

Externalism and Self-Knowledge: Content, Use, and Expression

Dorit Bar-On, UNC-Chapel Hill

1. Introduction

Suppose, as I stare at a glass in front of me, I say or think: There's water in the glass. The thought has the content *that there is water in the glass*; and it will be true if there is indeed water in the glass. According to external-world skepticism, even if my thought is true, I do not know that there is water in the glass, because my way of telling that this is the case does not allow me to rule out the possibility that I am only under some illusion. But surely I can know that I am **thinking** that there is water in the glass! Yet according to **external-content skepticism**, that is **not** something that I know, since my way of telling what I am thinking does not allow me to rule out the possibility that I am only under an illusion about the content of my thought.

Our confidence that we have secure basic knowledge of the contents of our thoughts and other propositional attitudes can be expressed as follows:

CSK: Our ordinary self-ascriptions of present contentful mental states (in speech or in thought) enjoy a special security relative to other contingent claims we make. Although they do not rely on observation, evidence, or inference, they typically represent genuine **and** privileged knowledge we have of the contents of the relevant states.

CSK is intended to capture the commonsense belief that, when we make ordinary self-ascriptions of contentful mental states, we have a special kind of **non-evidential knowledge** of the content of the mental states we are. That knowledge is privileged at least in the sense that each of us can have this special kind of secure knowledge, though only with respect to her **own present** states of **mind**.

Using a strategy analogous to the external-world skeptic's, the external-content skeptic may try to undermine CSK by presenting us with alternative contents that would be distinct in

reality but appear indistinguishable to us. Such alternatives are precisely what content externalism is supposed to have provided. A strong version of content-externalism says, roughly:

EXT: The contents of our thoughts (and other mental states) depend for their individuation on the nature of our physical or social environment. Whether an individual is in a mental state with one content rather than another depends in part on relationships between that individual and her extra-mental environment.

Thus, an individual cannot, e.g., harbor any *water*-thoughts if she has never been in any causal contact, direct or indirect, with the substance water. If, for example, she has only ever been in contact with a substance that looks, tastes, feels like water, but is not water – call this counterfeit substance “*twater*” – then she could not have *water*-thoughts, though she could have *twater*-thoughts.¹

Can CSK be maintained in the face of EXT? The content skeptic will argue otherwise.

Let us take the following Self-Ascription of a Contentful state

SAC: I am thinking that there’s water in the glass.

EXT entails that for SAC to be true, I must have had some kind of causal contact with water. I can, of course meet this condition without knowing that I have met it. However, if SAC itself is to constitute something I know, then it seems that I do have to ascertain that I have satisfied the relevant condition – I must have some **basis** for believing that my thought is indeed a *water*-thought, as opposed to some competing alternative. But how can I have such a basis without studying my relations to my extra-mental environment? Yet CSK implies that I **can** have knowledge – indeed, privileged knowledge – of what I am thinking – without such study.

Elsewhere, I argue that the content-skeptic’s challenge to ordinary knowledge of content depends on combining the possibility of **mental content counterfeits** – contents that are distinct in reality but may appear indistinguishable to a thinker – with what I call the **recognitional conception** of content self-knowledge. On the recognitional conception, though content self-knowledge (and self-knowledge more generally) is not based on any ordinary observation,

evidence or inference, it nevertheless does rely on our special ability to recognize characteristic features of our present mental states ‘from the inside’. Once this conception is added to the idea that, on any given occasion, things in the mental realm could be different from the way they appear to anyone, **including** oneself, then our ordinary **self**-knowledge should seem in principle as vulnerable to a form of skepticism that is analogous to external-world skepticism. Yet this seems to fly in the face of the commonsense conviction. If we want to adhere to the commonsense conviction, then we have two options. We can deny that there is a genuine appearance/reality gap in the mental realm, at least in the case of one’s own present states of mind. Or we can deny that ordinary self-knowledge is epistemically mediated by recognitional access, however secure.²

In this paper, I would like to canvass the second option, that of rejecting the recognitional conception. However, while rejecting the recognitional conception may help rebut content skepticism, the price may be a **deflationary** view of self-knowledge. Ordinary mental self-ascriptions are not open to epistemic doubt, says the deflationist, but they also do not represent genuine knowledge. In sections 2 and 3 below, I sketch a position on content self-knowledge that purports to meet the threat of external content skepticism without deflating self-knowledge. On the view I favor, ordinary self-ascriptions of contentful mental states do not rest on recognitional judgments. They enjoy a special kind of **immunity to error**, which explains why they seem invulnerable to content skepticism. However, this immunity to error is consistent with their being instances of genuine knowledge. So even if externalism is true, ordinary self-knowledge as represented by self-ascriptions of contentful states could withstand deflation. Now, however, a certain **inflation** of ordinary knowledge of content may seem in order. For if we combine content-externalism with the claim that we **do** have genuine knowledge of the content of our present mental states, it looks as though ordinary knowledge of content is very powerful indeed. From the simple, effortless knowledge of what we are thinking, combined with the knowledge that content-externalism is true, it seems that we can derive knowledge of the external world

determiners of our thought contents by pure reflection. But surely, external-world knowledge cannot be obtained in that way. In Section 4, I discuss the 'inflation argument' and explain why, once we reject the recognitional conception of content self-knowledge, the inflation of ordinary self-knowledge should be no more warranted than its deflation.

2. The Security of Self-Knowledge: Immunity to Error

Ordinary present-tense mental self-ascriptions that are not made on the basis of therapy, scientific findings, self-interpretation, or careful reflection – what are often called "avowals" – seem to enjoy a special security. As I ascribe to myself in the normal way, say, a desire for a glass of water, I seem to **know** that I want water. I do not seem vulnerable to certain skeptical doubts. My self-ascription seems secure as regards the **content** of my desire – its being **water**, rather than something else, that I want. It also seems secure as regards it being a desire that I feel for water, rather than just a thought (or some other attitude) involving water. Indeed, the security of content-assignment seems to be but a special case of a more general security we enjoy when making mental self-ascriptions (including non-intentional, or so-called 'phenomenal' self-ascriptions) in the normal way.³ At least as commonsense would have it, our avowals do not seem as readily vulnerable to skeptical hypotheses as our claims about the extra-mental world. What could explain this?

As mentioned earlier, it is often thought that if content externalism (EXT) is true, then the commonsense view cannot be right, since ordinary avowals **would** become vulnerable to skeptical hypotheses much in the same way as our ordinary perceptual reports are. But it can be argued that raising a skepticism about our ordinary knowledge of content on the model of external world skepticism requires conceiving of our knowledge of our own present mental states on the model of our knowledge of our extra-mental world. In particular, it requires thinking of our ordinary avowals of thoughts, desires, hopes, intentions, etc., as representing what may be called 'recognitional judgments'. When I judge in the normal way that there is a red car in front of my house, or that George is raking the leaves outside, I typically rely on the characteristic

recognizable features of various extra-mental things (red things, cars, George, the raking of leaves). So, if I discover that it is not a **red** car, or not a red **car** that is in front of my house, I can still adhere to the judgment that there is a car in front of my house, which only appears to be red, or that there is something red, which only appears to be a car, and so on. Similarly, it might be thought, when I self-ascribe a present mental state, I typically rely on characteristic recognizable features of my present states of mind. So, if I self-ascribe: I am hoping that there's some water in the glass, I may retreat to the judgment that I am in some state, though it only appears to be a state of believing, or that my state has some content, though it only appears to be the content that there's some water in the glass.

On this **recognitional conception**, my ability to **tell** knowledgeably that my present state has a certain content c , crucially involves (and depends on) my ability to **tell apart** states with content c from states with other contents that are available to me as a thinker through some internally accessible 'identifying marks' borne by mental states. Now, assuming that there is no **conceptual** guarantee that, as I examine my own mind, things must be as they appear to me to be, we may wonder: could there not be counterfeits of my present state of mind? Specifically, though my state's content appears to me to be c , could I not **really** be in a state with a different content, c' ? And, unless I am able to rule out these counterfeits, how can I **know** – and in an especially secure way, no less – that I am in state with content c ?

Notice that, so far, the skeptical worry is presented without any appeal to a particular view regarding the nature of the content of mental states. The worry arises, it seems, as soon as we combine the idea of a potential gap between appearances and reality in the mental realm with the idea that self-knowledge rests on recognitional access. The first idea introduces the possibility of **counterfeits contents** – contents that appear to a thinker other than what they are. The second invites the expectation that having self-knowledge is a matter of ruling out counterfeits by **recognizing** our states **as** having particular contents. Now, it is often assumed that **internalist** views of mind are somehow better placed than externalist views to answer the

‘internal world’ skeptic. But I submit that, if this is so, it cannot be **simply** because such views insist that mental states and their contents depend for their nature entirely on factors internal to the subject who is in the mental state. After all, we **can** be subject to illusions about various of our **non**-mental internal states whose presence and character we may also identify through characteristic features. (A seasoned diabetic patient might perhaps on an occasion mistake a change in blood pressure for a drop in the level of insulin in her blood, for example.) In order to be able to offer an adequate response to the ‘internal world’ skeptic, an internalist would have to explain what it is about internal **mental** goings-on that, in our own case, affords us knowledge of them, which, furthermore, does not seem as vulnerable to skepticism as our knowledge of contingent extra-mental goings-on. As long as the internalist adheres to a conception of self-knowledge that is both objectivist and recognitional, it is not clear what explanation she can offer for this commonsense view.⁴

If, in telling what state of mind we are in at present we depend on recognizing characteristic features of our internal states, then it seems that skeptical alternatives abound. If ordinary self-knowledge requires ruling out these alternatives, then it would seem that its privileged status must be illusory. Perhaps, though, it is the recognitional conception that we should let go. I think (though I will not argue this here) that the recognitional conception misrepresents the character and status of ordinary self-knowledge. However, recognition (broadly understood) may seem to be our best model for ‘non-theoretical’ knowledgeable telling regarding empirical matters of fact. And if we are to eschew the requirement that a subject must be able to tell apart recognitionally the content a present state of hers has from other candidates, we may wonder what would allow the subject’s self-ascriptions of content to constitute a kind of knowledge at all, let alone an especially secure kind of knowledge. Thus, we must face the **deflationist** challenge:

If ordinary self-ascriptions of present mental states are not open to doubt in view of alternatives, this is not at all because they represent an especially secure form of knowledge, but because they represent **no** knowledge at all.⁵

The deflationist challenge is best met by offering a view of what allows ordinary mental self-ascriptions to be both especially secure and represent knowledge, even though they do not rest on recognitional determination. I cannot here develop such a view in detail. However, in what follows I would like to outline the non-recognitional view I favor, highlighting those features that I think would enable it to meet the charge of deflation, as well and the charge of inflation (which will be addressed in section 4).⁶

I propose that we understand the knowledge I have of what content a present mental state of mine has on the model of the knowledge I (often) have that it is **I myself** who am instantiating a certain property. In "Self-Reference and Self-Awareness," Sydney Shoemaker makes an observation regarding certain self-ascriptions that use a first-person device such as 'I':

The statement "I feel pain" is not subject to error through misidentification: it cannot happen that I am mistaken in saying "I feel pain" because, although I do know of someone that he feels pain, I am mistaken in thinking that person to be myself. ...

If I say "I feel pain" or "I see a canary," I may be identifying for someone else the person of whom I am saying that he feels pain or sees a canary. But there is also a sense in which my reference does not involve an identification. My use of the word "I" as the subject of my statement is not due to my having identified as myself something of which I know, or believe, or wish to say, that the predicate of my statement applies to it.⁷

Shoemaker offers the notion of **immunity to error through misidentification** to characterize the special epistemic status of self-ascriptions of the kind he describes in these passages (a notion also discussed by Gareth Evans and others⁸).

The following features of judgments that are immune to error through misidentification (IETM, for short) are crucial for our purposes. First, such judgments are not in general

incorrigible. "I see a canary," "I am sitting down," "My legs are crossed," can all be plausibly said to be IETM, yet no one would be tempted to take them as incorrigible. Furthermore, the class of judgments that are IETM includes **nonmental** self-ascriptions as well as judgments that are not self-ascriptions (both Shoemaker and Evans think that demonstrative thoughts such as "This is moving fast" are IETM). Second, when an ascription is IETM, then, even if it can be mistaken, it is not open to a certain **kind** of error: an error of mistaking one object for another because of a mistaken identification. Third, whether or not an ascription is IETM depends not on its semantic content, but rather on its "epistemic pedigree". One and the same (semantically individuated) ascription can be IETM or not, depending on the **basis** on which it is made. On occasion, I might tell that I am sitting down by seeing in a mirror, or on a video screen, someone who is sitting down, and whom I take to be me (perhaps mistakenly). In that case, my self-ascription will **not** be IETM, even though it is IETM when it is made in the normal way. Finally, if I issue a self-ascription "I am F" that is IETM, I can be legitimately said to **know that it is I myself** who has the property F. However, this knowledge I have is not based on recognition of myself as the 'right' subject of my ascription.

When a self-ascription of the form "I am F" is IETM then, although my self-ascription may be false, there is no room for me to think: **Someone** is F, but is it **me**? If I come to think **in the usual way**: "I have a toothache," or "I'm standing in front of a tree," or "I see a canary," there is no room for me to doubt whether it is **me** who has the relevant properties without doubting that **someone** has them. This is because in such cases I have no grounds for thinking that **someone** has the relevant properties over and above, or separately from, any grounds I might have for thinking that **I** have them. With respect to many self-ascriptions I do have such independent grounds. Often, I tell how much money I have in my bank account, for example, by consulting the bank teller's screen. I receive information that **someone** has (say) \$500 in her account, and taking that information to be about **my** account, I come to think that it is **me** who has \$500. If in

fact the information is about someone else's account, I would be right to think that **someone** has \$500 in her account, though wrong to think that **I** have \$500 in my account.

As I said, being IETM is only being immune to a certain **kind** of error – an error about **who** is F. But this immunity does not reflect a special success in identifying or singling out recognitionally the right individual of whom to predicate F. Quite the opposite. Immunity to error through misidentification reflects the **absence** of recognitional identification. When I say or think in the ordinary way: "I see a canary," the reason I cannot go wrong in identifying myself is that my ascription in no way relies on a separate recognition that **someone** sees a canary coupled with a judgment I make that that someone is me, based on characteristic features I observe. Correlatively, if I can be said to have knowledge that it is me who sees a canary, this knowledge should not be seen to depend on the correctness and justification of the independent judgments that **someone** sees a canary **and** that that someone is me. In the normal case, my only epistemic route to the belief (or judgment) that someone sees a canary is via my belief that **I** see a canary.⁹

Neither Shoemaker nor Evans are inclined to deny that in such a case I have **knowledge** that **I** see a canary. Although I do not come by the self-ascription via the recognition that it is **me** as opposed to someone else, who is sitting down, the self-ascription still represents something I know about myself. How can that be? On Evans' analysis, self-ascriptions that are IETM represent knowledge that we gain about ourselves in a distinctive way. Evans notes that we possess two general capacities for gaining information about some of **our own** states and properties. First, we possess "a general capacity to perceive our own bodies" (which includes "our proprioceptive sense, our sense of balance, of heat and cold, and pressure"), and second we also have a capacity for determining our own "position, orientation, and relation to other objects in the world ... upon the basis of our perceptions of the world."¹⁰ It is the exercise of these capacities that gives rise to self-judgments that are IETM. For, when a subject gains information that I am F (for the relevant range of F's) in one of these ways

[t]here just does not appear to be a gap between the subject's having information (or appearing to have information), in the appropriate way, that the property of being F is instantiated, and his having information (or appearing to have information) that **he** is F; for him to have, or to appear to have, the information that the property is instantiated just is for it to appear to him that **he** is F.¹¹

When it comes to self-judgments that are IETM, Evans is suggesting, the only reason I have for thinking that someone is F is that I think that I am F. This is why in such cases it will make no sense to wonder: "Someone may be F, but is it I who am F?". In these cases, there is no way for me to doubt that **I** am the one who is F without thereby doubting that **someone** around is F. To put it somewhat awkwardly, in these cases, there is no stable judgment of appearances concerning who is F that I can hold onto so as to raise a doubt about how things really are. I can, of course, be wrong that I am F. But Evans denies that this by itself is sufficient to show that my grounds for thinking that I am F must factor out into the grounds I have for thinking that someone is F and the grounds I have for thinking that I am that someone. By the same token, he would deny that I could only be **entitled** to the belief that **I** am F, if I have separate grounds for the judgment that I myself am the individual who is said to be F.

Now what about counterfeits? To vary the example, suppose I say or think: "I am sitting down." Couldn't it be, for all I know, that it is my identical twin who is sitting down? Well, that possibility would seem entirely irrelevant in this case, given that I do not tell that it is me who is sitting down by telling myself apart from other potential candidates. Since in the normal case, my self-ascription does not rest on any recognitional judgment identifying me as the individual who is sitting down, there seems to be no room for doubt on my part as to whether it is someone else, not me, who is sitting down. We could say, if we like, that the possibility that it is my identical twin is the one sitting down is in this case not a 'relevant alternative.' What is crucial to observe, however, is that what allows this alternative to be ruled out of court is a negative feature of the epistemic situation: the fact that the judgment is **not** made on a certain kind of basis. The

judgment does represent a piece of knowledge I have about myself, viz., that **I** am sitting down. But my knowledge does not depend on my having separate grounds for the judgment that the person (whom I take to be) sitting down is me, in virtue of my recognition of some characteristic features of the relevant individual. If I am indeed sitting down, my judgment will be true, and I can be said to be entitled to it, at least inasmuch as defeating alternatives are not in the running.

If we accept Evans' analysis, then we have an interesting and powerful model of epistemically secure self-ascriptions, which I think could serve us well in understanding the epistemic status of ordinary content self-ascriptions (and mental self-ascriptions more generally). On this model, there are self-ascriptions that exhibit the following features:

- (i) Though not absolutely infallible, they are **immune to a certain kind of error and doubt** (concerning the identity of the subject of ascription)
- (ii) They can represent **knowledge** that something is the case (viz., that **I** am F)
- (iii) The knowledge is **non-recognitional**, in the sense that it does not require one to have a recognitional basis, or grounds for the relevant judgment (the judgment that I am F does not rest on a recognitional identification of myself as the one who is thought to be F)

Let us now turn to the case of present-tense of contentful mental self-ascriptions (avowals, as I referred to them earlier). We have seen that an "I"-ascription that is IETM can represent knowledge I possess concerning the subject component of the ascription, namely, that I myself am F, even though I do not make any recognitional identification of myself as the one who is F. This is an idea I want to apply to the **ascriptive** component of avowals (i.e., the component in which I ascribe a present mental state to myself). Specifically, my suggestion is this:

When making a mental self-ascription of the form: "I (am) M(ing) (*that*) *c*" in the ordinary way, my self-ascription may constitute knowledge I have that *c* is my state's content, even though the content is not for me a **recognitional target**.

Suppose I avow: "I am nervous about this dog," or "I am hoping that you'll join us tonight." My self-ascription does not rest on a recognitional judgment to the effect that an internal state of mine (which I recognize as being of a certain kind) has some content, which content I take to be *c*, on the basis of the way it appears to my "inner eye" (or my internal scanner). If the self-ascription represents my knowledge that I am in state *M* with content *c*, this will not be knowledge that rests on any separate recognitional judgment to the effect that it is *c* (rather than *c'*) that is the content of my state. For, I do not have reason to think that my state has some content or other that is separate from, or independent of, my reason for thinking that I am in a state with content *c*. (The analogous claim in the case of judgments that are IETM was that I do not have reason to think that someone is *F* that is separate from or independent of my reason for thinking that I myself am *F*.)

Content ascriptions to others do typically rest on independent recognitional judgments. I can tell by the way you behave that you are scared of **something**, and I look around to figure out what it is you are scared **of**. I conjecture, based on the direction of your gaze, and other signs, that it is the dog you are scared of. Still, I can sensibly wonder whether it is the cat, not the dog, that you are scared of, while not questioning that you are scared of **something**. (I can, of course, also question whether you are scared of anything at all.) Even content **self**-ascriptions can rely on recognitional judgments. I can sometimes determine that **it is my neighbor** that I am annoyed at through observation of my own behavior (how I react to her presence, the tone of voice I use, etc.). I can also reach this conclusion after reviewing the circumstances, consulting with friends, or submitting to costly therapy. But this is not my **normal** way of telling that I am annoyed **at my neighbor**. I submit that, when avowing being annoyed at my neighbor, I do not rely on "internal" recognition of the intentional object of my annoyance, either. I do not normally tell who it is I am annoyed at by recognizing some characteristic features of my internal state.¹²

Avowals of contentful states are among the paradigm examples of ascriptions that are IETM; they are immune to error regarding **who** it is that is in the relevant contentful state. I want

to suggest that, in addition, avowals are also **immune to error through content misascription**. When assigning content c to a present state of mine in the ordinary way, I am not subject to a recognitional error as to what content my state has, precisely because my content self-ascription is not the upshot of (nor does it rest on) my judgment that I have a state with **some** content, coupled with an independent recognition that **that content is c** , rather than c' . If my self-ascription is false, this will not be due to a mistaken recognitional judgment on my part – due to my failure to identify correctly the content based on some characteristic recognizable features. Correlatively, if my self-ascription is true, my knowledge that I am in a state with a certain intentional content c will in no way rest on my having recognitional grounds for judging that it is c (and not c') that is the "right" content of my state. (The analogous point about self-ascriptions that are IETM was that, when they are false, this is not due to an error of misidentifying the ascription's subject. Correlatively, when they are true, the subject's knowledge that she is F in no way depends on her having separate grounds for judging that it is she, rather than someone else, who is F.)

It is worth highlighting the fact that the present proposal pertains to the **epistemology** of avowals, not their semantics. If I avow: "I hope dinner will be served soon," my avowal has the semantic content that I hope that dinner will be served soon. However, epistemically speaking, it is not grounded in a recognitional identification of my state or its content (or of the state's bearer). Just as it would be a mistake to regard "I am F" as epistemically based on "b is F" and "I am b", so I think it would be a mistake to regard "I am in state M with content c " as epistemically based on "I am in some state with some content" and "That state is M, and its content is c ." When ordinarily saying, or thinking: "I am in state M with content c ," I do genuinely ascribe to myself a specific kind of state with a specific content. But my avowal is immune to errors of misascription (as well as misidentification), since it does not rest on recognitional judgments concerning the state or its content. And I am suggesting that just as an "I"-ascription of the form "I am F" that is immune to error through misidentification can represent my knowledge that I am F, so an "I"-

ascriptions of the form “I am in state M with content *c*” that is immune to errors of misascription (i.e., an avowal) can represent my knowledge that I am now in a state with that content.¹³

3. The Security of Self-Knowledge: A Neo Expressivist Account

We saw that, on Evans' analysis, although self-ascriptions that are IETM do not rest on recognitional identification of their subject, they are nonetheless "epistemically underwritten," in that they are obtained through distinctive capacities for gaining information about some of our own properties and states. A **true** self-ascription such as “I am sitting down,” when made in the ordinary way, conveys information about **me** which is obtained through my putting to use these distinctive capacities. This is presumably part of what, on Evans' view, allows me to know that **I** am F, even in the absence of a recognitional identification of myself. If we are to insist, as against the deflationary view, that avowals can represent knowledge we have about our mental states and their contents, we would do well to find an appropriate analogue for this feature in the case of avowals.

So our question is this: if avowals do not rest on recognitional judgments concerning the content of one's present state, what is it that allows avowals to represent knowledge that we are in a state with a particular content? My proposed answer is that, here too, we can appeal to a distinctive capacity that we deploy when issuing avowals. The proposal is as follows:

When ascribing to myself a state with content *c* in the normal way, I exercise an **expressive** capacity: the capacity to **use** content *c* (rather than content *c'*) in the course of expressing, or giving voice to, my present state. When avowing, say, a present thought that is crossing my mind, I tell what content the thought has by telling **it**, i.e., that content. If I "identify" the content *c*, it is not by recognizing *c* to be the content I take my thought to have, but rather by using **that** content to articulate my thought.

I shall explain this expressivist proposal in connection with the case of content self-ascription with which we began, and then indicate briefly how it can be generalized.

Thus, let us return to

SAC: I am thinking that there is water in the glass.

Understood purely as a self-ascription of a presently entertained thought, SAC is maximally secure. For, when making it, I cannot misidentify the subject of my ascription, but I **also** cannot misascribe to myself a specific thought that there is water in the glass. Indeed, on a natural reading of self-ascriptions of this special kind, i.e., self-ascriptions of presently entertained thoughts, they are **self-verifying** – they make themselves true.¹⁴ Now contrast SAC with: "I am thinking something boring," which clearly does not seem self-verifying. What makes "I am thinking that there is water in the glass" self-verifying is the fact that the very act of self-ascribing the thought summons up the ascribed thought, as it were; it inevitably involves the thought "passing through" my mind, thereby rendering my ascription true. For, in ascribing the entertained thought, I directly spell it out. Now, the use of explicit articulation as a way of 'latching onto' content is not restricted to avowals. We commonly assign intentional content to others' intentional states by displaying, or spelling out the assigned content. (Contrast "He is thinking about what I was just thinking" with "He is thinking about the dinner he's about to miss.") However, in the case of ascribing content to another's state through articulating it, one still relies on a recognitional judgment to the effect that the other's state has that content. One's assignment of content to the other's state relies on recognition of characteristic features exhibited by the others' behavior, for example. Not so in the case of present-tense self-ascriptions. (We return to this below.)

Self-verifying avowals such as "I am thinking that there's water in the glass" present a special limit case. When the state ascribed is simply one of entertaining a certain thought, all I need to do in order to succeed in **truly** ascribing to myself the entertaining of a thought with content *c* is to entertain a thought with that content, which I am bound to do if I ascribe the thought explicitly. Barring slips of tongue (or their analogues in thought, if there are such), as I articulate the content, I cannot but entertain the very thought I ascribe to myself, which is why my avowal is self-verifying. This is not true of all avowals of contentful states. However, although

avowals of contentful states are not in general self-verifying, they do typically articulate their content. When avowing, say, the **hope** that there's water in the glass, I do not obliquely describe my hope as when I report: "I am hoping for the same thing you are hoping for." Rather, I explicitly articulate the content *that there is water in the glass*, in the course of giving voice to my wish. The explicit articulation of the state's content does not represent an independent assessment of a recognizable feature of my first-order state. Rather, it is the upshot of my putting to use the very same resources I would put to use in saying, or thinking "If only there was water in the glass" **hopefully**. Although, semantically speaking, I am **ascribing** a state with this content to myself, epistemically speaking, I am doing no more than **using this** content to give articulate voice to my state, which is what I sometimes do when I say (or think): "If only there was water in the glass!" Thus, I maintain, when I **avow** a state, as opposed to ascribing it to someone else, or even ascribing it to myself indirectly, on some specific epistemic basis, explicit articulation obviates the need for assigning the content on a recognitional basis.¹⁵

A self-ascription that is IETM, such as "I am sitting down," is **about** a particular individual in the world, me. Semantically speaking, it **specifies** a particular individual who is said to be sitting down. However, if such a self-ascription represents my knowledge that **I** am sitting down, this is not because it rests on a recognitional identification of myself as the one who is sitting down. Rather, it is (partly) because in issuing the self-ascription I put to use certain capacities for gaining information, which I have only with respect to certain of my own states. I am making an analogous claim regarding self-ascriptions that are, as I put it, immune to content misascription. An avowal such as "I am hoping that there's water in the glass" ascribes to me a present state with a particular content. Semantically speaking, it specifies a state I am said to be in as well as a content that state is said to have. However, it does not rest on my recognitional judgment to the effect that my state has that content. If it represents my knowledge that what I am hoping is that there's water in the glass, this is (partly) because in issuing the self-ascription, I put to use a distinctive capacity that I have only with respect to certain of my own states. The

capacity in question is one I put to use when **avowing** my mental states, as opposed to reporting my findings about my present mental conditions. It is a capacity that only I have with respect to my mental states: the expressive capacity to give articulate voice to present states of mine, or to **speak from** them.

It is important to note the following difference between self-ascriptions that are merely IETM and self-ascriptions that are, in addition, immune to errors of misascription. A self-ascription that is IETM still rests on a recognitional judgment concerning one's present bodily state. It's just that the self-ascription is not based on a recognitional judgment concerning the identity of the **bearer** of the state (viz., oneself). By contrast, on my proposal, avowals of one's present mental state do not rest on recognitional judgments concerning one's present state either. Here my concern is specifically with content: I am claiming that an avowal of a contentful state does not rest on a separate judgment regarding **what** content one's state has which is based on one's 'internal' recognition of some characteristic features exhibited by the state. This is what renders it immune to errors of content misascription as well as misidentification. On my proposed **neo-expressivist** account, when I avow a present hope, or wish, or belief, my avowal is secure not because it represents an especially secure belief—a belief that is grounded in the special recognitional access I have to the ascribed state. I do issue a genuine self-ascription of a state with a particular content, but what I express in so doing is the ascribed mental state itself; I use the particular content in the course of directly giving voice to my hope, wish, or belief.¹⁶ And this is what renders my avowal especially secure.

But if this is so, one may worry that, even if avowals are not vulnerable to "internal-world" skepticism, because immune to errors of misascription as well as misidentification, avowals are poor candidates for genuine knowledge in the first place, in which case the deflationist challenge has not been met. The present worry can be perspicuously couched in terms of the familiar conception of knowledge *that p* as requiring (minimally) a **warranted belief that p**. We can see the deflationist about self-knowledge as someone who maintains that ordinary

self-knowledge (so-called) fails to meet one or the other minimal conditions on possessing knowledge in general. Given the neo-expressivist account I have sketched here, the deflationist may argue, first, that avowals will fail to represent **beliefs** that subjects have regarding their present contentful states. And, second, even if avowals do represent self-beliefs on this account, they cannot represent **warranted** beliefs.

I can here only offer a few brief remarks.¹⁷ Let us consider the first worry first. The worry is that since, on the neo-expressivist account, avowals serve directly to express the mental state that the subject ascribes to herself, they cannot express the avower's higher-order belief or judgment that she is in the mental state. But notice that this presupposes that an avowal's expressing one mental state (say, the avower's hope that *p*) excludes its expressing another mental state (the avower's belief that she hopes that *p*). This exclusivity claim can be denied. Perhaps an avowal expresses **both** the state ascribed and the avower's judgment that she is in that state, though, of course, that judgment is not recognitionally based. But even if we accept the exclusivity claim, and adhere to the 'belief' requirement on knowledge, we need not accept that avowing subjects cannot meet this requirement. Even if, when I avow being in some mental state *M* I do not express my higher-order judgment that I am in the state, it may still be correct to say that I do believe that I am in the ascribed state. Semantically speaking, my avowal **says that** I am in *M*. So, even if when issuing it I do not express a judgment I entertain at that moment that I am in *M*, I can still be said to **hold-true** that I am in *M*. I see no reason to deny that believing in the sense of holding-true is sufficient for satisfying the 'belief' requirement on knowledge.

What about the other requirement on knowledge? It seems fair to say that, on the neo-expressivist account, when I avow, I do not have **justification** for holding-true the self-ascription I issue, at least not in the traditional sense. For, my self-ascription is not based on any inference, evidence, or observation, and, on the neo-expressivist account, it is not grounded in any 'internal recognition' either. However, this does not mean that I cannot still be warranted in my belief, in the sense of being **entitled** to hold-true the self-ascription. The model of judgments that are

immune to error through misidentification may suggest the following possibility. A self-ascription arrived at through the use of a distinctive, non-recognitional capacity – e.g., “My legs are crossed” – can reliably present us with information about some matter – its being **me** whose legs are crossed. As I make the self-ascription, I can plausibly be said to hold-true that it is me whose legs are crossed. Moreover, given that my self-ascription is in no way epistemically grounded in any judgment identifying me with the individual whose legs I think are crossed, no defeating alternatives to its being **me** whose legs are crossed would seem to be in the running. In such a case, we could say that I am **entitled** to hold-true the relevant judgment (viz., that **my** legs are crossed) **by default**. The absence of defeating alternatives can be regarded as the epistemic cash-value of the judgment’s immunity to error, which, in turn, is a consequence of the epistemic character of the judgment (as opposed to its semantic content). If a self-ascription such as “My legs are crossed” can represent my knowledge that **my** legs are crossed in this way, then I think we can similarly see how an avowal such as “I am hoping there’s water in the glass” can represent my knowledge that **what** I am hoping is that there’s water in the glass.

This is not the place to lay out the full positive story of what allows so-called ordinary self-knowledge to constitute a species of genuine knowledge that is non-recognitional yet privileged. But I believe that enough has been said so far in support of the commonsense intuition that ordinary content self-knowledge as captured by CSK is not readily threatened by a form of skepticism that appeals to content externalism as expressed by EXT. Furthermore, by offering a characterization of the items of basic self-knowledge as genuine self-ascriptions that enjoy a special kind of epistemic immunity to error, I have shown how averting the threat of content skepticism **need not** land us in a **deflation** of ordinary content self-knowledge. For, I have suggested one way to see ordinary assignments of content to our mental states as representing genuine self-knowledge even though they are protected from skeptical alternatives.

4. Could I Know There Is Water Around, Just By Knowing What I'm Thinking?

Now, however, we may have to face what I will call “the problem of **inflation**.” If it is true that I can have privileged, ‘effortless’ knowledge of the content of my thoughts, and if, moreover, I know that externalism is true, then it seems that I could obtain specific information about the external world surrounding me by simply consulting the contents of my present thoughts. Thus, consider the following simple argument (henceforth referred to as "ABC"):

(A) I am thinking that there's water in my glass.

(B) If I am thinking that there's water in my glass, then there is water in my environment.

(C) So, there is water in my environment.

This is, of course, a valid argument. Moreover, (A) simply ascribes to me an episode of thinking in which I might engage, and the ascription may well be true on some occasion, and (B) is presumably true, if externalism is. McKinsey, Brown, and Boghossian (MBB henceforth) think that a problem arises, once we recognize that premises (A) and (B) are both available to me without any empirical study of my environment.¹⁸ We can agree that my ordinary knowledge of what I am thinking does not depend on my examining my relations to my external environment. And MBB think that, inasmuch as externalism is established through philosophical thought experiments, I can know (B) also without empirical investigation of my environment. So it looks as though I can come to know (C) purely through inference from something I know by 'effortless' self-knowledge, coupled with armchair philosophical reflection. If this seems absurd, then we may again think that either content-externalism (EXT) or the doctrine of content self-knowledge (CSK) must go.

Before explaining where I think the above reasoning goes wrong, let me make two preliminary remarks. First, I have said that ABC is incontestably valid; but its validity, of course, does not guarantee the validity of MBB's inflation argument, which is not about *water*-thoughts and water, but about various things we can be said to **know** and about how we know them. Knowledge is notoriously intensional. One can know something under one description but not

under another. I can know that I am talking to Clark Kent without knowing that I am talking to Superman. And I can know that my ring is made of gold, without knowing that my ring is made of a metal with atomic number 79. Since the inflation argument appeals to knowledge claims, we will have to pay careful attention to the way we construe those claims. We will need to be clear on what exactly we are said to know when we know A, B, and C. For example, if I know in the ordinary way that I am thinking that there's water in the glass, do I thereby necessarily know that I am thinking a thought that is causally linked to water **as opposed to twater**? If I know by inference from A and B that there's water in the glass, does that mean I know in that way that there is a liquid with the chemically distinguishing marks of water (i.e., the chemical structure H₂O, as opposed to XYZ) in the glass?

Second, in discussing MBB's reasoning, I will be conceding something I do not quite believe, namely, that content externalism, as a general metaphysical thesis about the nature of content, can be established through pure philosophical reflection, unaided by any empirical assumptions or information.¹⁹ We should note, however, that the general thesis only tells us that at least some of our thoughts depend for their individuation on the presence of substances with specific underlying natures in our environment. Yet even if we accept that the general thesis of content externalism can be established through armchair reflection, it does **not** follow that knowledge of specific externalist conditions can be obtained in this way. (The significance of this will emerge later.)

Let us begin by considering an argument that is similar in form to ABC:

ARGUMENT I There's water in my glass.
 If there's water in my glass, then there's H₂O in my glass.
 So, there's H₂O in my glass.

I know the first premise of this argument through unaided observation. Furthermore, I can know it **without** knowing anything about H₂O. On the other hand, I can know the second premise by inference from chemical theory. But, since the conclusion logically follows from the premises, I

can know the conclusion by inference. This means I can have knowledge that there is H₂O in my glass without conducting any specific tests of the sample of liquid in front of me. But my ability to learn that there is H₂O in my glass through inference does not impugn the status of “There is H₂O in my glass” as a piece of empirical knowledge. Thus, Argument (I) illustrates the fact that what is indisputably a piece of empirical knowledge **can** be arrived at purely on the basis of a valid argument, even when the canonical way of coming to know it is different in kind from the way either of the premises is known. So, in general, there is no absurdity in the possibility of coming to know something that is normally known through an empirical study of a certain kind, by inference from premises that are not known in that way.

However, you may protest that there is an important disanalogy between ABC and Argument I. Even if the premises and conclusion of Argument I are known in different ways, these different ways all involve empirical observation or study. So, if I were to come to know that my glass has H₂O in it on the basis of knowledge of the premises of Argument I, this would not show that I can achieve a piece of empirical knowledge about the world through logical inference from non-empirical knowledge. What seems disturbing about ABC, in contrast, is that it can be apparently used to show that one can know contingent facts about the external world just by putting pure philosophical reflection on the nature of thought together with a bit of ordinary knowledge of what we’re presently thinking about. So consider next

ARGUMENT II This man is George.

 If this man is George, then he is of biological origin B.

 So, this man is of biological origin B.

The first premise of argument II will be typically known in a non-theoretical way, through acquaintance, whereas its conclusion is not typically knowable in that way; it requires empirical investigation. What about the second premise? The second premise of Argument II, like that of Argument I, relies on a general background theory (biological theory). In this case it may be even clearer than in the case of Argument I that **knowing** the second premise requires more than

pure inference from the relevant theory. It requires specific study of the relevant empirical domain. Argument II illustrates, further, that in some cases, to obtain a conditional premise that is knowably true, and which could carry us to the relevant conclusion by Modus Ponens, we have to conduct the same kind of empirical investigation that is normally required to establish the conclusion. All that we can infer from biological theory alone is that George, being a man, must have a specific biological origin. But we cannot know what is biologically essential to being George without studying the biological facts about George and his history. If we stuck with what we could infer from biological theory, we could only legitimately use as a second premise

If this man is George, then he is of some specific biological origin.

From this and the first premise of Argument II all we could infer is

This man is of some specific biological origin.

Once we see precisely what it is that can be known through knowledge of biological theory alone, we would not be tempted to use Argument II to show that knowledge of biological theory alone (plus everyday acquaintance) can teach us surprising things about our friends. All we are in a position to learn by logical inference from biological theory with respect to people of our acquaintance is that there is a specific biological origin that is essential to them. But this much knowledge is already secured by our acceptance of biology as a general theory that is applicable to all human beings.

It might still be pointed out that neither premise of Argument II (even on the weakened version) is knowable without empirical study; this, again, makes for an apparent difference between Argument II and ABC. So now consider the following argument:

ARGUMENT III I am feeling nervous right now.

If I am feeling nervous right now, then my brain is in functional state N.

My brain is in the functional state N.

As in the case of ABC, the first premise of Argument III represents an instance of ordinary self-knowledge. To know it, I do not have to study my extra-mental world; I do not need to study my

brain or my internal functional organization, either. (I may not even realize I have a brain!) And, very much like ABC's second premise, the second premise of Argument III represents application of a general philosophical theory of the nature of mental states to a particular type of case. But assessing Argument III, along the lines suggested by our assessment of Arguments I and II will help us see something that ABC obscures.

Suppose we accept (implausibly, in my opinion) that functionalism, as a metaphysical theory about the nature of mental states can be established through pure, empirically innocent philosophical reflection. It still does not seem true that **particular** functionalist claims, about what is required for being in a particular type of mental state, can be established without empirical investigation. So if we insist that the second premise of Argument III **is** to be knowable without such investigation, we are at most entitled to a rather weak understanding of that premise. For the most we could derive from knowledge of the general thesis of functionalism alone, with no further empirical study, is that, if I am in some particular mental state that I ordinarily describe as being nervous then my brain is in some appropriate functional state, about which I cannot say much of anything, except that it is whatever state fits the functionalist profile for being nervous. In other words, we are not entitled to understand "N" in the second premise as providing a substantive characterization of the relevant functional state in informative terms (by which I mean terms that are not available to us simply in virtue of our ordinary understanding of the term "nervous"). But then it can be seen that there isn't here a piece of substantive empirical knowledge that I can be said to come by through the use of a logical argument alone. Or, if there is, it is empirical knowledge that already informs my knowledge of the argument's second premise, in which case the fact that we can derive it should not seem in any way surprising or absurd. This seems to me a reasonable response to anyone who wanted to use Argument III as part of showing the absurdity of combining functionalism about mental states with the idea that we ordinarily know our mental states without checking our brains.

We can now return to ABC and apply to it some of the observations made in assessing arguments I-III. Recall that MBB want to use (arguments like) ABC so as to raise a special problem for those who want to adhere to both externalism (EXT) and content self-knowledge (CSK), given that the premises of ABC are knowable non-empirically. This means that we must understand these premises as representing distinctive instances of ordinary content self-knowledge and of externalism, respectively, and as offering things we can come to know without studying our environment. I submit, however, that there is no proper construal of the relevant knowledge claims on which the first premise of ABC is known through **ordinary** content self-knowledge, its second premise is simply derived from externalism with **no** further empirical study, and its conclusion presents something that **cannot** be plausibly thought to be knowable purely on the strength of such premises. Intuitively, the problem I see is that, to get their *reductio*, MBB must either implicitly and illicitly ‘beef up’ what is involved in, or entailed by knowing A, or ‘water down’ what is claimed in B. In either case, we don’t get a stable reading of ABC that supports the inflation argument.

To see what I have in mind, let us go back to the second of my preliminary remarks. I pointed out (see above, p.) that ABC’s second premise (B), just like the second premises of Arguments II and III, is an **application** of a general theory to a particular case. The general thesis of content externalism does not tell us which of our concepts exhibit dependence on the environment, nor does it tell us which environmental conditions (if any) are relevant to the individuation of particular types of concepts, or whether they obtain in the case of individual thinkers. We can get the right application of externalism, so as to have a ‘conditional bridge’ between A and C, **assuming** the key concept in the self-ascribed thought mentioned in A, namely, the concept *water*, is a natural kind concept (as opposed to, say, a fictional concept), and assuming, further, that it is the presence of a substance with the chemical structure of **water** (i.e., H₂O) that is required for possession of that concept. But what we are assuming needs to be established, and I don’t see how it could be established without any empirical investigation. How can it be established what kind of concept the *water* concept is without studying the cognitive

organization of users of that concept? And how do we determine what particular substance is relevant to possession of that concept absent a study of the users' environment and history? Thus, even if we accept that the general externalist thesis can be established in an empirically innocent way, there is no reason to suppose that its application to the relevant case (in the form of premise B, properly understood) can be established in such a way.

This takes us back to my other preliminary remark, and to our earlier discussion of the character of ordinary content self-knowledge. Given the foregoing, it seems as though, to be explicit, we should replace B with something like

B': If I am using a natural kind concept in thinking that there's water in my glass, and that concept is the *water*-concept, as opposed to some other candidate, then the substance water exists in my environment.

C would then be replaced by

C': The substance water exists in my environment.

To get to C', we would then need to replace premise A with

A': I am using a natural kind concept in thinking that there's water in my glass, and that concept is the *water*-concept, as opposed to some other candidate.

Does A' represent something I can know simply in virtue of knowing A? If my rejection of the recognitional conception of ordinary content self-knowledge is correct, then it should be clear that my knowledge of what I am thinking neither requires nor affords knowledge that I have met the externalist conditions on having thoughts with a particular content. I have argued that my knowledge that I am thinking a *water*-thought is, as regards content, not even like my knowledge that the man in front of me is George. For, my knowledge of content is not grounded in my recognizing the content 'from the inside' through distinctive marks that purport to separate it from other candidate contents. On the non-recognitional, neo-expressivist view I advocate, my ordinary knowledge that I am thinking that there's water in the glass is grounded in my use of **that** particular content and no other in the course of expressing the thought I am having. Such

knowledge neither depends on my knowing **which concepts I am employing** in thinking the thought, in the sense of knowing their specific individuation conditions, nor does it allow me to infer this information (thereby guaranteeing that I can know it apriori, given my knowledge of A).²⁰ Thus, A' is not a premise that I can know simply in virtue of knowing that I am thinking that there's water in my glass.

I would argue, further, that there is no other empirically innocent way for me to know A'. Knowing A' would require me to engage in a cognitive study of the concepts I employ and to discover their individuation conditions. If externalism is true, the required investigation cannot be exhausted by an inspection of my mind 'from the inside'. It requires determining what environmental conditions have occasioned and controlled my use of the relevant concepts. Quite possibly, to know that I am using the *water* concept (as opposed to other candidates), I'd need to know already about the presence of water in my environment! This means that here, as in the case of Argument II, I'd need to have some way of learning the truth of ABC's conclusion that is independent of knowing its premises.²¹

Is there any information about my environment that I could derive from my knowledge of externalism alone without any specific investigation into the nature of my concepts or my environment? Well, suppose (contrary to what I said above) that I could somehow determine through introspective reflection that the concept *water*, which is an ingredient in the thought I ascribe to myself in A, is a natural kind concept. Then, given externalism, I could obtain

B": If I am thinking that there's water in my glass, then there is in my environment a substance with some particular underlying nature.

From B" together with A, we can derive:

C": There is in my environment a substance with some particular underlying nature.

On the (dubious) assumptions that establishing externalism requires no empirical study, and that I can determine by pure reflection which of my concepts are natural kind concepts, it may not be so implausible that we can know non-empirically that our extra-mental environment must contain

some substances with underlying natures which constrain our thoughts.²² (After all, the literal truth of externalism seems to require that.) Notice, first, that this would be a far cry from the claim that externalism commits us to the possibility of knowing without any empirical investigation that our environment contains **water** (as opposed to twater), for example, or else to the impossibility of ordinary knowledge of the contents of our thoughts. But even setting this aside, if we replaced B with B", it would become unclear what role is to be played in the inflation argument by the idea that we have a special, non-empirical way of knowing the contents of our thoughts. My ability to tell non-empirically that right now I have a thought with the specific content **that there's water in my glass** would need to play no role in enabling me to infer C". There is any number of much less specific or committal premises that could get me to C" (given the right instance of externalism), including the very generic premise that I have some natural kind concepts. But in that case, the inflation argument would have nothing specifically to do with the combination of the content externalism and the doctrine of content self-knowledge.

The inflation argument we originally attributed to MBB appealed to my ability to know in a special, non-empirical way exactly **what** it is I am thinking – e.g., that there's water in the glass – combined with my alleged ability to know by pure reflection the specific externalist conditions on having thoughts with this **particular** type of content. Yet none of the versions of the argument we have considered have succeeded in advancing a *reductio* against the combination of EXT and CSK along these lines. We saw that, to develop such a *reductio* the inflationist must presume that we can know by pure reflection what conditions have to obtain for the exercise of particular concepts (specifically, the ones deployed in the self-ascription displayed in step A). Yet, even if we grant that externalism can be established without any empirical study, there is no reason to think that the externalist conditions necessary for possession of specific concepts can be established absent such a study. More importantly, we saw that, for the inflationist reasoning to go through, it must be assumed that ordinary content self-knowledge can afford us rather substantial knowledge of our concepts and their individuation conditions. Given

my earlier rejection of the recognitional conception of ordinary content self-knowledge, this assumption should be seen as highly implausible. Finally, we saw that if the inflationist weakens the premises of the ABC argument, so it becomes more plausible that they **are** knowable without empirical study, it becomes correspondingly more dubious that the inflationist argument can reduce to absurdity the combination of externalism and the doctrine of ordinary content self-knowledge. I conclude that, even if content externalism is true, we can know the contents of our thoughts in a way that withstands a skeptical threat based on externalist alternatives. At the same time, given the character of content self-knowledge, knowing the contents of our thoughts does not afford us easy knowledge of the external world. So inflation of self-knowledge is no more warranted than its deflation.²³

REFERENCES

- Bar-On, Dorit. 2000. "Speaking My Mind." *Philosophical Topics* 28, no. 2, pp. 1-34.
- and Douglas Long. 2001. "Avowals and First-Person Privilege." *Philosophy and Phenomenological Research* 62, no. 2 (March 2001): 311-35.
- and Douglas Long. "Knowing Selves: Expression, Truth and Knowledge." in *Privileged Access: Philosophical Accounts of Self-Knowledge*. edited by Brie Gertler. Ashgoode Epistemology and Mind Series. (Forthcoming.)
- , forthcoming A. *Speaking My Mind: Expression and Self-Knowledge* Oxford University Press.
- , forthcoming B. "Externalism, Skepticism, and the Recognitional Conception of Self-Knowledge"
- Bernecker, Sven. 1998. "Self-Knowledge and Closure." Pp. 333-49 in Ludlow and Martin (eds.).
- Boghossian, Paul A. "What the Externalist Can Know A Priori." Pp. 271-84 in Wright *et. al.* 1998. (Originally published in *Proceedings of the Aristotelian Society* 97 (1997): 161-75).

- “Content and Self-Knowledge.” Pp. 149-74 in Ludlow and Martin (eds.) 1998.
 (Originally published in *Philosophical Topics* 17 (1989): 5-26).
- Brown, Jessica. “The Incompatibility of Anti-Individualism and Privileged Access.” Pp. 185-93
 in Ludlow and Martin (eds.) 1998. (Originally published in *Analysis* 55 (1995): 149-56).
- Brueckner, Anthony. 1998. “What an Anti-Individualist Knows A Priori.” Pp. 197-205 in
 Ludlow and Martin (eds.). (Originally published in *Analysis* 52 (1992): 111-8).
- 2000. “Externalism and the Apriority of Self-Knowledge.” *Analysis* 60, pp. 132-
 136.
- Burge, Tyler. 1979. “Individualism and the Mental.” *Midwest Studies in Philosophy* 4, pp. 73-
 122. Pp. 21-83 in Ludlow and Martin (eds.).
- 1988. “Individualism and Self-Knowledge.” *Journal of Philosophy* 85, no. 11, pp.
 649-63. Pp. 65-79 in Quassim Cassam (ed.).
- 1996 “Our Entitlement to Self-Knowledge.” *Proceedings of the Aristotelian Society*
 96, pp. 91-116. Pp. 239-64 in Ludlow and Martin (eds.).
- 1998. “Memory and Self-Knowledge.” Pp. 351-70 in Ludlow and Martin (eds.).
- Davies, Martin. 1998. “Externalism, Architecturalism, and Epistemic Warrant.” Pp. 321-61 in
 Wright *et. al.* (eds.).
- Dretske, Fred. “Externalism and Self-Knowledge.” (unpublished ms.)
- 1995. *Naturalizing the Mind*. Cambridge, MA.: MIT Press.
- Evans, Gareth. 1982. *The Varieties of Reference*. Oxford.: Clarendon Press.
- Falvey, Kevin, and Joseph Owens. 1994. “Externalism, Self-Knowledge, and Skepticism.” *The*
Philosophical Review 103, no. 1, pp. 107-37.
- Fricker, Elizabeth. 1998. “Self-Knowledge: Special Access Versus Artefact of Grammar – A
 Dichotomy Rejected.” Pp. 155-206 in Wright *et. al.*.
- Macdonald, Cynthia. 1998. “Externalism and Authoritative Self-Knowledge” in Wright *et. al.*

- McKinsey, Michael. 1998. "Anti-Individualism and Privileged Access." Pp. 175-84 in Ludlow and Martin (eds.). (Originally published in *Analysis* 51 (1991): 9-16).
- McLaughlin, Brian P., and Michael Tye. "Externalism, Twin-Earth, and Self-Knowledge." Pp. 285-320 in Wright *et. al.*
- Peacocke, Christopher. "Our Entitlement to Self-Knowledge: Entitlement, Self-Knowledge and Conceptual Redeployment." Pp. 265-303 in Ludlow and Martin (eds.). (Originally published in *Proceedings of the Aristotelian Society* 96 (1996): 117-58).
- Pryor, James. 1999. "Immunity to Error Through Misidentification." *Philosophical Topics* 26, no. 1 & 2, pp. 271-304.
- Putnam, Hilary. 1975. "The Meaning of 'Meaning'." Pp. 215-71 in Putnam, *Mind, Language and Reality*. Cambridge; Cambridge University Press, 1975.
- Raffman, Diana. 1998. "First-Person Authority and the Internal Reality of Beliefs." Pp. 363-70 in Wright *et. al.*
- Shoemaker, Sydney. 1994. "Self-Reference and Self-Awareness." Pp. 80-93 in Cassam (ed.). (Originally published in the *Journal of Philosophy* 65, no. 19 (1968): 555-67).
- Villanueva, Enrique, ed. 1998. *Philosophical Issues, 9: Concepts* edited by. Atascadero, CA. Ridgeview Publishing Company.
- Warfield, Ted. 1998. "Self-Knowledge and Privileged Access Are Compatible." Pp. 215-34 in Ludlow and Martin (eds.). (Originally published in *Analysis* 52 (1992): 232-7).
- Wright, Crispin. "Self-Knowledge: The Wittgensteinian Legacy." Pp. 15-46 in Wright *et. al.* (eds.).
- Wright, Crispin, Barry C. Smith, and Cynthia Macdonald (eds.). 1998. *Knowing Our Own Minds*. Oxford.: Clarendon Press.

NOTES

¹ The *loci classica* are Putnam (1975) and Burge (1979).

² I discuss the recognitional conception and its role in generating external content skepticism in [forthcoming B], as well as in Chapter V of my *Speaking My Mind: Expression and Self-Knowledge*.

³ Dretske (unpublished) makes the interesting suggestion that, while we are practically infallible in the assignment of content to our present states, we tell **that** we are in a given mental state (or in any mental state at all) only using normal channels of obtaining knowledge, and the knowledge we have is no more secure than ordinary empirical knowledge. I think this view departs from our commonsense conception. (For what I take to be the commonsense conception of the security of avowals see Bar-On 2000, and Bar-On and Long, 2001.)

⁴ I develop the foregoing line of thought in my [forthcoming B].

⁵ Boghossian (1989), Wright (1998), and Fricker in Wright (1998) discuss the view that ordinary self-knowledge (so-called) involves 'no **cognitive** achievement'. Since it may be possible for a thinker's judgment to represent genuine knowledge even if it involves no **cognitive** achievement, I prefer to put the deflationist challenge in overtly epistemic terms. (We could attribute to the deflationist the claim that ordinary mental self-ascription represent no **epistemic** achievement.)

⁶ I motivate and develop the view I favor in *Speaking My Mind: Expression and Self-Knowledge* [forthcoming A].

⁷ Shoemaker (1994: 82-3).

⁸ See Evans (1982). For a recent discussion of this notion and for references see James Pryor (1999). I should note that I use "ascription" where others sometimes use "judgment," where an ascription can be made in speech or in thought.

⁹ It is useful to distinguish in this connection between a 'thin' or referential notion and a 'thick' or recognitional notion of identifying an object. The former is a semantical notion; the latter is an epistemic one. Compare Evans (1982: 218).

¹⁰ Evans (1982: 220, 222).

¹¹ Evans (1982: 221).

¹² I would also argue that I do not tell that it is annoyance that I feel via recognizing that I am in some state and taking that state to be that of annoyance, based on its characteristic qualitative features. Here, however, I focus on the content aspect of our ordinary present-tense self-ascriptions of contentful mental states.

¹³ As suggested by the previous footnote, I maintain that avowals are immune to errors of misascription as regards the state ascribed, as well as its content (when the state ascribed is a contentful state). Yet, on my view, avowals in general are not infallible – they **can** be false. (So-called self-verifying self-ascriptions are an exception; see below.) How can that be? The short answer is that, while avowals can be false, when they are, the falsity is not due to the avowing subject **mistaking** her state or its content for some other state or some other content. (For my account the kinds of failure that can issue in false avowals, see Chapter 8 of my *Speaking My Mind: Expression and Self-Knowledge*.)

It should also be kept in mind that one **can** issue a self-ascription that is **not** immune to errors of misascription (or even to errors of misidentification). We sometimes do ascribe present mental states to ourselves on various indirect (‘third-person’) bases. As emphasized earlier, the kind of immunity to error that concerns us here is not a matter of the semantics of the relevant ascription, but a matter of its epistemology – it has to do with the basis on which it is made (or lack thereof).

¹⁴ For an analysis of self-verifying thoughts that is designed to be consistent with content-externalism, see Burge (1996).

¹⁵ I think this can help clarify, and vindicate, the intuition that content self-ascriptions in general, and not only self-verifying self-ascriptions of entertained thoughts, have a self-referential or reflexive element. See, e.g., Burge (1979: 70), and Peacocke in Ludlow and Martin (1998).

¹⁶ My neo-expressivist account departs from traditional expressivist accounts of avowals in several ways, the most crucial one for present purposes being that it insists that avowals differ **semantically** from, e.g., natural expressions of mental state, in that an avowal makes a **genuine self-ascription** of a state.

¹⁷ For a lengthier discussion, see Bar-On and Long (forthcoming).

¹⁸ I am here extracting a single argument that is based on papers by McKinsey and by Brown (both reprinted in Ludlow and Martin (1998)), as well as by Boghossian (reprinted in Wright (1998)), abstracting

from unnecessary complications and infelicities. For relevant discussions of the argument, see especially the articles by Brueckner, McLaughlin and Tye, and Ted Warfield in Ludlow and Martin (1998), and the articles by Bernecker, Burge, Davies, and Raffman in Wright (1998). See also replies to Boghossian in Villanueva (1998).

¹⁹ For a critical examination of this claim, see Falvey and Owens (1994), Brueckner in Ludlow and Martin (1998) and Brueckner 2000.

²⁰ There is, then, a certain disanalogy between the case of knowing George by acquaintance (see Argument II above) and having ‘expressive knowledge’ of content. I can know George as someone I met a few days ago (or as someone who meets some other, perhaps visual, description) without knowing him as someone who fits some other descriptions. Similarly, I can know what I am thinking as *that there’s water in my glass*, without knowing it as meeting certain other conditions. However, on my neo-expressivist view, knowing the content ‘as’ *that there’s water in my glass* is not a matter of bringing it under any description – it is a matter of using **that** content (which, I maintain, is a legitimate way of **specifying** a particular content). I believe (though I cannot argue this here) that this difference means that, **even if** we accept that knowledge, in general, obeys some Closure Principle, (expressive) knowledge of A would not support inferential knowledge of A’.

²¹ For relevant discussion, see the articles by Brown, and by McLaughlin and Tye, in Ludlow and Martin (1998).

²² See the conclusion of Brueckner in Ludlow and Martin (1998).

²³ I wish to thank Katya Hosking, Janine Jones, Douglas Long Bill Lycan, Ted Parent, and Zena Ryder for helpful exchanges on issues discussed in this paper. Special thanks to Carol Voeller, Ram Neta and Dylan Sabo for valuable comments on earlier drafts, and for editorial suggestions.