was taken to be epistemic in the bootstrapping argument, a *semantic relation* is needed in the incompatibilist argument. In order to avoid terminological confusion consider instead:

(9*) $K(p)$ by relying on Fr

where the *relying-on* relation is a semantic dependency relation. The amended claim is then that the incompatibilist (MC) argument fits the template comprising (9*), (17) and (12). That is to say, both the bootstrapping argument and the incompatibilist argument are such that one of the premises is known as a result of a source of knowledge instantiating a certain property in an argument with the conclusion that that source instantiates that property. This structural similarity between the two arguments stems from the fact that both fit the same template.

Let me explain. In Section II epistemic relevance was defined as follows:

(ER) If belief producing source r generates S 's belief that p , then F is an epistemically relevant property of r if and only if: had r not instantiated F , then S 's belief state, as produced by r , would not qualify as a knowledge state K with content p

The source that produces S 's belief that (13) she has the concept **water** is introspection. S knows (13) a priori in virtue of introspection instantiating various epistemically relevant properties such as being reliable and producing a true belief on this occasion. But these are not the properties to which the content of S 's putative knowledge that (15) water exists pertains. S supposedly acquires a priori knowledge that her external environment instantiates the property of containing water. According to the kind of semantic externalism needed to sustain the conditional in (14) that water exists if S has the concept **water**, S 's external environment is a belief-individuating source in the sense that the content of S 's belief that she has the concept **water** is at least in part individuated in terms of that property of her external environment.²⁵ Consequently, if belief states are individuated by their contents, then

²⁵ Remember to keep causation and individuation separate. On a given occasion, the presence of water in the vicinity of S may play a causal role in the formation of S 's belief that water is abundant, but to say that her belief is individuated in terms of her external environment containing water implies (on the current view) that S sustains a pattern of causal-historical interactions with water.

S 's state of believing that she has **water** is individuated in terms of her external environment containing water. Put differently, if the content that S has the concept **water** constitutively depends in part on whether her external environment contains water, then so does being in a belief state with that content.²⁶ Now *semantic relevance* can be defined counterfactually as follows:

(SR) If belief individuating source r individuates S 's belief that p , then F is a semantically relevant property of r if and only if: (i) had r not instantiated F , then content p would be unavailable as the content of S 's belief state, and (ii) were r to instantiate F , then p would be available as the content of S 's belief state.

Thus in our example, containing water counts as a semantically relevant property of S 's external environment by (SR): S could not be in the state of believing that she possesses the concept **water** if her external environment had failed to contain water. For if water were absent in S 's external environment, then she could not possess **water**, and so she could not be in a belief state with the content that she possesses that concept. But if water were present in that environment, then S could be in a belief state with the content that she possesses **water**. If S believes that she possesses **water**, then presumably that concept constitutes part of the content of that belief, and S must possess whatever concepts constitute the contents of her beliefs. Maybe the fact that the content of this particular belief-ascription refers to **water** rather than water means that S need not actually possess that concept, or at least imposes much weaker constraints on such concept-possession, but we could pick a different example, e.g. S believes that water is wet. It may also be that incomplete possession of **water** suffices for S to be in belief-states with contents constituted in part by that concept, but if S 's external environment had lacked water, then S could not even incompletely possess **water**.

²⁶ Thus Burge (2010) characterizes external individuation of content in terms of relations of constitutive dependency between the individual and her wider environment. Note also that talk about belief is notoriously ambiguous between the content of a belief and the state of having a belief with that content. Following Burge (1982), we shall hereafter take it that if the contents of belief are wide then so are the corresponding belief states themselves.

Containing water is, however, an epistemically irrelevant property by (ER), because *S*'s external environment is what instantiates that property yet is presumably not what generates *S*'s belief that she has the concept **water**. In this case, *S*'s external environment is a belief-individuating source rather than a belief-producing source. Still, *S* could not know that she possesses **water** if her external environment had lacked water. The reason being that if *S*'s external environment were unfavourable in the envisaged manner, no content composed in part by **water** would be available as the content of one of *S*'s knowledge states. Hence, if knowledge states are individuated at least in part by their contents in just the way belief states are, then *S* could not be in a knowledge state with the content that she possesses **water**. This means that just as we distinguished in Section II between belief- and knowledge-producing sources, we can distinguish between belief- and knowledge individuating sources. Consider then this take on *semantic relevance*:

(SR*) If knowledge individuating source *r* individuates *S*'s knowledge that *p*, then *F* is a semantically relevant* property of *r* if and only if: (i) had *r* not instantiated *F*, then content *p* would be unavailable as the content of *S*'s knowledge state, and (ii) were *r* to instantiate *F*, then *p* would be available as the content of *S*'s knowledge state.²⁷

In the incompatibilist argument, *S*'s external environment is a knowledge-individuating source that instantiates the semantically relevant* property of containing water. In contrast, in the bootstrapping argument, producing a true belief that the table is red is a semantically irrelevant* property, because *S*'s visual

²⁷ Two comments are in order. (i) Another reason *S* could not know that she possesses **water** if her external environment had lacked water is simply that she would then lack that concept. But the factivity of knowledge does not render the property of containing water semantically relevant*. Moreover, take *S*'s knowledge that her glass contains no water. Obviously, if *S*'s external environment had lacked water, then we can say that the content of that state would still be true. It would just not be the content of a state of knowledge. (ii) The assumption that knowledge states are individuated at least in part by their contents might arguably be challenged, but it would presumably still be true that knowledge entails belief, and so a notion of semantic relevance as pertaining to knowledge states could be fleshed out via content-individuation of the corresponding belief states.

perception is not a knowledge-individuating source that individuates S 's knowledge that the table is red. Indeed, as S 's visual perception is not a belief-individuating source either, that property is not even semantically relevant by (SR).

Let's now revisit our template. In our example, $p = S$ has the concept **water**, F = containing water, $r = S$'s external environment, $q = \text{if } (p) S \text{ has the concept } \mathbf{water}$, then (Fr) S 's external environment contains water. Our problem was that when S puts her introspective knowledge that (13) she possesses **water** together with her armchair knowledge that (14) water exists if she possesses that concept, then (Transmission) seems to furnish S with a priori knowledge that (15) water exists. But we can now explain in a way fully consistent with semantic externalist commitments why (Transmission) fails in the incompatibilist argument (13) – (15): knowledge of the first premise (13) relies on a source of knowledge-individuation r instantiating a semantically relevant* property F in an argument with the conclusion (15) that source r instantiates property F . The argument is thus afflicted by *semantic circularity*:

(SC) An argument is semantically circular if and only if (i) S knows at least one of the premises p , q by relying on knowledge-individuating source r instantiating semantically relevant* property F , and (ii) the conclusion that r has F , distinct from p , q , can be deduced from the conjunction of the premises p & q (where it is possible that $q = (p \Rightarrow Fr)$)

Note that (SC) resembles our notion of epistemic circularity (GEC) in that both render arguments circular just in case a premise is known in part as a result of a source of knowledge instantiating a given property in an argument with the conclusion that that source has that property. That is to say, arguments are circular when they import such a property into the content of a knowledge state that partially obtains as a result of that source having that property.

Now, the thought, corresponding to the one in the bootstrapping argument, is then that banning such semantically circular knowledge blocks the incompatibilist argument. More precisely, the proposal is that the semantically circular nature of the putative knowledge of (15) undercuts the epistemic support that (13) and (14) would otherwise provide for (15). That seems eminently plausible: if, as argued in Section II,

(GEC) undercuts the epistemic support that a set of premises would otherwise provide for a conclusion, then (SC) also constitutes a similar kind of undermining defeat. Accordingly, (SC) triggers failure of (Transmission): *S* cannot acquire knowledge of (15) by competently deducing (15) from (13) (and (14)) of which *S* already has water-dependent knowledge. Again, the dependency claim is important: the competent deduction of (15) from (13) and (14) would transmit knowledge had *S*'s knowledge of (13) been independent of whether her external environment instantiated the semantically relevant* property that features in (15).²⁸

IV. Conclusion

We have developed a template for diagnosing two distinct yet structurally related kinds of circularity such that any argument that meets this template, including our bootstrapping and incompatibilist arguments, trigger failure of (Transmission) due to the undermining nature of that circularity. The bootstrapping argument is afflicted by epistemic circularity: it illicitly imports epistemically relevant properties of knowledge-producing sources into the contents of knowledge states that obtain in virtue of those sources having those properties. The incompatibilist argument is afflicted by semantic circularity: it illicitly imports semantically relevant properties of knowledge-individuating sources into the contents of knowledge states that obtain by relying on those sources having those properties. We then argued that epistemic and semantic circularity undercut the epistemic support that the premises in those two arguments would otherwise provide for the conclusion so that knowledge fails to transmit across the relevant entailments. Importantly, our diagnosis is consonant with process reliabilism and other epistemological externalist views, which permit the possibility of basic perceptual knowledge. As these views are faced with the bootstrapping problem, their proponents can consistently adopt our template. Likewise, our diagnosis is consonant with a natural-kind-dependent thought view and other semantic externalist views, which allow for the possibility of basic introspective

²⁸ It should be pretty straightforward how to apply our template to the other exemplifications of the (MC) form that were contemplated in Section I. For instance, in order to handle architectualist arguments *S*'s internal cognitive architecture will count as a knowledge-individuating source, and the property of having a certain syntactic articulation might qualify as semantically relevant*. In general, while some such exemplifications pertain to a priori entailments outward away from *S*, others concern a priori entailments downward into *S*'s underlying cognitive machinery.

knowledge. As supporters of these views are confronted with the incompatibilist challenge, they can also safely embrace our template.²⁹

References

- Alston, William. 1986. 'Epistemic Circularity', *Philosophy and Phenomenological Research* 47, 1-30.
- Bergmann, Michael. 2004. 'Epistemic Circularity: Malignant and Benign', *Philosophy and Phenomenological Research* 69, 709-27.
- Boghossian, Paul. 1998. 'What the Externalist Can Know A priori', in *Knowing Our Own Minds*, (eds.) Wright, Crispin, Smith, Barry, & Macdonald, Cynthia, Oxford: Oxford University Press, 271-284.
- Brown, Jessica. 2003. 'The Reductio Argument and Transmission of Warrant', in *New Essays on Semantic Externalism, Scepticism, and Self-Knowledge*, Nuccetelli, Susana (ed.), Cambridge Mass.: MIT Press, 117-130.
- . 2004. *Anti-Individualism and Knowledge*. Cambridge: MIT Press.
- Brueckner, Anthony. 2013. 'Bootstrapping, Evidentialist Internalism and Rule Circularity', *Philosophical Studies* 164, 591-597.
- Burge, Tyler. 1988. 'Individualism and Self-Knowledge', *Journal of Philosophy* 85, 649 - 663.
- . 1993. 'Content Preservation', *The Philosophical Review*, 102, 457-488.
- . 1996. 'Our Entitlement to Self-Knowledge', *Proceedings of the Aristotelian Society, New Series* 96, 91-116.
- . 2003. 'Reply to Martin Davies, 'Externalism, Self-Knowledge and Transmission of Warrant'', in *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge*, (ed.) Frápolli, M. & Romero, E., Stanford: CSLI Publications, 261 - 268.
- . 2010. *Origins of Objectivity*. Oxford: Oxford University Press.

²⁹ I would like to thank an anonymous referee for this journal for extremely helpful comments on an earlier version of this paper, as well as [...].

- Cohen, Stuart. 2002. 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research* 65, 309-329.
- Davies, Martin. 2003a. The Problem of Armchair Knowledge', in *New Essays on Semantic Externalism and Self-Knowledge*, (ed.) Nuccetelli, Susana, Cambridge Mass.: MIT Press, 23-56.
- . 2003b. 'Externalism, Self-Knowledge and Transmission of Warrant', in *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge*, (ed.) Frápolli, M. & Romero, E., Stanford: CSLI Publications, 99-124.
- . 2008. 'Two purposes of arguing and two epistemic projects' in *Minds, Ethics, and Conditionals. Themes from the Philosophy of Frank Jackson*, (ed.) Ravenscroft, Ian. Oxford: Oxford University Press, 337-384.
- Evans, Gareth. 1979. 'Reference and Contingency', *The Monist* 62, 161-189.
- . 1982. *The Varieties of Reference*, Oxford: Oxford University Press.
- Korman, Daniel, Z. 2006. 'What Externalists Should Say About Dry Earth', *The Journal of Philosophy* 103, 503-520.
- Kornblith, Hilary. (2009) 'A reliabilist solution to the problem of promiscuous bootstrapping', *Analysis* 69, 263-267.
- McGinn, Colin. 1989. *Mental Content*, Oxford: Blackwell.
- McDowell, John. 1984. 'De Re Senses', *Philosophical Quarterly* 34: 283-294.
- McKinsey, Michael. 1991. 'Anti-Individualism and Privileged Access', *Analysis* 51, 9-12.
- . 2003. 'Transmission of Warrant and Closure of Warrant', in *New Essays on Semantic Externalism and Self-Knowledge*, (ed.) Nuccetelli, Susana, Cambridge Mass.: MIT Press, 97-116.
- McLaughlin, Brian. 2003. 'McKinsey's Challenge, Warrant Transmission, and

- Skepticism', in *New Essays on Semantic Externalism and Self-Knowledge*, (ed.) Nuccetelli, Susana, Cambridge Mass.: MIT Press, 79-96.
- McLaughlin, Brian & Tye, Michael. 1998. 'Is Content-Externalism Compatible with Privileged Access?', *Philosophical Review* 107, 349-380.
- Pryor, Jim. 2004. 'What's Wrong with Moore's Argument?', *Philosophical Issues* 14, 349-378.
- . 2007. 'What's Wrong with McKinsey-Style Reasoning', in Goldberg, Sanford (ed.), *Internalism and Externalism in Semantics and Epistemology*, Oxford: Oxford University Press, 177-200.
- Titelbaum, Michael. 2010. 'Tell me you love me: bootstrapping, externalism and no-lose epistemology', *Philosophical Studies* 149, 119-134.
- Tucker, Christopher. 2010. 'When Transmission Fails', *Philosophical Review* 119, 497-529.
- Vogel, Jonathan. 2000. 'Reliabilism Leveled', *Journal of Philosophy* 97, 602-623.
- . 2008. 'Epistemic Bootstrapping', *Journal of Philosophy* 105, 518-539.
- White, Roger. 2006. 'Problems for Dogmatism', *Philosophical Studies* 131, 525-557.
- Wright, Crispin. 2000. 'Cogency and Question-begging: Some Reflections on McKinsey's Paradox and Putnam's Proof', *Philosophical Issues* 10, 140-163.
- . 2002. '(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell', *Philosophy and Phenomenological Research* 65, 330-348.
- . 2003. 'Some Reflections on the Acquisition of Warrant by Inference', in *New Essays on Semantic Externalism, Scepticism, and Self-Knowledge*, Nuccetelli, Susana (ed.), Cambridge Mass.: MIT Pres, 57-77.

- . 2004. 'Warrant for Nothing (and Foundations for Free)?', *Aristotelian Society Supplementary* 78, 167-212.
- . 2007. 'Perils of Dogmatism', in *Themes from G. E. Moore. New Essays in Epistemology and Ethics*, (eds.) Nuccetelli, Susana and Seay, Gary. Oxford: Oxford University Press, 25-48.
- . 2011. 'Frictional coherentism? A comment on chapter 10 of Ernest Sosa's *Reflective Knowledge*', *Philosophical Studies* 153, 29–41.