

De se attitudes and computation

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Hamilton Fairley, N. (2021) De se attitudes and computation. *Theoria*, 87 (1). pp. 207-229. ISSN 1755-2567 doi: 10.1111/theo.12241 Available at <https://centaur.reading.ac.uk/81918/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/theo.12241>

Publisher: Wiley

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

De Se Attitudes and Computation

by

NEIL HAMILTON FAIRLEY

University of Reading

Abstract: There has been debate between those who maintain that indexical expressions are not essential and those who maintain that such indexicals cannot be dispensed with without an important loss of content. This version of the essentialist view holds that thoughts must also have indexical elements. Indexical thoughts appear to be in tension with the computational theory of mind (CTM). In this case we have the following inconsistent triad:

- (i) *De se* thoughts are essential.
- (ii) *De se* thoughts are indexical, they have a (Kaplanian) character.
- (iii) Computations can only take the syntactic *type* into account, they cannot take tokens into account.

If (iii) is correct, then it seems we cannot make sense of a thought which uses a character such that its referent could vary from tokening to tokening. I argue that (iii) need not cause a problem, while maintaining the CTM. I claim that computations need not be sensitive to the features of a tokened symbol in the way that character demands. This job may be performed by a non-modular part of the mind. Resolving the triad in this way provides a reason to accept that indexicals in thought are possible.

Keywords: de se, computational theory of mind, mental files, Recanati, Perry, concepts, thoughts, context sensitivity

1. Introduction

ESSENTIAL INDEXICALITY in natural language is the “dominant” position in the literature (Cappelen and Dever, 2013, p. 2). Essential indexicality is the claim that certain indexical attitudes cannot be replaced by non-indexical attitudes without losing an important feature of that attitude (Recanati, 2012, p. 33). A second position is *de se* essentialism, which states that certain attitudes that one has about oneself (the *de se* attitudes) cannot be replaced by attitudes of a different kind without a loss or change in what that attitude is able to explain. This view is often defended on the basis that *de se attitudes* pose a unique problem for theories of propositions.

A prior version of this article was erroneously published pre-acceptance and retracted. Retraction: Neil Hamilton Fairley, *De Se Attitudes and Computation*, *Theoria* 2019. <https://doi.org/10.1111/theo.12175>

De se attitudes are often taken to have their own kind of content. *De se* contents are contents that make a perspectival or first-personal reference to the person entertaining that content and are usually expressed in English with “I”. *De se* attitudes are often taken to need an indexical element. If this is correct, then it seems that *de se* thoughts (thoughts tokened when having a *de se* attitude) will also be essentially indexical. When one has a *de se* attitude one has a *de se* thought, understood as a sentence in Mentalese, e.g., I AM IN READING, and the thoughts have their own kind of content which makes a reference to the individual thinking it.

I will focus on *de se* essentialism as including an indexical element, and will use the term to reflect this, because this view seems promising for the essential indexicalist and has come under pressure in recent debates (see Ninan, 2015; Cappelen and Dever, 2013; Lewis, 1979; Perry, 1979; Millikan, 1990; Bermudez, 2017). This is a view that has been targeted by *de se* sceptics who deny that there is anything essential about *de se* attitudes. *De se* essentialism can provide a useful test case for the possibility of indexicals in thought, provided that *de se* attitudes are also ones that include an essential indexical element.

A key problem with *de se* essentialism understood as involving indexicals is that it seems incompatible with the computational theory of mind (CTM). It has been pointed out that there is a tension between these two views, as computations struggle to ascribe different tokens of the same type of symbol different referents, but this is what indexicals require (Ball, 2010).¹ (Throughout this article I will use “computation” as short for classical or Turing style computation.) For instance, different tokens of “now” will refer to different times. For symbols in Mentalese to be indexicals they need a context of use to determine their content. So, each new token of an indexical type will take the context of use into account, potentially yielding a different content in each case. So, each token of this type is treated as a token rather than an instance of a syntactic type. However, it seems that computations can only treat symbols as an instance of a given syntactic type.

While this tension seems to make the CTM and *de se* essentialism incompatible, which would be a serious problem, I will argue that *de se* essentialism can be compatible with the CTM. The cases where a tension might arise, for instance in a mental module dedicated to processing syntax for input or output of linguistic items, do not threaten to make the two theories incompatible with one another. Instead, I will argue that we should not expect computation alone to determine the reference of our thoughts.

¹ When it comes to homonyms and proper names, it may be that the best way to understand these on a computational theory of mind is to give each one an identifying subscript. So “bank” might become BANK1 and BANK2, for the financial institution and riverside respectively. In these cases, we no longer have an ambiguity that the context needs to deal with.

It follows that there need not be a tension between *de se* essentialism and the CTM. I will claim that the inconsistency depends on a certain understanding of the CTM according to which computation is the only means of processing information and reasoning. I will follow Fodor (2000) in denying that we should accept such a strong view of the CTM. Instead, the CTM need only commit one to the claim that some processes and reasoning undergone by the mind are computational.

On this weaker understanding of the CTM we do not need to maintain that all the processes in the mind are computational. This opens space for non-computational processes to occur in the mind. This means that even if we allow that computations cannot deal with indexicality, we do not require computation to do so to have *de se attitudes* that are indexical. I will conclude that this gives us reason to accept that context sensitivity is possible at the level of thought.

2. Inconsistent Triad

The following three claims are widely held, and it is these claims which Ball (2010) states are incompatible with one another (in an argument to be discussed in section 2).

- (i) Having a *de se attitude* is essential to grasping propositions in a *de se* way.²
- (ii) *De se thoughts* are indexical, in the sense that they have a context-sensitive (Kaplanian) character (Kaplan, 1989, p. 506).
- (iii) Computations are only sensitive to the syntactic *type* of a symbol; they are not sensitive to other properties that the tokens that they operate on may have (Ball, 2010).

For a *de se* thought, a thought entertained when in a *de se* attitude, to be a *de se* thought, one must realize that the thought is about oneself. Evans adds two necessary conditions for thoughts to meet in order for them to be ones that we have about ourselves.

First, one must recognize the thought's significance to the "various special ways he has (as every person does) of gaining knowledge about himself" (Evans, 1982, p. 206).³ So if one entertains a thought that is about oneself but does not appreciate that it can contribute to knowledge of oneself, then they do not realize

2 Perry (1979/2010, pp. 374–375) expresses this as certain belief states being necessary to access propositions in a particular way. Perry (1977), Ninan (2016) and Lewis (1979) discuss a similar view. Catañeda (1968, pp 446–447) also discussed similar cases which motivated much of the ensuing discussion on *de se essentialism*.

3 Recanati (2012/2017) makes a similar point about the SELF file depending on various epistemically rewarding relations that we have to ourselves.

that that thought is about them. A second necessary condition is that one “must realise how to *act upon* propositions” (Evans, 1982, p. 206).

Perry’s arguments are to the effect that we cannot explain some behaviour without positing a *de se* attitude. The argument’s form is that if two thoughts express the same proposition then they should have the same effect on the person’s behaviour (all else being equal). When we have a *de se* thought, other non-*de se* thoughts do not have the same effect on the person’s behaviour (all else being equal). Therefore, these two thoughts are not interchangeable, and the *de se* attitude is essential.

To motivate *de se* essentialism, Perry introduces a case in which he is following a trail of sugar in a supermarket, trying to tell the shopper that their sugar bag is leaking. As he searched for the shopper:

the trail [of sugar] became thicker. But I seemed unable to catch up. Finally it dawned on me. I was the shopper I was trying to catch. I believed at the outset that the shopper with a torn sack was making a mess. And I was right. But I didn’t believe that I was making a mess. That seems to be something I came to believe. And when I came to believe that, I stopped following the trail around the counter, and rearranged the torn sack in my cart. My change in beliefs seems to explain my change in behaviour. (Perry, 1979/2010, p. 366)⁴

Perry then wants to suggest that there is something special about this change in belief. Without the use of “I” in characterizing the belief, it seems that we no longer have an explanation for Perry’s behaviour:

When we replace [I] with other designations of me, we no longer have an explanation of my behaviour and so, it seems, no longer an attribution of the same belief. It seems to be an *essential* indexical. (Perry, 1979/2010, p. 366)

The evidence for this claim is that we can explain some changes in behaviour only by shifting from a non-*de se* expression (such as “a shopper”) to a *de se* expression (“I”). When Perry thinks of the person making a mess as “a shopper” he continues to search for them. It is only when he realizes “I am making a mess” that his behaviour changes. It seems that even if Perry believed “Perry is making a mess” he would still need to have a belief to the effect that “I am Perry” for this to make a difference to his behaviour.⁵

4 See also Perry (1977) for further examples.

5 Recanati (2017, p. 180) defends a way of inferring, from the premises that **a** is *F* and the premise that **b** is *G* to the conclusion that something is both *G* and *F* without invoking a premise to the effect that **a** = **b**. This is referred to as trading upon identity. However, this move requires that two mental “files” be associated with one another, in that they both fall under a third file based on multiple epistemically rewarding (ER) relations. It seems that the self would have a file that is based on multiple ER relations. Even if the self file (presumably the relevant one for *de se* attitudes) was not based on multiple ER relations (which seems implausible), it seems that trading on identity still requires some association between two files, and this may be enough to maintain Perry’s point that there needs to be some sort of belief to the effect that “I am Perry”.

Perry's view has recently been targeted as being a special case of Frege problems (Cappelen and Dever, 2013, p. 59). Just as one learns that Hesperus is Phosphorus, so Perry learns that Perry is I. In which case, *de se* attitudes do not pose a unique problem that would motivate their essentialism.

To illustrate this, consider a variant of the messy shopper case:

Pushing my cart down the aisle I was looking for CK (Clark Kent) to tell him he was making a mess. I kept passing by Superman, but couldn't find CK. Finally, I realized, Superman was CK. I believed at the outset that CK was making a mess. And I was right. But I didn't believe that Superman was making a mess. That seems to be something that I came to believe. And when I came to believe that, I stopped looking around and I told Superman to clean up after himself. My change in beliefs seems to explain my change in behavior. (Cappelen and Dever, 2013, p. 61)

This seems analogous to Perry's case and if there is no difference between them then the problem posed by *de se* attitudes is not sufficient to single them out as essential.

Here I will give a summary of Ninan's response to this worry, to help motivate (i). Both (i) and (ii) are important for there to be an apparent tension with (iii). Without (ii) there is no need for indexicality and without (i), or some equivalent, e.g., *de nunc* attitudes, there is no need for indexical attitudes either. So, it is important to give some consideration to each of these points, though my focus will be on a tension between (i)–(iii). Ninan argues that *de se* attitudes create a conflict between two desiderata that we should have for a theory of propositions, beyond those generated by Frege cases. These desiderata are Agreement and Explanation.⁶ Two agents *x* and *y* Agree when they have beliefs that are equivalent to one another (Ninan, 2016, p. 100). For instance, if each is in a belief state such that they believe that the sky is blue, then they agree. If one believes that the sky is not blue, then they do not agree. However, there is an understanding of Agreement which concerns their desires as well. These will be cases in which two people have the same desires, relative to the case at hand (Ninan, 2016, p. 100). So if we are watching a sports match and desire that the same team win, we have Agreement in terms of our desires relative to that case.

A condition for Explanation is that two agents Agree on a relevant set of beliefs and desires. As they want the same outcome and believe the same things then they should perform the same action, all other things being equal (Ninan, 2016, p. 102). Explanation can be understood as a conditional, so that if two agents Agree and all else is equal then they will perform the same action.

Ninan outlines his problem as follows:

⁶ I use capital letters to mark these technical uses of these terms.

AGREEMENT would seem to entail if two agents agree on how things are and on how they would like them to be, then those agents have the same beliefs and desires, i.e. they believe and desire the same propositions. And EXPLANATION tells us that if two agents have the same beliefs and desires, then, other things being equal, they will behave in the same way. But then it follows from those two claims that if two agents agree on how things are and on how they would like them to be, then, other things being equal, they will behave in the same way. But, as I shall now argue, this final claim is false. If my argument is sound, it follows that AGREEMENT and EXPLANATION are inconsistent ... (Ninan, 2016, p. 103)

To show this, consider a case in which Ninan is being chased by a bear, and you are watching this. Ninan believes I AM BEING CHASED BY A BEAR, and that curling into a ball is the best way to avoid being mauled by a bear. Watching, you may think NINAN IS BEING CHASED BY A BEAR and believe that curling into a ball is the best way to avoid being mauled by a bear. You both desire that Ninan not be mauled by a bear. So, you and Ninan seem to have met the conditions for Agreement. It follows that you should both perform the same action because of Explanation. So, you will both curl into a ball. But this prediction seems false (Ninan, 2016, p. 104). Ninan will curl into a ball; you will probably do something else. From this it follows that either Agreement or Explanation is false. This is a problem of *de se* attitudes that does not seem to be reducible to a Frege case.

An objection is that there is not Agreement in Ninan's case. Does the observer have the same belief as Ninan does, if one is expressed indexically (Ninan's) and one is not (the observer's)? Ninan assumes that there is Agreement. "Now if, as I am assuming, we count as agreeing that I am being chased by a bear in virtue of my having this *de se* belief and your having the corresponding *de te* belief, then it would seem to follow from AGREEMENT that these beliefs have the same content" (Ninan, 2016, p. 103). This assumption is perhaps something that Ninan's opponents would want to deny. If Ninan and the observer do not Agree, then Ninan's counter-argument does not work. There is no Agreement and therefore no Explanation either. So, there is no contradiction between the two.

However, if one adopts the perspective of the sceptic, then it appears there should be Agreement, since for the sceptic there is no substantive difference between Ninan's belief and the observer's belief. There is nothing special about the *de se* attitudes on the sceptic's view. The thought that "I am being chased by a bear" could just as easily be replaced by the thought "Ninan is being chased by a bear" without a change in Explanation. In this case, the two agents in Ninan's case Agree with one another. But then they should act in the same way, in virtue of Explanation. This seems implausible. So, arguing that Ninan and his observer do not Agree seems to undermine the sceptical position.

If *de se* essentialism is correct, then Ninan's belief state and the observer's belief state can have the same content. Where they differ is in terms of how the

contents are accessed. This also has the advantage of preserving a level of content at which *de se* attitudes and non-*de se* attitudes express the same content.

Claim (ii) is that these *de se* attitudes make use of a context-sensitive character to determine their reference. I take (ii) to be a part of this version of *de se* essentialism. This does not seem to be an uncommon assumption, and including an indexical seems to be essential to *de se* attitudes being represented in natural language (see, for example, Recanati, 2012, p. 33). However, it is not clear whether this should also be a requirement at the level of thought. In this article my main aim is to show that indexicals are possible at the level of thought, with some reason to think that there are such indexicals. Proving indexicality would go beyond the scope of this article.

It should be noted that many endorse (i) without also endorsing (ii). Notable examples include Peacocke (1981, pp. 189–190), Evans (1982, pp. 35, 159, 205; 1981, pp. 280, 284) and McDowell (1984, pp. 288–289); who argue that reference is determined by a (*de re*) sense instead of by character. In this article I will briefly defend (ii) but I will not engage too much with these views as my aim is to show how the CTM can be compatible with indexicals, rather than to show that (i) entails (ii). For the purposes of this article, I will nevertheless assume that one can maintain both (i) and (ii). This is significant as it allows for the possibility of context-sensitive thoughts on a CTM. Whether that view is correct is a discussion for a later time.

The indexicality that I have in mind is Kaplan's context-sensitive characters, which are non-constant. He claims that "Indexicals have a *context-sensitive* character. It is characteristic of an indexical that its content varies with context" (Kaplan, 1989, p. 506). These context-sensitive characters require a context of use to get a content. The important point here is that these indexicals do not have a fixed reference across all possible contexts of use; rather their reference depends on the context of use in which the indexical is tokened. It should be noted here that though *de se* attitudes might be indexical in Kaplan's sense, they may not actually change their referent in many, if any, contexts of use. In this sense, the *de se* attitudes will not be practically indexical (Ball, 2015, p. 354). A practical indexical is one that "will very often link different tokens to different referents, even if we hold the speaker fixed. Call terms ... that shift reference in this way practically indexical" (Ball, 2015, p. 354). By contrast, practical non-indexicals will not shift their referents in this way. Given that an individual's SELF concept will not shift its referent across contexts of use, these thoughts are not practically indexical.⁷ However, this does not mean that they are not indexical at all. They may still depend on the context of use to get a reference.

(iii) is the point that computations will interact with all instances of a given syntactic type in the same way (Ball, 2010). So while computations operate over

⁷ Here I follow the convention of using capitals to show when a concept is being spoken about.

particular instances of a syntactic type, they operate on tokens; they are not capable of treating tokens of the same type differently depending on the context in which they occur. This is because computations can only track syntactic properties. They cannot consider non-syntactic properties like the context of use. So, computations can only pay attention to a semantic type indirectly, by paying attention to syntactic type. I will refer to this as the computation acting on a symbol *qua* type, rather than *qua* token. Treating a symbol *qua* token would involve paying attention to its non-local properties so that its semantic content is not just a product of its syntactic type.

Claims (i) and (ii) are of interest to me as, when taken together, they support an underdeterminacy of content at the level of thought. This kind of underdeterminacy occurs when a given representational item S makes an explicit reference to the context of use to establish a content or a reference. If this underdeterminacy exists at the level of thought, then it seems that there is some context sensitivity at the level of thought. If (i) and (ii) are true, then this underdeterminacy is essential.

However, Ball has made an argument to the effect that the CTM is incompatible with indexical thoughts. The challenge can be put as follows. Computations can only treat a symbol as an instance of a syntactic type. Indexicals allow for different tokens of the same syntactic type to be of different semantic types. So, computation cannot make sense of how indexicals function in thought. The conclusion seems to be that we should not make use of indexicals if we have a CTM.

In what follows I will present Ball's argument for the incompatibility of these theories. I will then present some possible responses to Ball's argument. My preferred response is to allow a non-modular, non-computational, general system in the mind to ascribe reference. If this response is correct, then Ball's argument depends on a false premise and is therefore unsound. A CTM need not rely only on computations to function. Two competing views will also be touched on. These are the names only view and an indexing view. I will argue against both.

3. Argument 1 for Incompatibility

Claim (iii) is the claim that computations are only sensitive to the syntactic *type* of a symbol; they are not sensitive to other properties that the tokens that they operate on may have. Here is an argument taken from an unpublished paper by Ball. It is not clear that he fully endorses the argument and his published work seems to show a more concessionary position according to which whatever can be done by an indexical account of attitudes can be done by a non-indexical account (see Ball, 2015 for details). However, the argument presents a serious worry for the triad of views that I wish to endorse, so I will take it seriously here.

I will refer to this as Argument 1. Argument 1 draws on the syntactic nature of the CTM to argue that the CTM is incompatible with indexicality.

Argument 1 can be put as follows:

- (1) If there are constituents of the language of thought that have an indexical character, then the computational theory of mind is not true.
- (2) The computational theory of mind is true.
- (3) Therefore, there are no constituents of the language of thought that have an indexical character (1, 2, MT)

It seems to follow from this argument that the CTM is incompatible with indexicals in thought, so (i)–(iii) are incompatible. Kuczynski (2007, p. 237) makes a similar point but argues that this means we should dispense with the CTM. As my aim is to show that there can be indexicals in thought, I will focus on Argument 1 as it provides a better foil to this position.

The idea is that computations must treat all instances of a syntactic type in the same way. So, any types which require different treatment are not possible. The key premise to defend is (1); if (1) is false then Argument 1 does not work. (1) is plausible in so far as the Type Sensitivity Constraint (TSC) is true (see below). The TSC captures what is important about (iii) in a more rigorous way. Thus, if (1) is true, it shows how a tension can be created between (i), (ii) and (iii).

The TSC follows from the CTM's reliance on mechanical processors to perform operations. A processor reacts in a mechanical way based on the shape of the symbol it encounters. It may copy or delete a symbol or type a new one depending on the symbol and its program. A second support for the TSC comes from the claim that if it were false then we would be liable to produce errors that we are not, in fact, liable to commit. These are errors of equivocation. I will elaborate on this later. I will first give a definition of the TSC and explain why it creates a tension with (ii).

If the TSC is true, then we have an argument for the truth of premise 1. The TSC is that:

There is some way of typing mental representations such that (i) if two representations differ as regards their content, then they are of different types; and (ii) computational mechanisms such as those postulated to account for reasoning are sensitive to type identities and difference among representation tokens. (Ball, 2010, p. 5)

For a given symbol type T, its tokens can express a different content each time, or not.⁸ If not, then there is less motivation to think that T is an indexical symbol.

⁸ The possibility that they refer to different things need not be practical here, so it may be that the symbol always refers to one individual during that person's lifetime. This does not mean that the symbol could not change its referent, just that it need not change its referent to still be encountering problems.

This claim gets additional plausibility when we consider that indexicals do not refer as types (Recanati, 2012, p. 57).

If the tokens do express a different content, then the TSC is violated. Condition (iii) is not met as the CTM entails the TSC. It seems that if the TSC is violated, then we are no longer using a CTM (Ball, 2010, p. 13). This seems to require us to choose between indexical essentialism and the CTM. Ultimately, I will claim that this is not so, but I will first present some reasons for endorsing the TSC.

One way to defend the TSC is to look at the nature of computation in the CTM (Ball, 2010, p. 4). A CTM will operate according to computations. “A computation ... is a formal operation on syntactically structured representations. Accordingly, a mental process, qua computation, is a formal operation on syntactically structured mental representations” (Fodor, 2000, p. 11). These computational processes are only sensitive to the syntactic *type* of the symbol that they are operating on.⁹ This is because the computations use a “processor” which has set reactions to syntactic features (Pinker, 1994, p. 75). In which case, the processor will be unable to take other features into account including features that it has because it is the particular token that it is.¹⁰

If the processor could take non-local features of a token into account, such as the context in which it occurs, then we are no longer dealing with a computation that can plausibly be attributed to humans when they engage in certain kinds of reasoning. It does not seem plausible that a human’s non-demonstrative reasoning processes, such as abduction, operate over a human’s entire belief system. However, if a given syntactic type *could* have its content changed depending on the context, then any number of contextual features might prove to be relevant in determining that content. So, a process that determined the content of such a symbol would need to be able to check for relevant beliefs and to be able to treat the symbol accordingly in the future. This does not seem to be the sort of thing that a computational processor can do. Instead, the processor seems to be acting in a manner akin to an intentional agent who takes several features into account. At this point, one of the key advantages of the CTM is lost as we can no longer account for some rational inferences without already making use of rational agency. This, at least, seems to be a part of the rationale behind Argument 1.

So, it seems plausible that, on a CTM, if two symbols could differ in content, then they are of different syntactic types. We should expect that if two symbols have different content then this should be reflected in how those symbols are

9 There are others who support this claim; for instance, Müller (2008, p. 119) states that “At the syntactical level, the states or physical objects are taken to be tokens of a type (e.g. charge/no charge) and are manipulated according to algorithms”. This suggests that it is their type that is important here.

10 While Argument 1 is, I think, novel, endorsement of something like the TSC on computation is not. In which case, this argument seems quite important.

treated. For instance, the belief that there is a postman at the door will (hopefully) elicit different reactions to the belief that there is a murderer at the door. So, these two contents will be treated differently. If a computational process is going to acknowledge this difference it will have to perform different operations over them. If computations are operations over different syntactic types, then one way for the process to acknowledge the different content is to make use of different syntactic types.

A second point in defence of the TSC is that if we were to deny it, then we would be liable to commit errors that we are not in fact liable to commit. For instance, “I should be disposed to infer from *that is a cat* and *that is a dog* to *something is both a cat and a dog*” (Ball, 2010, p. 13). The point is that if we could have a mental symbol whose reference varies with the context of use then the computations involved would not always (if at all) be able to track this. The processor would come across “that” and treat it as though it had the same content in each instance. In which case, we would get an error of equivocation as presented here. However, given that we do not make this sort of error as often as this theory would seem to predict, it seems that this is evidence against indexicalism in a CTM.

This point is particular to certain sorts of indexicality, and it may not apply to those cases in which we do not have practical indexicality, as is the case with *de se* attitudes. This is because there would not be a prediction of an error of equivocation. Given that the symbol type for the self would always be in the same context of use relative to what it refers to whenever a given individual uses it, i.e., the symbol type will always have the same content across all the contexts that it is likely to be used across by a given individual, then we would not predict that there would be these errors. This assumes that the “self” symbol is the same for each person, but even if it were not and each person has their own symbol, then the symbol is still not practically indexical. (This is not to say that it is impossible for there to be cases in which the prediction would change. For instance, some sort of brain transplant might have this effect. However, in this case it is not obvious that the prediction of equivocation would be incorrect.)

Given that the mental “I” that will appear in the person’s belief state is not practically indexical, it may seem like an unfair case to apply to Argument 1. However, Ball seems explicit in taking his considerations to apply to such thoughts as well. He says that:

the reason that the thought Perry expressed by “I am the messy shopper” was cognitively distinct from the thought that he expressed by “John Perry is the messy shopper” is because the two thoughts have different characters; in particular, the “I” thought has a context-sensitive character. So, the “I” thought is indexical. I will argue that this explanation cannot be correct if thinking is a matter of computation. (Ball, 2010, p. 4)

This suggests that Argument 1 is taken to apply to indexicals in general, and therefore to “I” thoughts as well, when these thoughts are indexical. These non-practically indexical symbols will still require that a computation makes use of non-local properties of that symbol, so these cases still seem incompatible with a CTM. So, it seems we cannot have (ii) if we have (iii).

3.1 A response to Argument 1

In this section I will consider some responses to Argument 1. One response is a modification of the TSC. An alternative TSC might allow for indexicals and still be consistent with the CTM. The proposed change goes as follows:

*Type Sensitivity Constraint** (TSC*) There is some way of typing mental representations such that (i) if two representations differ as regards their content, then they are of *different non-indexical types*, or of the *same indexical type*; and (ii) computational mechanisms such as those postulated to account for reasoning are sensitive to type identities and difference among representation tokens of non-indexical types, and sensitive to type identities and differences as well as other associated information among representation tokens of indexical types. (Ball, 2010, p. 13)

Using TSC* might suggest that indexical symbols can have additional features added to them to allow the processor to take the context of use into account. One is to add an additional feature to each of these symbols, such as a subscript. The example that Ball considers is to give a time stamp to each tokening of NOW to allow the processor to track it appropriately: “for example, we could imagine representations prefixed with a sort of temporal quantifier or operator (AT 2:00, THE MEETING STARTS NOW), or representations that link the NOW to a time via a sort of parenthetical (THE MEETING STARTS NOW [2:00])” (Ball, 2010, p. 14). This allows the computation to take the context into account.

This proposal is like adding a subscript to the logical form of the thought. So, in the case of *de se* attitudes an equivalent proposal would be to add a subscript that identifies the person in question. Rather than having a mental “I”, one might instead have a mental “I_{NN}” or a symbol that has a subscript that is appropriate for the individual. Here, the subscript types the referent of the symbol to allow it to get content, rather than having an indexical character.

However, we no longer have indexicals in this case. These are not types whose tokens could express different contents or which refer only because of a contextual relation to the referent; they are now different types, each with a fixed content across contexts of use. Sticking to the time example:

now it looks like for the purposes of computation, we have two distinct types: the *nows* tokened at 2:00 are type-distinct from the *nows* tokened at 3:00 ... In general, any computational system that has the resources to avoid making the bad inferences in question will end up treating the alleged indexicals as of distinct types. (Ball, 2010, p. 14)

Each symbol is a distinct type because they include the time to which they refer as a part of their symbol. In this case, we no longer have indexicals. It seems that a similar move can be made with “I”. We could have I@X, where X is the agent referred to.

There are two ways that one can take this. One is to treat these cases as indexed symbols. Another is to treat them as names. In either case the context sensitivity is removed and (ii) is denied. So, the TSC* does not present an alternative to the TSC which also allows for genuine indexicals. In which case, it seems that we still have (1). However, this now presents us with a problem. Why maintain that (ii) is true, when we could instead accept a name or an indexed view of *de se* attitudes? I will deal with these objections in turn.

4. Alternatives to (ii)

I noted a couple of alternative positions that one might take when confronted with Argument 1. One alternative is that in place of having (ii) we instead have mental proper names. This is a position discussed by Millikan. I will first argue that this is not available to someone who accepts (i). I will then consider the indexed alternative and present some arguments against that view. That these alternatives do not work as well as indexical views provides us with some reason to think that these thoughts should include indexicals. Other candidates do not seem to have the properties that *de se* attitudes have, as I will argue in what follows.¹¹

4.1 Millikan’s view: a names account of *de se* attitudes

Millikan is critical of the idea of indexical elements in thought. This is illustrated in the following quotation:

Do I succeed in identifying the content of various tokens of my mental “@RM,” that is, do I succeed in reidentifying myself, only because I grasp for each token of “@RM” independently that it bears a certain adapting relation to me? Isn’t it more reasonable that my mental “@RM” is simply a mental proper name? I take different tokens of “@RM” to refer to the same not because of their individual contexts, not as a result of some relation each of these tokens independently bears to me, but simply because they are tokens of the same type. (Millikan, 1990, p. 11)

Here, Millikan expresses her disapproval of the idea of an indexical element in thought. She finds it implausible that we would have a context-sensitive token when a name would seem to fulfil the same task of referring to oneself without the need to take the context of use into account. In which case, it seems that we

¹¹ Proving that there are indexical thoughts goes beyond the scope of this article. That would require an understanding of how thought and language relate to each other. It would also require ruling out all alternative possibilities, which would also go beyond the scope of this article.

should reject (ii). On the view that she presents (though does not ultimately endorse), “my mental ‘I,’ my ‘@RM,’ is not an indexical” but a mental proper name (Millikan, 1990, p. 11). She describes this as a Millian name so that the only semantic contribution it makes is its referent (Cumming, 2016, S2.1). (This does not seem to be the view that Millikan settles on. Instead, she opts for the view that either the self does not need to be represented at all when we are thinking about what to do, or if it does then it need only be a [Millian] name.)¹²

The challenge that Millikan presents here can be put as follows. Why should we prefer practical non-indexicality when we could make use of a proper name for ourselves that would fulfil the same function? Recall that a non-practical indexical will not change its referent across contexts of use with a fixed speaker but will still select its referent in virtue of its relationship to the context.

This is pressing as having practical non-indexicality seems to come with additional commitments that we are not certain we should be making. For instance, there might be some circumstances in which it could become practical indexicality, e.g., if my “self” symbol were to appear in your head. It also seems to add an additional mechanism of saturation, where saturation is the process that “takes place whenever the meaning of the sentence includes something like a ‘slot’ requiring completion or a ‘free variable’ requiring contextual instantiation” (Recanati, 2004, p. 7).¹³ In this instance, that would involve assigning a value to the indexical. Then the question is: what warrants this additional mechanism?

4.2 Responses to Millikan

One response is to say that these *de se* attitudes are only properly expressed in natural language by using indexicals, such as “I”.¹⁴ Even if this true, however, it does not seem to follow that the thoughts that we have must also use such an indexical component. That seems to suppose that natural language and thoughts are analogous in, at least, this respect. That seems to beg the question.¹⁵ We must assume that thoughts are like natural language, but this is what we are trying to find out. This will not do.

12 Millikan (1990, p. 12) writes: “my intentions are not designed to guide anyone’s actions but my own. Hence they have no need to explicitly represent me. I do not have to take into account variations in whose head a token of ‘@RM’ appears, nor variations in whose action it is supposed to guide. But, once again, this inarticulateness in how the self is represented has nothing to do with indexicality”. This leaves it open that the self is represented inexplicitly, and there may be a parallel with Perry (1986), whose view I will not discuss here. At present, it appears Millikan opts for any inexplicit representation to occur by default. There is no one else to whom this representation could apply.

13 If one takes a Kaplanian view then saturation we must still find a content for that character.

14 This is denied by some defenders of the *de se*, e.g., Bermudez (2017, p. 12), who claims that we can use our own names as deferred uses of “I”.

15 Ball (2010) makes a similar point, to which I am sympathetic.

If one believes that natural language is the medium of thoughts, rather than the language of thought, then supposing that natural language and thoughts are similar will not be begging the question (Fodor, 1975). This kind of view can be attributed to Carruthers (2002, p. 21). This is not, however, an option that I will be considering because I am assuming that the medium of thoughts is the language of thought.

A point against Millikan's understanding of *de se* attitudes is that in the case of a name, it seems that it is possible for a name's reference to be forgotten, or to fail to refer in some cases, unlike in the case of *de se attitudes*. Millikan's name would be somewhat different in this case, in that it seems very difficult to imagine cases in which one has a *de se* thought that fails to refer (Anscombe, 1975, pp. 22–23). In fact, *de se* thoughts seem to be incapable of failing to refer. If someone can token them, then they can refer to themselves with it, and have done so (Anscombe, 1975, p. 28). The same does not seem to be true of names. This seems to be so regardless of the theory of names that one picks. These disanalogies between the way *de se* contents work and the way that names work give us reason to resist Millikan's point. If names do not have the properties that *de se* attitudes have, then this gives us reason to believe that names are not well suited to playing a role in *de se* attitudes.¹⁶

For this point to be convincing it is necessary to prove that "I" can have these features as well. It seems that it can. Tokening "I" seems to make it refer. It is hard to conceive of a tokened mental "I" failing to refer. If it is tokened, then it will be referring to the thinker that tokened it. For these reasons it seems that we can justify making use of an indexical account, rather than one of names, to account for *de se* attitudes.

This gives us some reason to accept that *de se* attitudes cannot be properly accounted for by names, contra Millikan. This goes some of the way to supporting the claim that non-indexical forms of reference would struggle to have the features that we need to account for *de se* attitudes, which supports (ii).

4.3 Indexing *de se attitudes*

The second alternative to the view that I have defended is that rather than having a SELF concept whose reference is determined each time it is tokened, we instead index the concept to fix its reference.¹⁷ So we may have several demonstratums, which lets us list the demonstratums as:

¹⁶ A reviewer has pointed out that there may be a special class of names that are incapable of failing to refer, thus meeting the condition of *de se attitudes*. Supposing that there is such a class of names, it is still not clear that they will do the required work, as one must realize that the person referred to by that name is oneself. It is not clear that names provide a good account of how this is so, given that Perry could believe that "John is making a mess" without realizing that he is making a mess. I discuss a similar point in section 3.3.

¹⁷ I would like to thank Dr Nat Hansen for this point, in discussion.

first demonstratum, second demonstratum, ... (some of which may be null) as features of a context. We then attach subscripts to our demonstratives and regard the n -th demonstrative, when set in a context, as rigid designator of the n -th demonstratum of the context. Such a rule associates a character with each demonstrative. (Kaplan, 1989, p. 528)¹⁸

To give an illustration, this is the difference between having a case such as “I” and “ I_{NN} ”, where the latter has been indexed to an individual, NN.¹⁹ This view avoids the worry that the names view faced, because it seems that this sort of indexing requires either a perceptual or anaphoric link to the object that it refers to. They will be linked to their referent in *de se* attitudes via the agent’s awareness of themselves and it seems that they will not fail to refer or have the referent be forgotten as names do. But these indexed components also do not seem to be indexical any more (Ball, 2010, p. 14). This indexing view seems to relinquish (ii).

There is, however, an issue with this account. We do not have a guarantee that one is aware that the object that is indexed with “ I_{nn} ” is oneself. To see this, consider again the case in which Perry is searching for the messy shopper, except that in this case, there are mirrors at either end of the counter so that as he pushes his trolley down the aisle he sees a reflection of the messy shopper going down the aisle on the other side (Perry, 1979/2010, p. 371). Perry points and says “I believe that he is making a mess” (Perry, 1979/2010, p. 371). Suppose that here we have a case in which the “he” is indexed to become “ he_1 ”. In this case, Perry is not aware that he has indexed himself as the referent of “ he_1 ”. This shows that indexing oneself is not sufficient to realize that one is referring to oneself.²⁰

For the indexing account to be a useful account of *de se* attitudes one will presumably have to know that they are the NN. In which case, they seem to have to judge that “I am the NN”. Doing this involves a non-indexed first-person concept, the “I”. In which case, they must have already been able to think of themselves in a *de se* way prior to having the indexed first-person concept “ I_{nn} ”. So, the self-awareness of *de se* attitudes cannot be accounted for by mental indices.²¹ Attempts to do so only move the question. So, the prospects for an indexing view of *de se* attitudes do not seem to be promising.

To conclude this section, Argument 1 claims that if the CTM is true, then there cannot be mental indexicals at the level of thought. This is on the basis that a CTM uses a processor that can only react to symbols by taking their syntactic type into account. However, an indexical view would require taking the symbol

18 A subject need not represent this relation to themselves.

19 Peacocke (1981, pp. 189–190) seems to endorse a similar view.

20 Similar problems are possible in the case of names.

21 I would like to thank an anonymous reviewer for their helpful comments in this regard.

qua token into account.²² This seems to suggest that the CTM cannot allow for mental indexicals. In this case, claim (ii) seems to be false when (iii) is true. I have also considered some ways in which one might handle *de se* attitudes in the absence of (ii) but have rejected them as they seem unable to capture what is essential to *de se* attitudes. In the following section I will attempt to show that this threat to (ii) can be avoided without giving up on the CTM.

5. Rejecting Premise (1)

I will deny that premise (1) is true. This premise stated that if there are constituents of the language of thought that have an indexical character, then the CTM is not true. Denying (1) undermines Argument 1 while preserving the CTM. I will argue that the CTM can allow for non-computational processes to occur outside of mental modules. This means that one can allow that computations are not able to treat symbols as anything other than instances of a given type without taking this to conflict with (ii).

To do so I do *not* want to claim that there is a module that deals specifically with *de se* attitudes which then must give an output to a general reasoning centre to get content. Rather, I want to claim that the problem that Argument 1 raises would only apply to those *de se* attitudes when they are processed within a module. This might happen when someone wishes to express an “I” thought, and this must go through a speech processing module, or a syntax module. But such instances need not be problematic. That the module is not able to treat symbols *qua* token does not matter, so long as a general reasoning capacity can do so. In this case, it does not matter that the module cannot treat the tokens of a given type differently; the general reasoning capacity can.

However, this is also not enough. It cannot be the case that the non-modular part of the mind saturates the indexical and removes this context-sensitive part. To do so would be to deny (i) or (ii). In which case, the solution will not have been successful. If we take saturation to be a process of replacing indexicals with

22 A possible line of argument would be that it seems to be strange to have evolved a computational system which has an indexical element that does not change its referent. Why not have a non-indexical element that would not change its reference, in which case we have something that seems to be computationally simpler and does not have any unnecessary complexity (having only a content, rather than a character and a content)? My response to this is to reject the idea that the mind is computational all the way through so that what is computationally simplest is not the deciding factor in our understanding of the mind. Furthermore, it is not obvious that we can account for *de se attitudes* without using indexicals. So, this complexity may be necessary.

non-indexical concepts, then it seems that an account of how an indexical thought can refer without being saturated will be required.²³

5.1 Restricting the scope of the CTM

Here I will suggest that the CTM can accommodate cases of indexicals without requiring that it attempts to infer their content in a computational manner. Instead, the symbol can be moved out of the module (if it was ever there at all) to be dealt with non-computationally. This move requires endorsing a modular theory of mind according to which we have domain-specific modules (i.e., modules that are dedicated to a particular task and kind of information) and a general information processing part of the mind that can deal with other processes that are not dealt with by modules. This general part of the mind will be able to draw on relevant and available information to whatever its task is at that time. It need not operate computationally. I will first give some motivation for this understanding of the CTM. I will then go on to deny that (1) is true.

The rationale for this theory of mind stems from the Frame Problem.²⁴ This is important because the theory of mind proposed above should not be motivated only by Argument 1. That would make the solution *ad hoc*. This problem begins with the idea that when we are making some, typically non-demonstrative, inferences we try to make use of premises that are relevant to that task. A relevant proposition can be understood as one that “if it *were* attended to *would* affect the estimated subjective probability of the belief [derived by this reasoning]” (Fodor, 2010, p. 116). The problem is that this property is not one that can be had in virtue of the syntactic properties of a belief or proposition (Fodor, 2010, p. 124). Relevance cannot be determined on a syntactic basis because the relevance of a given belief will vary depending on the hypothesis that one is seeking to prove but the syntax will remain fixed. This means that computations cannot be responsible for selecting the relevant premises for these sorts of inferences and tasks. So, computations cannot be responsible for these inferences, yet these are inferences that we make. This supports the case for there being a general, non-computational domain of the mind.

We can, however, still accept computational processes that do not have to deal with the Frame Problem. These will have to be performed over areas where the range of possible information is already constrained, so that the selection of

23 This will not follow from all versions of saturation, just the one that I have mentioned here. If one wished to maintain a Kaplanian account of saturation, on which the character interacts with a context of use to provide a content without changing the representation, then the indexical element would not need to be replaced.

24 At least, it stems from Fodor's version of the Frame problem; see Kamermans and Schmitts (2004) for other variations of the Frame problem.

relevant information is not done as the result of a computation (Fodor, 2000, p. 64). Instead, “the information available to perform a task depends on which task it is; and the constraints in virtue of which this is so are ‘architectural’” (Fodor, 2000, p. 63). This has the result that “modules are informationally encapsulated by definition”, so they do not have to select the information that they make use of (Fodor, 2000, p. 63). This means that we can preserve a CTM without having to claim that the mind is computational through and through.²⁵

Now that a modular CTM is on the table, it is possible to give a way of denying (1) and therefore deny that (iii) causes a problem. (iii) can be true while (1) is false for this reason. We can make indexical reference to an individual by making use of a non-modular, and therefore non-computational, system. This can take features of the context of use into account and giving a reference to indexicals that occur in thought on a theory of mind that is still computational.

Here, I want to concede that computations do need to meet the TSC but deny that a CTM needs all its processes to occur computationally. Instead, only the modules operate computationally. So, the idea that we might have thoughts with a character need not be incompatible with the CTM. So, the antecedent of (1) can be true, there can be indexical components of Mentalese, while the consequent is false, there can still be a true CTM. (It may be that all a computational module can do here is provide a semantic constraint, which constrains which propositions could be expressed by a term without determining a proposition; see Harris, 2018, p. 6 for discussion.)

This way of responding to Argument 1 can also avoid predicting errors that we are not, in fact, liable to commit. Recall the worry that indexical symbols in thought would predict that we are liable to errors in equivocating that we do not, in fact, commit. For instance, “That is a cat, that is a dog, therefore that is a cat and a dog”. This prediction can be avoided by avoiding a commitment to the CTM. If the process is not computational then it is not a requirement that all instances of a given type be treated as though they had the same content. So, each instance of “that” can be treated as a token with its own content. So, this account need not predict this kind of error.

In sum, this solution requires a modular account of mind. If the mind were computational through and through, then the CTM would be incompatible with indexicals. Instead, there needs to be some component of the mind that does not function in a computational way. The Frame Problem also gives us reason to think that there are such parts of the mind as this, so this is not an *ad hoc* move.

25 Alternatively, one might want to dispose of the CTM altogether. In this case, we also do not have an inconsistent triad as (iii) ceases to be relevant. However, in this case one needs to propose an alternative theory of mind that can accommodate indexicals.

From here, we can allow that the context can be taken into consideration in determining the content of a proposition.

One response that might be made here goes as follows. I have presented a dilemma between the CTM and *de se* essentialism. To resolve this, I have opted to dismiss the CTM in favour of *de se* essentialism. But why not go the other way and deny *de se* essentialism and keep the CTM? There are a couple of things to say in response to this. One is that I do not think that this is what I have done. I have not argued that the CTM should be disposed of. I have just argued that the scope of computations should be restricted to modules, and that the mind is not only modular. Of course, one might say that this makes little difference. Why not keep a thoroughly computational mind and lose *de se* essentialism? One reason is that the CTM must deal with the Frame Problem, and it does not seem able to do so if it is thoroughly computational. Another is that the arguments in favour of (i) are compelling, and until a flaw is found in them it seems preferable to maintain (i), with (ii), and limit the thoroughness of computation in a CTM.

Another option would be to lose modules and computation entirely and claim that the general processing system does the work formerly attributed to mental modules. Responding to this objection in full would go beyond the scope of this article. However, there are still cases of what appears to be the encapsulation of certain systems, which suggests that a modular account is still possible. For instance, contour interpolation seems to be a modular process (Keane, 2018).

One objection that might be made here is that I have not offered a solution to Argument 1 as the general reasoning processes are difficult to describe or understand in any detail, whereas it is an advantage of the CTM that its processes can be described. In response to this, I would like to emphasize that the move to general processing is motivated by the Frame Problem. While there is a price to pay in terms of which processes we can understand computationally, this is a price that we are already paying.

Also, this understanding of the CTM is not novel and is well known from Fodor, who is one of the chief proponents of the CTM. In his words: “when I wrote books about what a fine thing CTM is, I generally made it a point to include a section saying that I don’t suppose that it could comprise more than a fragment of a full and satisfactory cognitive psychology” (Fodor, 2001, p. 1). A similar claim is that “Even if input systems are domain specific, there must be some cognitive mechanisms that are not” (Fodor, 1983, p. 101). From this we get the idea that the CTM need not be expected to be exhaustive of our understanding of cognition. The TSC still applies to those processes that are computational, but it is a mistake to think that it must apply to all the processes that occur in a CTM.

While this means that we do not fully understand how *de se* attitudes should work, it seems that if we adopt (i)–(iii) then however they work they cannot work

computationally. This does not create an incompatibility with the CTM as it does not need to process *de se* thought computationally. So, it is possible to have (i)–(iii) without the risk of incompatibility. The incompatibility only arises if we assume that a CTM must be computational through and through. We should not make this assumption. Furthermore, my aim was to show how it is possible to maintain a CTM while also allowing for indexical thoughts in the form of *de se* attitudes. I have presented a means of doing that.²⁶ This view differs from that introduced by Argument 1 as indexicals are a possibility on this view.

6. Conclusion

To conclude, I have presented an argument that a CTM can be compatible with indexical symbols, contrary to Argument 1. We can allow for this by denying that the CTM requires that all processes, inferences, and so on, need to be done via computation. This version of the CTM also has independent motivation from the Frame Problem. This suggests that we can have context-sensitive thoughts, in the form of *de se* attitudes, where these attitudes are both essential and essentially indexical, and the CTM. This opens up an important dialectical space for future discussion.

Acknowledgements

I would like to acknowledge the feedback that I received on earlier drafts of this article from the Postgraduate Research group at the University of Reading, Prof. Borg and Dr Hansen. I would also like to thank the feedback from participants of the 2018 Context, Cognition and Communication conference and the anonymous reviewers whose insightful comments greatly improved the quality of this article. Dr Ball was also kind enough to allow the use of his draft. Finally, I would like to thank the University of Reading for funding this research.

References

- ANSCOMBE, G. E. M. (1975) “The First Person.” In S. Guttenplan (ed.), *The Philosophy of Mind: Wolfson College Lectures 1974*, pp. 22–36. Oxford: Clarendon Press.
- BALL, D. (2010) “Can There Be Indexical Thoughts?” Draft, www.st-andrews.ac.uk/~db71/ball_indexical_thoughts.pdf.
- BALL, D. (2015) “Indexicality, Transparency, and Mental Files.” *Inquiry* 58(4): 353–367.

²⁶ This is important as there may be other reasons to think that indexicals play an important role in a CTM, such as playing a role in unifying information from modules, a discussion of which lies beyond the scope of this article.

- BERMUDEZ, J. L. (2017) *Understanding I: Language and Thought*. Oxford: Oxford University Press.
- CAPPELEN, H. and DEVER, J. (2013) *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford: Oxford University Press.
- CARRUTHERS, P. (2002) "The Cognitive Functions of Language." *Behavioural and Brain Sciences* 25: 657–726.
- CATAÑEDA, H.-N. (1968) "On the Logic of Attributions of Self-Knowledge to Others." *Journal of Philosophy* 65(15): 439–456.
- CUMMING, S. (2016) "Names." In E. N. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*. Stanford: Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2016/entries/names/>.
- EVANS, G. (1981) "Understanding Demonstratives." In H. Parret and J. Bouveresse (eds), *Meaning and Understanding*, pp. 280–303. Berlin: Walter de Gruyter.
- EVANS, G. (1982) *The Varieties of Reference*. Oxford: Oxford University Press.
- FODOR, J. A. (1975) *The Language of Thought*. Cambridge, MA: Harvard University Press.
- FODOR, J. A. (2000) *The Mind Doesn't Work that Way*. Cambridge, MA: MIT Press.
- FODOR, J. A. (2001) "Language, Thought and Compositionality." *Mind and Language* 16 (1): 1–15.
- FODOR, J. A. (2010) *LOT2: The Language of Thought Revisited*. Oxford: Clarendon Press.
- HARRIS, D. W. (2018) "Semantics without Semantic Content." *Mind and Language*, <https://semanticsarchive.net/Archive/2IzNzE4Z/DanielHarris-SemanticsWithoutSemanticContent.pdf>.
- KAMERMANS, M. and SCHMITTS, T. (2004) "The History of the Frame Problem", <https://staff.fniwi.uva.nl/b.bredeweg/pdf/BSc/20032004/KamerMansSchmits.pdf>.
- KAPLAN, D. (1989) "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals." In J. Almog (ed.), *Themes from Kaplan*, 481, 564. Oxford: Oxford University Press.
- KEANE, B. P. (2018) "Contour Interpolation: A Case Study in Modularity of Mind." *Cognition* 174: 1–18.
- KUCZYNSKI, J.-M. (2007) *Conceptual Atomism and the Computational Theory of Mind: A Defense of Content-Internalism and Semantic Externalism*. Amsterdam and Philadelphia, PA: John Benjamins.
- LEWIS, D. (1979) "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88: 513–543.
- MCDOWELL, J. (1984) "De Re Senses." *Philosophical Quarterly* 34(136): 283–294.
- MILLIKAN, R. (1990) "The Myth of the Essential Indexical." *Nous* 24(5): 723–734.
- MÜLLER, V. C. (2008) "Representation in Digital Systems." In A. Briggles, K. Waelbers and P. Brey (eds), *Current Issues in Computing and Philosophy*, pp. 116–121. Amsterdam: IOS Press.
- NINAN, D. (2016) "What Is the Problem of *De Se* Attitudes?" In M. Garcia-Carpintero and S. Torre (eds), *About Oneself: De Se Attitudes and Communication*, ch. 4. Oxford: Oxford University Press.
- PEACOCKE, C. (1981) "Demonstrative Thought and Psychological Explanation." *Synthese* 4(2 Demonstrative and Indexical Reference, Part II): 187–217.
- PERRY, J. (1977) "Frege on Demonstratives." *Philosophical Review* 86(4): 474–497.
- PERRY, J. (1979) "The Problem of the Essential Indexical." In A. P. Martinich (ed.), *The Philosophy of Language*. 5th ed. Oxford: Oxford University Press.

- PERRY, J. (1986) "Thought without Representation." *Proceedings of the Aristotelian Society, Supplementary Volumes* 60: 137–151.
- PINKER, S. (1994) *The Language Instinct: How the Mind Creates Language*. New York, NY: HarperCollins Publishing.
- RECANATI, F. (2004) *Literal Meaning*. Cambridge: Cambridge University Press.
- RECANATI, F. (2012) *Mental File*. Oxford: Oxford University Press.
- RECANATI, F. (2017) "Cognitive Dynamics: A New Look at an Old Problem." In M. Ponte and K. Korta (eds), *Reference and Representation in Thought and Language*, pp. 179–194. Oxford: Oxford University Press.