



HAL
open science

Dynamic Files, Transparency, and Reference Failure.

François Recanati

► **To cite this version:**

| François Recanati. Dynamic Files, Transparency, and Reference Failure.. 2015. hal-01104735

HAL Id: hal-01104735

<https://hal.science/hal-01104735v1>

Preprint submitted on 23 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamic files, transparency, and reference failure

Response to Onofri, Ninan, and Ball

François Recanati

I

Mental files do a number of things for us, corresponding roughly to the roles which Frege assigned to ‘senses’. They determine the reference of expressions : an expression refers to what the file associated with it refers to (at the time of tokening). They account for cognitive significance : Frege’s Constraint states that if a rational subject can believe of a given object both that it is F and that it is not F (as happens in so-called ‘Frege cases’), then the subject thinks of that object under distinct modes of presentation (distinct files). They also account for coreference *de jure*: if two token singular terms **a** and **b** are associated with the same file, it is presupposed that they corefer (if they refer at all) and ‘trading upon identity’ (TI) becomes valid : one can move from ‘**a** is F ’ and ‘**b** is G ’ directly to ‘there is an x which is F and G ’, without needing to invoke an identity premiss.

Onofri argues that mental files can’t simultaneously play the FC role (satisfying Frege’s Constraint) and the TI role (enabling coreference *de jure*). To play the first role, they must be fine-grained, but to play the second role they must be coarse-grained, given that cross-modal and cross-temporal integration of information licenses TI. For example, if I remember that a certain object was F and, upon encountering it again, notice that it is G , I can infer : something that was F is now G . This is TI, yet the mental file deployed in the first premiss (to the effect that the remembered object was F) is a memory file, while the mental file deployed in the second premiss (to the effect that the seen object is G) is a perceptual file. These files are distinct types of file, and we can easily imagine Frege cases involving them, i.e. cases in which the subject does not realize that the object he sees is the same he previously encounters and still remembers. Or think of a case in which I see and touch a certain glass : if cross-modal integration of information occurs as it should, I will trade upon identity and infer, from ‘that (seen) glass is yellow’ and ‘that (touched) glass is cold’, that the glass is yellow and cold ; but again, one can easily imagine Frege cases in which distinct files are deployed, one for each modality, without the identity being presupposed. The subject might wonder whether the yellow glass he sees is the cold glass he touches. So Frege’s constraint dictates that there be two distinct modes of presentation, while the subject’s trading upon identity shows that, for him, there is only one file, which integrates visual and tactile information. So : one file, or two files ? It seems that two (fine-grained) files are needed to the FC role, while a single, coarse-grained file is required to account for TI. The suggestion that a single entity accounts for both FC and TI is therefore dubious and the whole framework threatens to collapse.

Ninan comes to a similar conclusion, based on a different type of example :

Suppose that, at time t_1 , I look at my watch and think to myself, *It's 1:00pm now* . Later (at t_2) I come to think, *Actually, it wasn't 1:00pm then; I forgot to change my watch after the flight*. Intuitively, I changed my mind between t_1 and t_2 about whether it was 1:00pm at the time in question. But my initial thought employs a *now* -

file μ_1 , while my later thought employs a *then*-file μ_2 . According to Recanati, μ_1 and μ_2 are distinct files or modes of presentation. But then why should this count as a change of mind? Normally, if I ascribe a property to an object o while thinking of it under mode of presentation m_1 and then later ascribe an incompatible property to o while thinking of it under a distinct mode of presentation m_2 , I do not thereby count as changing my mind. For example: I might think in the morning that Phosphorus is bright, but then think in the evening that Hesperus is not bright. If I am employing two distinct modes of presentation on these two occasions, this does not count as a change of mind. The reason for this seems to be that the two contents in question are not incompatible in the appropriate sense. (pp. 2-3)

The subject changes her mind only if she thinks of the object under the same mode of presentation throughout; otherwise there is no conflict, hence no change of mind has to be posited. But if the mode of presentation is the same, we can no longer equate modes of presentation with files *and* maintain that files are based on fine-grained ER relations to the object of thought; for in this case as in the others, the ER relation changes while the mode of presentation remains one and the same.

As Onofri and Ninan acknowledge, this problem is not new. It traces back to Evans' discussion of cognitive dynamics in 'Understanding Demonstratives' and *The Varieties of Reference*. I devote a chapter to the problem in *Mental Files*, and I offer the following solution. A file can be *converted into* a distinct, more inclusive file based on more ER relations than the initial file. Through incremental conversion, files grow new information links. Trading on identity across deployment of distinct ER relations is made possible by the *composite* nature of the inclusive file. So in cross-temporal or cross-modal cases, a single file is used in the train of thought, and that file, in contrast to the initial, pre-conversion file, is based on several ER relations.

Onofri thinks the solution does not work, for the revised theory faces a dilemma. Is the post-conversion file (the *inclusive file*) the same file as the pre-conversion file (the *initial file*)? If the answer is positive, files become coarse-grained and cannot play the FC-role. If the answer is negative, the mental file account of coreference *de jure* is made redundant. What makes two token singular terms coreferential *de jure* is no longer the fact that they are associated with the same file, but the fact that the (distinct) files they are respectively associated with are related through the mechanism of conversion. It is the mechanics of conversion that does all the work. As Onofri puts it,

On this account, what is it that really explains the rationality of trading on identity? Clearly, not that the same file is involved in the inference! On the contrary, the correct explanation will now be as follows: A has produced a distinct file B through conversion, and this makes it the case that the inference is rational (perhaps together with other factors). For our purposes, it doesn't matter whether this explanation is correct. What matters is that all the explanatory work is done by the mechanism relating the files, rather than their alleged sameness. It would now be unnecessary to include the assumption 'A and B are the same file' in a complete explanation of trading on identity: file sameness has become explanatorily useless. (p. 10-11)


Onofri concludes that my account 'faces the following dilemma: either the files involved in trading-on-identity inferences are literally the same, in which case files are not modes of presentation; or they are distinct but related through cognitive mechanisms like conversion, in which case the appeal to file sameness plays no role in explaining rational inferences' (p. 11).

But I think the objection rests on a confusion. In TI the *de jure* coreferential

occurrences correspond to deployments of *the same file*, namely the inclusive file (post-conversion). So TI is still accounted in Fregean fashion, by appealing to sameness of sense (or sameness of file). This is compatible with the fact that the single file which is deployed twice in the TI train of thought is an inclusive file resulting from the conversion of an initial file, and therefore distinct from that initial file. So I maintain that the pre- and post-conversion files are *distinct* files, but that does not entail the unwelcome consequence advertised by Onofri : that ‘appealing to sameness of file/mode of presentation now becomes explanatorily useless in accounting for trading on identity’ (p. 10). The files which must be ‘the same file’ to preserve the account of TI in terms of file identity are *the files associated with the two token singular terms* in the TI train of thought. The terms are coreferential *de jure* because they are associated with the same file. That file is an inclusive file resulting from the conversion of an initial file, and therefore distinct from that initial file. I conclude that the distinctness of the pre- and post-conversion files does not make the appeal to file identity redundant in my account of TI.

The source of Onofri’s confusion is the assumption that the pre- and post-conversion files *are* the files respectively associated with the (*de jure* coreferential) token singular terms in cross-modal or cross-temporal TI-licensing trains of thought. If this were true, Onofri’s objection would go through. But I reject the assumption : I deny that the initial file *is* deployed in cross-modal or cross-temporal TI. The illusion that it does stems from the diachronic nature of the reasoning in the sort of example Onofri focuses on. Thus consider the following case. Remembering a certain object from a previous encounter, and remembering it as *F*, I think : ‘**that thing** was *F*’. Then I look around, recognize the object in the vicinity, and think : ‘**it** is *G* now’. Onofri would presumably say that the first premiss involves a deployment of a pure memory file (the initial file : a ‘memory demonstrative’), while the second premiss involves a deployment of the inclusive file (a ‘recognitional demonstrative’ resting on memory *and* perception). On that way of understanding the example, the two token singular terms are coreferential *de jure* (TI is licensed), yet they are associated with distinct files, namely the initial file and the inclusive file it converts into when the subject recognizes the object. So if we accept that there are such examples, Onofri’s objection goes through. But I deny the legitimacy of that ‘diachronic’ understanding of the TI train of thought. For logic purposes, a train of thought has to be construed as synchronic (Kaplan 1989b : 584-85). Of course, we are free to stipulate that the subject had the initial, memory-based thought, ‘**that thing** was *F*’, at time t_1 , before perceiving and recognizing the remembered object, and only later – at t_2 – thought ‘**it** is *G* now’. But in this case I hold that the train of thought *occurring* at t_2 involves two premisses : one is the overt premiss ‘it is *G* now’, where ‘it’ is associated with the inclusive file, and the other one is what the subject has retained of the initial thought. What the subject has retained is not the initial thought itself, but a variant that *results from updating the initial thought (the thought held at t_1) through conversion of the initial file into the inclusive file*. Updating here is *necessary* to retention. The initial file is no longer available at t_2 , so the thought in which it occurs (the thought held at t_1) cannot be directly recruited into the train of thought that takes place at t_2 . At t_2 , after recognition, the initial file has been converted into the inclusive file, so only the inclusive file remains and can feature in the ongoing train of thought.

If we make the associated files explicit by adding subscripts to the singular terms, we get the following representation of the TI-train of thought in the example I have just discussed:



conversion



- a) JOHN* IS NOT JACK
- b) JOHN* IS IN THE ROOM
- c) THERE IS ONLY ONE PERSON IN THE ROOM
- d) JACK IS NOT IN THE ROOM

To sum up : In diachronic cases, updating is a prerequisite for enrolling a previous thought into the current train of reasoning. Once updating has been factored in, there no longer is any temptation to regard the premisses of the reasoning as involving distinct mental files or modes of presentation. This provides a solution to the difficulty raised by Ninan in the passage I already quoted :

Suppose that, at time t_1 , I look at my watch and think to myself, *It's 1:00pm now* . Later (at t_2) I come to think, *Actually, it wasn't 1:00pm then; I forgot to change my watch after the flight*. Intuitively, I changed my mind between t_1 and t_2 about whether it was 1:00pm at the time in question. But my initial thought employs a *now* -file μ_1 , while my later thought employs a *then* -file μ_2 . According to Recanati, μ_1 and μ_2 are distinct files or modes of presentation. But then why should this count as a change of mind? Normally, if I ascribe a property to an object o while thinking of it under mode of presentation m_1 and then later ascribe an incompatible property to o while thinking of it under a distinct mode of presentation m_2 , I do not thereby count as changing my mind. (pp. 2-3)

Change of mind occurs only when there is diachronic disagreement between the subject and his former self. Disagreement requires a common content for the parties to disagree about, and that common content is reached through conversion in diachronic cases. I analyse Ninan's example as follows :

t_1 : a°) *It's 1:00pm now*

conversion



t_2 : a) *It was 1:00pm then*
b) *Actually, it wasn't 1:00pm then; I forgot to change my watch after the flight*

There is a clear inconsistency between the updated thought 'it was 1 : 00 then', which results from converting the file in the initial thought 'It's 1 : 00 now', and the new thought which supersedes the updated thought when the subject realizes that he forgot to change his watch: 'it *wasn't* 1 : 00 then'. The contents of (a) and (b) are 'incompatible in the appropriate sense' (Ninan p. 3) and they give rise to genuine disagreement : the subject has changed his mind.

II

As Ninan acknowledges, it is not the same thing to think of a time as *then* and to think it as *now*. John Perry pointed out that the actions one takes as a result of thinking one thing (e.g. 'the meeting starts now') are systematically different from the actions one takes as a result of thinking the other ('the meeting started then'). Given the constitutive connection between thought and action, we need to acknowledge fine-grained thought constituents which distinguish *now*-thoughts from *then*-thoughts even when the time referred to is the same.

Mental files based on specific ER relations are meant to play that role (the FC role). Yet to account for cross-temporal cases we have to make room for coarse-grained modes of presentation enabling the mixing of temporal perspectives (e.g. the simultaneous exploitation of memory and perception). Coarse-grained modes of presentation correspond, in my framework, to *composite files based on multiple ER relations* (and generated through incremental conversion). So it is true that the FC role requires fine-grained files, while the TI role demands coarse-grained files, but thanks to the mechanism of incremental conversion we can have both ; and it is a noteworthy feature of my account that it appeals to both types of file.

Since, to solve the problem raised by cross-modal and cross-temporal TI, I have to make room for coarse-grained files, based on multiple ER relations, in addition to the fine-grained files based on specific ER relations, it is natural to wonder whether one might not do everything with the coarse-grained files, and dispense with the fine-grained files altogether. Instead of starting with fine-grained files based on specific ER relations, and generating the coarse-grained files through the mechanism of incremental conversion, we might start from the other end. At the other end we find *maximally inclusive* files. I call them *encyclopedia entries* : files that are hospitable to information derived through *any* ER relation available to the subject. Even though I describe them as based on a ‘higher-order’ ER relation to the reference and therefore as still in line with my indexical model (which construes all files as based on ER relations), I agree with Ball (and Papineau) that they are not *practically* indexical : they are independent of any particular contextual relation to the reference and are therefore eminently stable, in contrast to the indexical files based on specific ER relations. They are what Perry calls ‘detached files’. Encyclopedia entries are more like names than they are like indexicals (and indeed, I want to say that encyclopedia entries are the type of file associated with proper names). The suggestion, then, is that we should start with these maximally inclusive files, independent of any specific contextual relation, and distinguish different stages in their evolution. This is how Papineau describes the situation. When I first encounter an object, I open an encyclopedia entry (a stable file) for that object. For the time being, my only relation to it may be demonstrative/perceptual, but soon I will be in a position to gain information from it through additional ER relations (e.g. testimony). So-called demonstrative files, on this view, are nothing but stages in the development of namelike files – the only files we need. No conversion needs to take place, because the encyclopedia entry is there from the start.²

Both Ball and Ninan concur with Papineau. They think we can do everything with lasting files based on indefinitely many ER relations. Ball shows in some detail that a ‘names-only’ system can be used to represent indexical information (about time, in his example), without any need for a specialised indexical file. If the two systems (the names-only system

² Here is a verbatim quote from Papineau : ‘When I first encounter some item perceptually, I open a potentially permanent file in which to accumulate information about that item. That file outlasts the original encounter, and the same file is reactivated when I remember the relevant item or re-encounter it. The information earlier acquired is thus automatically available on those later occasions, and can be added to when new facts are acquired, without any need for any multiplication of files. (See Papineau 2006.) On this view, the files that we open on first perceptual encounters, and in general on coming into any contact with any new item of thought, are *name-like*. They are designed to be permanent repositories of information about the item in question, and are not dependent on any particular sources of information about that object. In this respect they are akin to Recanati’s ‘encyclopaedic’ files, whose function is to gather information about some referent from whatever sources offer themselves.’ (Papineau 2013 : 168-69)

and the indexical system) are equivalent, as Ball suggests, why not go for Papineau's simpler view (only encyclopedia entries)? As Ninan puts it, 'once we have the stable relations and piles, what need is there for unstable relations and files proper?' (p. 5)

Onofri mentions that I have myself taken a step in that direction (in addition to the distinction between files and piles already in *Mental Files*). In 'Mental Files : Replies to My Critics', in the special issue of *Disputatio* dedicated to my book, I introduce a numerical index on files which corresponds to the fact that they are related by conversion.³ Co-indexing means that two files belongs to the same *sequence of files*, what we may call a 'dynamic file'. Why, then, should we not try to do everything with dynamic files (and their stages)? Note that encyclopedia entries are intrinsically dynamic : their function is to exploit all the ER relations available, so they are designed to grow new information links whenever possible. Instead of describing conversion as an operation external to files, we can take a more dynamic approach and view incremental conversion as internally generated in virtue of the function of encyclopedia entries.

I am all in favour of such a dynamic approach. Consider a descriptive name like Evans' 'Julius'. It is associated with a descriptive file, a type of file which has interesting properties and is worth investigating in its own right. 'Neptune' also started its life as a descriptive name, but acquaintance with the reference has led to conversion of the descriptive file initially associated with the name into a regular encyclopedia entry based on ER relations to the reference. So the descriptive file at the origin of the name may be viewed as a stage in the dynamic evolution of an encyclopedia entry. As Stalnaker pointed out, the same dynamic process may well happen with 'Julius' (Stalnaker 2003 : 198-99). Or consider the deferential file a person might associate with a name overheard in a conversation : when the person hears more about the referent of the name the deferential file will be converted into a regular encyclopedia entry. These examples show that we need an analysis with two levels : the local, static level, and the global, dynamic level. The global level is essential for understanding informational updating, but also for understanding communication, if it is accepted that dynamic files (sequences of files) can be interpersonal as well as intrapersonal (Prosser forthcoming).⁴

So I agree that we need dynamic files, and I agree that encyclopedia entries are intrinsically dynamic. As a result, I have no real objection to Papineau's picture, which Ball and Ninan endorse. What I reject is only the idea that we should *dispense with fine-grained files and with operations like conversion*. I think the best way to account for dynamic files themselves is by analysing them as sequences of fine-grained files. We need the fine-grained files, in particular, to make sense of dynamic phenomena like fusion and fission of files.

As Prosser points out,

The problems resulting from fission and fusion for modes of presentation are very similar to those that arise for psychological continuity theories of personal identity. For every type of solution on offer for personal identity there will be a corresponding type of solution for modes of presentation. My own view is that the best solution is given by stage theory (Sider 1996, 2001; Hawley 2001). (Prosser forthcoming : 18)

Assuming that persons are both dynamic continuants (space-time worms) and stages, what is basic? Because the dynamic view which takes continuants as basic has to be unpalatable

³ The index in question has nothing to do with the 'indexed files' I talk about in sections 3 and 4.

⁴ The analogy between informational updating and communication has been noted several times in the literature. See e.g. Gibbard 2012 : Appendix 1.

consequence that distinct individuals can be located in the same place at the same time (Lewis 1976), Sider suggests that stages should be construed as basic (Sider 1996, 2001 : chapter 5). That is the view I take in the case of mental files. I take modes of presentation (what plays the FC role) to be mental files in the fine-grained sense (file stages, as we may more aptly call them),⁵ and I take such files to undergo dynamic operations such as conversion, incremental conversion (absorption), fusion (file merging), and fission (file splitting). Sequences of files related by such operations are dynamic files. These operations do not preserve file identity in the strict, Leibnizian sense: in contrast to identity, dynamic continuity between files is an intransitive relation, just like dynamic continuity between person stages.

Let us consider an example. At t_1 , I see a certain object and open a demonstrative file DEM₁ about it : ‘that thing’. At t_2 , the object I have been in contact with since t_1 disintegrates, but the demonstrative file persists because, as a result of taking a certain drug, I hallucinate the continued presence of the object. There is an issue whether or not the file continues to refer after t_2 . I think it does not, since the presupposition of the demonstrative file are doubly violated.⁶ At t_3 a doubt occurs to me and I wonder whether the object I remember seeing at the beginning of the episode (t_1) is really the same as the object that I (mistakenly) take myself to seeing at t_3 . Rational doubts about identity necessarily involve two distinct mental files, and here the two files result from splitting DEM₁, which is replaced by a memory demonstrative (referring to the object initially seen) and a perceptual demonstrative (purporting to refer to the object currently seen). So there are three files in this example: DEM₁, the demonstrative file opened at t_1 and maintained until t_3 (despite the disappearance of the object at t_2) ; MEM, the memory demonstrative file deployed at t_3 ; and DEM₂, the new demonstrative file, also deployed at t_3 . Both MEM and DEM₂ result from splitting DEM₁ at t_3 . At t_1 , DEM₁ referred to the object o the subject was then in visual contact with. At t_2 , DEM₁ did not refer to anything. At t_3 , DEM₁ went out of existence, and two new files were deployed : MEM inherits the reference of DEM₁ at t_1 , so it refers to o (and inherits the information in DEM₁ at t_1), while DEM₂ fails to refer because the subject is hallucinating.

It seems to me that there is as much dynamic continuity between DEM₁ (as deployed between t_2 and t_3 , while hallucinating) and DEM₂ as there is between DEM₁ (as deployed at t_1) and MEM. Such dynamic continuity can’t ground identity, for if it did, we would have both DEM₁ = DEM₂ and DEM₁ = MEM. That would entail that MEM = DEM₂, which is impossible since rational doubts about identity (like the doubt harbored at t_3) necessarily involve two distinct mental files. This suggests that we should resist Ninan’s claim (following Frege in ‘The Thought’ and Evans in ‘Understanding Demonstratives’) that

the mode of presentation associated with today on day d is the same as the mode of presentation associated with yesterday on the day after d . If modes of presentation are mental files, then the mental file involved in a *today* –thought on d should be the same file involved in a *yesterday* –thought on the following day. (Ninan, p. 4)

The modes of presentation are ‘the same’ only in a weak, dynamic sense which does not correspond to Leibnizian identity. In the strict sense they are not the same : they are distinct

⁵ File stages, or static files, may be fine-grained or coarse-grained ; so talk of ‘fine-grained file’ (in contrast to dynamic file) is inappropriate.

⁶ The file presupposes both that there is an object to which the subject is attending, and that whatever the subject is attending to is the same thing she has attended to in previous deployments of the file. Both presuppositions are violated in this example. If the object had been ‘switched’ at t_2 , only the second presupposition would be violated, but the subject would still fail to refer (through confusion).

modes of presentation even though they are stages of one and the same evolving dynamic file.

Ninan could respond by pointing out that not all dynamic relations between stages matter ; what matters are only those that support ‘rational relationships’. As he puts it,

The notion of a mode of presentation was introduced in part to describe what we might call rational relationships between beliefs. Since a rational agent can believe that Hesperus is bright without believing that Phosphorus is bright, these beliefs must have different contents, and so the two modes of presentation for Venus must be different. Here we focus on the rational relationships between the beliefs of an agent at a single time. But we can also consider the rational relationships between two beliefs of an agent that are held at different times. The reason for thinking that my belief (*It is now 1:00pm*) at t_1 and my belief (*It was not 1:00pm then*) at t_2 involve the same mode of presentation is precisely that there is rational relationship between these two beliefs: they are incompatible, which partly explains why I count as having changed my mind between t_1 and t_2 . (Ninan 2014 : 4)

If we focus on that subset of dynamic relations that qualify as ‘rational relationships’ in Ninan’s sense, it is not so clear that they can’t ground identity for modes of presentation. Arguably, in my example from the previous section, there is such a relationship between DEM₁ at t_1 and MEM, but *not* between DEM₁ at t_2 and DEM₂. So perhaps we can say, on that basis, that DEM₁ (at t_1) and MEM are the same mode of presentation, while DEM₁ (at t_2) and DEM₂ are not the same mode of presentation (thus blocking the unwelcome consequence that MEM = DEM₂).

Now what is the special relationship which holds between DEM₁ at t_1 and MEM but not between DEM₁ at t_2 and DEM₂? It is something like this : A memory file inherits its reference from the perceptual file it derives from ; so they are bound to corefer (if they refer at all). Let us call that coreference *de jure* between file stages (CDJ for short). On that basis, we could say, following Ninan (and Evans), that the perceptual file and the subsequent memory file count as the same mode of presentation because they bear the appropriate rational relation (CDJ) to each other.

The situation is different with DEM₁ and DEM₂ because DEM₂ results from the *fission* of DEM₁. Whenever fission (or, for that matter, fusion) is involved, there always is the possibility that one of the twin files (or both) refer while the inclusive file fails to refer (Ball 2014 : 17). So the two files are *not* bound to corefer if they refer at all. Fusion and fission do not support CDJ. Or perhaps we should distinguish a strong and a weak form of CDJ :

Coreference *de jure* (strong): the two occurrences corefer if either of them refers.

Coreference *de jure* (weak) : the two occurrences corefer if both of them refer.

In cases of fusion and fission, the files only exhibit the weak form of coreference *de jure* : the two files are bound to corefer if they *both* refer. There is the possibility that one of the file refers while the other fails to refer, but it is ruled out that the two files both refer while referring to different things.

If this is on the right track, then we need three levels . We need the fine-grained files (or, rather, the stages) and the dynamic operations on them : conversion, fusion, fission etc. This gives us the notion of a dynamic file. Then we can distinguish two kinds of dynamic files : those which do and those which do not support ‘rational relationships’. The files that do may be called the proper dynamic files. It is at the level of the proper dynamic file that we can try to capture the notion of a dynamic mode of presentation which Ninan and Evans argue for, and we can do so in terms of strong CDJ. Strong coreference *de jure* between stages is a transitive relation, so it can ground the identity of (dynamic) modes of presentation of the sort

Evans and Ninan are after. In contrast, weak CDJ is not transitive. For example, in fusion or fission, an inclusive file A/B may bear weak CDJ to both of the twin files A and B, while A and B do not bear weak CDJ to each other (Pinillos 2011, Recanati 2012).

But it's not absolutely clear that dynamic files *ever* exhibit strong CDJ. There are five types of dynamic file to consider : (i) deployments of the same file at different times; (ii) conversion of one file into another (*now/then, here/there, today/yesterday ...*); (iii) incremental conversion (multiplication of ER relations); (iv) fusion ; (v) fission. Fusion and fission only support weak CDJ, as we have seen. There always is the possibility that one of the twin files (or both) refer while the inclusive file fails to refer. The same thing is true of incremental conversion. The incremented file may fail to refer through violation of the uniqueness presupposition (to the effect that the various ER relations converge on the same object), even if the initial file (pre-incremental conversion) succeeds in referring. More interestingly, in case (i) too we find weak rather than strong CDJ. That is what the example of DEM₁ shows : at t_1 DEM₁ refers to o , while it fails to refer at t_2 . So two deployments of the same file at different times are not bound to corefer if they refer at all ; they are only bound to refer to the same thing if they *both* refer (weak CDJ).

It seems, then, that simple conversion, illustrated by Ninan's example, would be *the only case* of a proper dynamic file, exhibiting strong CDJ. It would be an *exception* to the generalization that dynamic files only support weak CDJ. But do we really need such an exception ? One may be skeptical of any account that creates too big a gap between simple and incremental conversion. Is it true that simple conversion (e.g. Ninan's *now/then* example) supports strong CDJ while incremental conversion only supports weak CDJ ? I am not sure. It is well-known that simple conversion supports strong CDJ only if the thinker has managed to 'keep track of time' (Kaplan 1989a, Evans 1981, Perry 1997). Rip van Winkle is the case of a thinker who has not kept track of time, and whose temporal file about the previous day is confused. But the requirement that the thinker keep track of time introduces an artificial asymmetry between simple conversion and incremental conversion. The sort of confusion which may give rise to reference failure in cases of incremental conversion (or fusion) are cases in which the uniqueness condition is not satisfied : the various ER relations do not converge on the same object. That is why incremental conversion only supports weak CDJ. But if we ruled out such cases by adding a condition analogous to the time-tracking requirement, e.g. the requirement that the subject 'keep track' of the object throughout the dynamic operation, then incremental conversion would turn out to support strong CDJ as well.

I conclude that it is not obvious that *any* dynamic file (without the help of extra conditions, like the time-tracking requirement) supports strong CDJ. Perhaps only *synchronic deployments of the same file* do so. If that hypothesis is correct, then we cannot use proper dynamic files (based on a transitive relation) to ground the identity of dynamic modes of presentation, because there are no 'proper' dynamic files. Dynamic files only support weak CDJ, and weak CDJ is not a transitive relation.

The issue is far from settled, and I don't want to claim too much for my argument. But one conclusion at least can be drawn : to proceed any further — to ask the questions that need to be asked in order to settle the issue — we need a framework which makes it possible to represent the stages and the dynamic operations on them. Stages are basic at least in that sense.

III

The Fregean appeal to modes of presentation is meant to protect the subject's rationality. A subject may believe of an object both that it is F and that it is not F , and still be rational. From

an *internal* point of view there is no contradiction, if the subject has two distinct modes of presentation for that object (and does not realize they both determine the same object). Of course, this move only makes sense if the subject has internal or, better, *transparent* access to the modes of presentation. Reference is opaque : whether or not the subject successfully refers, and to what, depends to a large extent upon environmental matters the subject has no control over. The whole point of positing senses or modes of presentation is that the subject is supposed to have transparent access to *them*. What that means is that the thinker knows when two file deployments ('occurrences') are deployments of the same file, and when they are deployments of distinct files. This is in contrast to reference : the subject need not be aware that she is referring to the same object twice (or to distinct objects) when she is.

Since senses (files) are transparent, and they determine reference (at a time), it follows that whenever the same file is deployed twice in the subject's thought, she has a priori knowledge that she refers to the same object twice if she refers at all. This a priori knowledge of *conditional* coreference (what Perry 2012 calls 'coco-reference') is compatible with externalism and the opacity of reference. The subject does not know whether or not her files refer, and to what (opacity of reference), but, in virtue of the transparency of sense, she knows that two synchronic deployments of the same file are bound to corefer if they refer at all (transparency of sense). In chapter 10 of the book, I discuss (and rebut) alleged counterexamples to the transparency principle based on the phenomenon of 'slow switching'.

Ball rejects the principle of transparency and provides a counterexample, based on the phenomenon of 'recognized equivocation':

I judge:

(1) Barry Smith is a philosopher of mind and language who works in London.

I then judge:

(2) Barry Smith is a ontologist who works in Buffalo.

I infer from these premises (and no others) that someone is a philosopher of mind and language who works in Buffalo. It initially seems obvious to me that the inference is valid. (I suppose that the philosopher in question must split his time between the two universities.) But upon further reflection, I decide that no single person could have distinguished himself in two fields in this way. I judge that my information must be about two distinct people who share the same name. I conclude that my initial inference was invalid. I judge:

(3) Barry Smith is not an ontologist

(intending my use of "Barry Smith" to be de jure co-referential with the first premise of my initial inference). I then infer (from (1) and (3)) that some philosopher of mind and language is not an ontologist. I also judge:

(4) Barry Smith does not work in London

(intending my use of "Barry Smith" to be de jure co-referential with the second premise of my initial inference). I infer (from (2) and (4)) that some ontologist does not work in London. (Ball, pp. 15-16)

This example shows two things, according to Ball. First, transparency fails. ‘Knowledge’ of conditional coreference is posited to account for the train of thought which enables the subject to trade upon identity and conclude from (1) and (2) that someone who is a philosopher of mind and language lives in Buffalo. But that is only an illusion of knowledge (according to Ball) : the subject is actually mistaken and comes to realize that she is. She comes to realize that the two tokens of ‘Barry Smith’ in (1) and (2) are not actually coreferential : one refers to the London philosopher of mind and language, the other one refers to the Buffalo ontologist. Second, ‘the transitivity of de jure co-reference fails, and (as I will argue) fails in a way that is difficult for Recanati to accommodate’ (Ball, pp. 16). Ball writes :

It seems – at least by Recanati’s lights – that the occurrence of “Barry Smith” in (1) is de jure co-referential to the occurrence in (2), that the occurrence in (2) is de jure co-referential to the occurrence in (4), that the occurrence in (3) is de jure co-referential to the occurrence in (1), but that the occurrence in (3) is not de jure co-referential to the occurrence in (4). That is, it seems to be a case in which the transitivity of de jure co-reference fails, (pp. 15-16)

I think both objections misfire. As for transparency, I maintain that the subject has a priori knowledge of conditional coreference when she makes the judgments (1) and (2). As a result, her TI train of thought is valid, contrary to what Ball claims (p. 17). It is true that the subject realizes she was in error when she went through the initial TI train of reasoning (moving from (1) and (2) to the existential conclusion about a philosopher of mind and language living in Buffalo). But the error consists in deploying a file that fails to refer, namely the inclusive file fed with information from the two distinct Barry Smiths. That file rests on a confusion of two distinct individuals, and such files fail to refer because the ER relations on which the file is based do not converge on a single object. Confusion may be innocuous in contexts in which there is a dominant source for the relevant information in the file, but in Ball’s example it is not innocuous : the file is equally about Barry Smith the London philosopher of language and mind and Barry Smith the Buffalo ontologist, so it is about neither and fails to refer. It follows that the two occurrences (the two deployments of the confused file in (1) and (2) respectively) exhibit strong CDJ. They do ‘corefer if they refer at all’ : the condition is trivially satisfied since the antecedent of the conditional is false — the confused file which is deployed twice does not refer. TI is therefore valid, even though the conclusion (that there is a philosopher of mind and language who lives in Buffalo) is false.

The subject’s mistake, in this sort of case, is architectural (one file instead of two), and the subject’s rational reaction is to revise the architecture by splitting the file. That is what happens when the subject changes his mind (after the initial TI train of thought). He splits the initially inclusive file BARRY SMITH, which does not refer, into two distinct files : BARRY SMITH₁ and BARRY SMITH₂. It is these files that are deployed in (3) and (4). It follows that, contrary to what Ball claims, there is no (strong) coreference de jure between the occurrence in (1) and the occurrence in (3), or between the occurrence in (2) and the occurrence in (4). The subject does not take, nor intend, these occurrences to corefer (if they refer at all). He is aware that the initial, inclusive file does not refer, because it rests on confusion, but he takes the new files deployed in (3) and (4) to refer. When Ball says that the subject intends the occurrence of ‘Barry Smith’ in the judgment (3) to corefer de jure with the occurrence in (1), in the strong sense (the sense that matters to transitivity), this can only be true if the subject is *updating* his initial judgment (1) by substituting the new file BARRY SMITH₁ for the inclusive

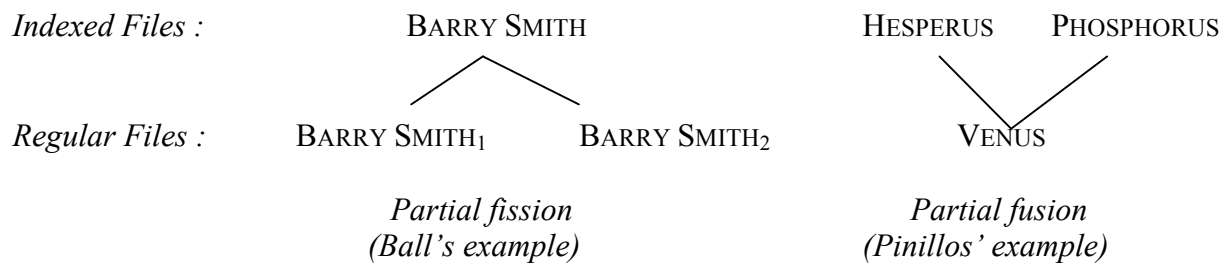
file that was initially deployed when the judgment was made.⁷

The subject himself might explain the situation as follows. (I use subscripts to indicate the files that are deployed on every occasion on which the name ‘Barry Smith’ is used ; ‘0’ stands for the initial, inclusive file.)

The Barry Smith story

When we discovered that Barry Smith₀ was both a philosopher of mind and language and an ontologist, and that he₀ lived both in London and Buffalo, we started wondering whether he₀ was a single person or whether there were two philosophers bearing the same name. Now, of course, we know that Barry Smith₁ is not an ontologist and that Barry Smith₂ does not live in London.

In this discourse, the three distinct files are synchronously deployed. The subject has split the initial file BARRY SMITH into two distinct files BARRY SMITH₁ and BARRY SMITH₂, but he has retained the inclusive file which he uses as an ‘indexed file’ to represent the perspective of his former self. (More on this in the next section.) As Ball points out, this is similar to the Pinillos example I discuss in *Mental Files*, but instead of partial merging (or, as I will now say, *partial fusion*), this is a case of ‘partial branching’ (Ball, p. 17) or *partial fission*. What makes the fusion or fission partial is the fact that the file which has been replaced is not discarded, but retained with a status of indexed file and a metarepresentational function. Diagrammatically, one can represent the difference (and similarity) between the two cases as follows :



In each case, three files are synchronically deployed : an inclusive file and two twin files. The inclusive file may be the input to fission or the output of fusion. In Ball’s example of partial fission, the twin files BARRY SMITH₁ and BARRY SMITH₂ are regular files at the time of thinking (after fission), and the inclusive file BARRY SMITH is retained with the status of indexed file to represent the pre-fission point of view. In Pinillos’ example of partial fusion, the inclusive file VENUS is the regular file resulting from fusion, and the twin files HESPERUS and PHOSPHORUS are retained as indexed files to represent the pre-fusion point of view.

Back to transparency. Ball thinks my transparency principle ‘underestimate[s] our power of rational criticism’ (p. 18), but it doesn’t. Transparency is synchronic : at any given time, certain relations of coreference de jure (weak and strong) hold between the files, and in virtue of these relations the subject is immune to certain errors. Certain patterns of noncoreference are excluded. But there are other errors that are still possible ; in particular, the subject can fail to refer by deploying files whose presuppositions are not satisfied. Diachronic self-criticism in such cases leads the subject to engage in *cognitive restructuring* : the system of files has to be revised through e.g. fusion or fission. That is what happens in

⁷ Or perhaps Ball views (1) and (2) as *utterances* which the subject accepts and on which he bases his judgment. Viewed as public utterances, it may be that (1) and (2) contains tokens that do not corefer. But viewed as judgments made by the confused subject, they can only involve deployment of the same (confused) mental file.

Ball's example of recognized equivocation.

IV

Ninan raises two further issues regarding the mental file framework. The first one concerns an argument I give for the framework. The other one concerns the imagination and other counterfactual attitudes.

In the book I compare the mental-file framework to the version of Lewis's centred content approach favoured by Ninan. On that version, the centred worlds that are used to represent the content of a belief state are centred on many different things : they are 'multi-centred'. The center contains, in addition to the Lewisian subject-at-a-time, a *res*-sequence : a sequence of entities to which the subject bears suitable acquaintance relations. In this way Lewis' egocentric descriptivism is avoided : the subject thinks of the objects in his environment as directly as he thinks about himself. I very much approve of this amendment to Lewis's framework, since my main criticism of the framework is its commitment to descriptivism.

The elements of the *res*-sequence in a given world correspond to the objects which the subject takes to exist. Since the same object can be thought to exist under distinct guises (e.g. as the man on the beach and the man with a brown hat in Quine's Orcutt example), what we need in the *res*-sequence, and what Ninan actually provides, are pairs of objects and acquaintance relations (with a given object possibly occurring in more than one pair). But I object that we also need to represent the case in which the subject thinks he bears acquaintance relations to an object which, in fact, does not exist. I give the example of the paranoid who is obsessed by 'that guy who keeps following me' and thinks of him singularly. To account for such cases, I suggest substituting mental files for the <object, acquaintance relation> pairs in the *res*-sequence. The file gives us the ER relation, and the object can be omitted since, from the file and the world, we can extract the object to which the file refers to in that world. By omitting the object, we also make room for the possibility that there is no object out there corresponding to the file. In this case the subject's belief is only 'pseudo-singular' : it involves a mental file, but one that fails to refer.

While remaining open-minded, Ninan says he is not convinced that the move is necessary, and suggests a disjunctive account of the empty case.

I am not necessarily opposed to Recanati's proposed amendment, but I wonder if it is required. For in describing his case, Recanati did something which we very naturally do when characterizing empty (putative) singular thoughts: he appealed to a description of the man that Lucy herself might give, viz. 'that guy who keeps following me'. But if Lucy's thought is really ultimately a descriptive thought, then no alteration of the sequenced worlds theory is needed, for we can say the following. Since Lucy is acquainted with herself via identity, there will be an element $\langle y_h, R_h \rangle$ of her *res* sequence such that y_h is Lucy and R_h is the identity relation. So if Lucy believes that there is a unique man who keeps following her, we can say that all of the sequenced worlds $\langle w', t', g' \rangle$ compatible with what Lucy believes at t in w relative to her *res* sequence $\langle \langle y_1, R_1 \rangle, \dots, \langle y_n, R_n \rangle \rangle$ are such that there is a unique man following g'_h around at t' in w' .

Of course, one might argue that, at some level of mental representation, empty (putative) singular thoughts should be treated no differently than genuine singular thoughts. But I think this is not so clear. A non-descriptive approach to genuine singular thoughts is possible because we can characterize the agent's state of mind by appealing to the object in the world that the thought is about, and to the way in which

the agent is related to that object. But since we lack these resources in cases of empty (putative) singular thoughts, all we have to go on are the qualitative properties that the agent believes are instantiated by the non-existent thing, along with the qualitative relations that the agent believes the non-existent thing bears to other existent things. From this perspective, a disjunctive treatment of these two kinds of cases seems quite natural. (Ninan, pp. 6-7)

But I think this is a step back in the direction of Lewisian descriptivism. Lewis treats all cases in which the subject thinks about an object he is acquainted with as a case in which the subject self-ascribes a certain relation to some object or other. The singular thoughts turn out to be object-general even if subject-singular. Ninan restores singularity for thoughts about objects but is willing to make an exception for nonexisting objects : in this case he suggests the thought is general even though, from the subject's point of view, it is singular. This disregards the phenomenology, the internal feel of the thought, which goes together with its functional profile. The profile in question is independent of whether or not the file refers ; it corresponds to one of the things people mean (or meant) by the 'narrow content' of the belief (what one uses in action explanation). The mental file story tries to capture that functional profile and is looking for generalizations which will presumably be missed if we opt for a disjunctive approach.

Ninan also wonders how the mental file framework handles attitudes such as the imagination. When we imagine something about an object, which mental file is involved ? We don't want to mix the things we believe about an object with the things we desire or merely imagine about them, so compartmentalization is needed. Forbes suggests that there should be compartmentalization *within* files : instead of storing predicates or 'conditions' in the file, one stores 'classified conditions', that is, pairs of attitude and predicate. As Forbes puts it,

A condition stands for something an object can satisfy, and the classifier is what specifies the subject's attitude toward a certain related proposition. Possible classifiers for conditions are 'believed to be true' or 'hoped to be true'. (Forbes 1990 : 538)

Suppose I believe that Clemens is not famous but I imagine that he is (Ninan's example): my CLEMENS file will contain both the information <BEL, not famous> and <IM, famous>. Another option, suggested by Ninan on my behalf, consists in having the attitudinal compartments first (belief box, desire box, imagination box), all populated with files. In this framework, the files will have to be multiplied since each box will contain its own files. My belief-file about Obama will contain such and such predicates, but my wish-file or my imagination-file about him will contain different things. As Ninan points out, the Obama files in the different compartments may be linked together. This suggests a generalization of my system of 'indexed files'.

In my book indexed files belong to a subject s_1 but they are indexed to another subject s_2 and are used to represent that subject's point of view (what he believes about the object). Though indexed to another subject, the file may still refer via its link to a regular file in s_1 's mind (a file which refers in the normal way). Some indexed files remain unlinked to regular files : they are 'free-wheeling'. A good example is the discourse I analysed above in connection with Ball's 'recognized equivocation' case :

The Barry Smith story

When we discovered that Barry Smith₀ was both a philosopher of mind and language and an ontologist, and that he₀ lived both in London and Buffalo, we started

wondering whether he_0 was a single person or whether there were two philosophers bearing the same name. Now, of course, we know that Barry Smith₁ is not an ontologist and that Barry Smith₂ does not live in London.

Here the subject (the speaker) knows that the inclusive file BARRY SMITH associated with all the singular terms subscripted with '0' does not refer. He deploys it for purely metarepresentational purposes, without thereby referring to anybody.

If we generalize indexed files by letting some be indexed to nondoxastic attitudes of the subject, rather than to other subjects, then we can apply the idea of a free-wheeling indexed files to the case in which I imagine something not about a real object, but about an object I merely imagine. In this case there will be an imagination-file about that object, unlinked to any of my regular files.

We can generalize the indexed file system in this way, but maybe we don't have to. Those who think that to imagine is e.g. to 'imagine seeing' think that the content of an imagining always involves the representation of an attitude (here, seeing). The subject of the represented attitude is the subject whose point of view is represented in the imaginative episode. It may be called 'the subject of imagination', possibly distinct from the imagining subject (Recanati 2007 : 203-10). In this framework, the alleged imagination file (a species of generalized indexed file) can be dispensed with in favour of an indexed file in the original, restricted sense : a file indexed to the subject of imagination. As usual such a file may or may not be linked to a regular file in the speaker's mind.

I don't know which option is better (generalizing indexed files, or positing a subject of imagination distinct from the imagining subject), but they share a common element : the idea of indexed files which can be linked to a regular file (and thereby refer) or remain unlinked (free-wheeling). Free wheeling indexed files are admittedly nonreferential. They don't carry information about any real object. Ninan complains that 'one can no longer characterize mental files as mental representations whose primary function is to carry information about objects. This characterization is true only of the files in the belief box' (the regular files). 'Is this a problem?', Ninan asks. I answer that it isn't. Even when they are free-wheeling, indexed files *simulate* reference (Recanati 2012 : 183, 2013 : 207-13). Because they do so, they are parasitic on the regular files — those 'whose primary function is to carry information about objects' (Ninan, p. 8) — so the primacy of reference is preserved.

References

- Ball, D. (this issue) Indexicality, Transparency, and Mental Files. *Inquiry*.
- Evans, G. (1981) Understanding Demonstratives. Reprinted in Evans 1985, pp. 291-321.
- Evans, G. (1982) *The Varieties of Reference*, ed. J. McDowell. Oxford : Clarendon Press.
- Evans, G. (1985) *Collected Papers*. Oxford : Clarendon Press.
- Forbes, G. (1990) The Indispensibility of Sinn. *Philosophical Review* 99 : 535-63.
- Gibbard, A. (2012) *Meaning and Normativity*. Oxford : Oxford University Press.
- Kaplan, D. (1989a) Demonstratives. In J. Almog, H. Wettstein and J. Perry (eds.), *Themes from Kaplan*, pp. 481-563. New York: Oxford University Press.
- Kaplan, D. (1989b) Afterthoughts. In J. Almog, H. Wettstein and J. Perry (eds.), *Themes from Kaplan*, pp. 565-614. New York: Oxford University Press.
- Lewis, D. (1976) Survival and Identity. In A. Rorty (ed.), *The Identities of Persons*, pp. 17-40. University of California Press.
- Lewis, D. (1979) Attitudes *De Dicto* and *De Se*. *Philosophical Review* 88 : 513-43.
- Ninan, D. (2008) *Imagination, Content, and the Self*. PhD dissertation, MIT.

- Ninan, D. (2010) *De Se* Attitudes : Ascription and Communication. *Philosophy Compass* 4 : 1-16.
- Ninan, D. (2013) Self-Location and Other-Location. *Philosophy and Phenomenological Research* 87 : 301-31.
- Ninan, D. (this issue) On Recanati's *Mental Files*. *Inquiry*.
- Onofri, D. (this issue) Mental Files and Rational Inferences. *Inquiry*.
- Papineau, D. (2013) Comments on François Recanati's *Mental Files*. In Salis (ed.) (2013), pp. 159-76.
- Perry, J. (1980) A Problem about Continued Belief. *Pacific Philosophical Quarterly* 61 : 317-32.
- Perry, J. (1997) Rip van Winkle and Other Characters. *The European Review of Philosophy* 2 : 13-40.
- Perry, J. (2000) *The Problem of the Essential Indexical and Other Essays*, expanded edition. Stanford : CSLI.
- Perry, J. (2012) *Reference and Reflexivity*, 2nd edition. Stanford : CSLI.
- Pinillos, A. (2011) Coreference and Meaning. *Philosophical Studies* 154 : 301-24.
- Prosser, S. (forthcoming). Shared Modes of Presentation. *Ms*, University of St Andrews.
- Recanati, F. (2007) *Perspectival Thought*. Oxford : Clarendon Press.
- Recanati, F. (2012) *Mental Files*. Oxford : Oxford University Press.
- Recanati, F. (2013) Mental Files : Replies to My Critics. In Salis (ed.) (2013), pp. 207-42.
- Salis, F. (ed) (2013). *Book Symposium on François Recanati's Mental Files*. *Disputatio* 5/36.
- Sider, T. (1996) All the World's a Stage. *Australasian Journal of Philosophy* 74 : 433-53.
- Sider, T. (2001) *Four-Dimensionalism*. Oxford : Clarendon Press.
- Stalnaker, R. (2003) *Ways a World Might Be*. Oxford : Clarendon Press.