

MENTAL FILING SYSTEMS: A USER'S GUIDE

HENRY CLARKE

How seriously should we take the idea that the mind employs mental files? Goodman and Gray (2022) argue that mental filing—a thinker rationally treating her cognitive states as being about the same thing—can be explained without files. Instead, they argue that the standard commitments of mental file theory, as represented by Recanati's indexical model, are better seen in terms of a relational representational feature of object representations, which in turn is based on the epistemic links a thinker bears to objects. This paper argues that this revision is misguided. Neither the representational property nor any basic role for epistemic links are needed for an adequate explanatory theory that makes use of the image of a mental filing system. A better alternative to the indexical model does posit files, albeit as causal-functional entities. This makes additional representational features redundant, and shows that epistemic links play a secondary role.

1. Introduction

The image of a mental filing system is frequently used to capture some manifest facts about how the mind works: in holding propositional attitudes, thinkers represent objects as having properties; many properties can be attributed to one and the same object in different representations; there is some way that the mind handles the fact that it is one and the same object to which different properties are attributed; this reflects how the thinker conceives of relations of identity and difference between the objects represented. Moreover, all this is rational from the perspective of representing the world truly. Talk of mental files is employed to give a theory of how these facts are realized. Long a feature of the philosophical literature on mind and language, this theoretical approach has recently risen in prominence and has undergone a good deal of scrutiny.¹

1. For an overview of uses of files in philosophy and psychology, see Murez and Recanati (2016: 265–72).

Contact: Henry Clarke <henry_clarke1@hotmail.com>

It is hard (if not impossible) to dispute the claim that there is such a thing as mental *filing*, mental states exhibiting these features, but one might well question whether the existence of mental *files* is required. One might doubt, that is, whether files play any substantive theoretical role. A different question is what, if there is a role for them to play, is required of them to play it—what sort of entities they are, what properties they have. Goodman and Gray have examined this question and concluded that filing can happen without files. They therefore return a negative answer to the first question: talk of such entities should be treated as merely metaphorical.² The second question then goes by the board. The basis for this comes in the form of two theses concerning how filing works: first, that filing involves a relational semantic property of object representations which they term *coordination*, and second, that coordination is grounded in epistemically rewarding relations (ERs) between object representations and objects. They recommend this as a better form for the explanatory commitments made by the standard application of the idea of a mental filing system, explicitly identified with Recanati's much-discussed indexical model of mental files (G&G: 205).

This paper puts the case for an almost diametrically opposed position: files do substantive work, filing does not involve a semantic notion of coordination, and epistemic relations play only a secondary role in determining the functional and normative profile of a filing system. On the view to be explored here, filing without files is incomplete because it leaves out the causal basis of the thinker having the distinctive inferential dispositions involved in filing. To distinguish this psychological notion from the semantic notion of coordination used by Goodman and Gray, I will talk in terms of the presence of *identity-presupposing dispositions* (IPDs).³ Positing files as mental particulars provides a good explanatory account of IPDs, given a suitable understanding of what a file as mental particular is. But once files are brought in to account for IPDs, the appeal to semantic coordination is redundant—plays no role in the activities of a mental filing system—and the appeal to ERs as grounding filing is inconsistent with how ERs are involved in a filing system that tracks objective identities.

2. Goodman and Gray (2022: 204–5): 'The guiding question for the paper is whether, once we move beyond the metaphors, there is any theoretical role for files. Our suggestion is that there is not. To put our view in a slogan: so-called "mental file" theory is committed to mental filing but not to mental files.' I cite this work as 'G&G' in what follows.

3. The language of 'presupposed identity' is frequently used by Recanati; see, e.g., Recanati (2012: 50, 88, 91–2, 140, etc.; 2016: x, 18, 29, 54, etc.). In Clarke (2018), I termed the psychological profile 'coordination', a choice of terminology that would be confusing in the present discussion. G&G mention this psychological use of 'coordination' and note that the discussion of mental files differs from theirs (see G&G: n31). This paper can be read as an attempt to spell out the difference and to show what hangs on it.

Section 2 lays out standard mental file theory and Goodman and Gray's recommended revision. Section 3 argues that standard mental file theory ought to posit mental files as causal-functional entities. Section 4 argues that causal-functional files render semantic coordination in thought redundant for explaining filing. Section 5 discusses the role of ERs, and argues that they play a secondary role which requires mental files to be in place. Section 6 concludes.

2. Filing without Files

I start with a brief summary of how Goodman and Gray set out what they term the 'standard' version of mental file theory (MFT), and their revisionary alternative.

Standard MFT can be stated as follows:⁴

- i) A mental file contains mental predicates. If two predicates are contained in the same file ('co-filed'), then this is linked to the fact that the thinker in question has the appropriate propositional attitudes and takes them to be satisfied by the same object.
- ii) A mental file refers to an object, at least if it is functioning correctly. The referential property of mental files equips them to be concepts, mental representations that are employed as the syntactic basis of thoughts. A thought is therefore about an object in virtue of the mental file it employs.
- iii) A file refers to an object by being related to the object through epistemically rewarding relations (ER). ERs are epistemic links between thinkers, or mental states, and objects. They are ways of gaining information about things, by being in epistemic contact with them. Perception is a paradigm of an ER: Lucy sees Honey the dog running from the bush into the muddy puddle, and so updates her file with Honey's current location, and the expectation that her paws will leave prints on the floor.
- iv) Files are based on the ERs through which they refer, and are individuated by them. A file is based on an ER either because it was opened as a new file to exploit an ER, or else because it is a file that was converted from one that depended on an ER of a given kind to an ER of a more expansive kind (e.g., one that includes both current perceptual information and remembered information). A file can therefore be updated over time, and so only contingently contains any collection of mental predicates.

4. See Recanati (2012: 35–88; 2016: 10–13).

- v) The inclusion of predicates in distinct files explains why a thinker may be ignorant of the fact that two objects she thinks of as non-identical are in fact identical (so-called 'Frege cases').

A thinker takes predicates to be satisfied by the same object when she is disposed to engage the relevant propositional attitudes in inferences which involve attributing two or more properties to the same object, where this does not involve a premise that explicitly asserts an identity involving that object. This is known as 'trading on identity'.⁵ The paradigm is an inference that conjoins two predicates and generalizes over the conjunction, without an identity premise: from 'a₁ is F' and 'a₂ is G' the thinker can infer 'There is something that is both F and G' without needing to entertain the identity 'a₁ = a₂'. As is by now broadly acknowledged, this sort of inference cannot be reduced to inferences with identity premises. Rather, the inferences *must* trade on identity at some stage. I am going to talk about these inferences in terms of *presupposed identity* to mark the fact that there is apparently an attitude held by the thinker to an identity, one that renders the inferences permissible if well-founded, that cannot be reduced to a representation of identity that works as an identity premise.

G&G note two problems with standard MFT:

Containment: If a file is a concept (point ii), then a file is a constituent of a mental state. But the mental states seem to ground the containment of predicates in files (point i). So there is a danger of 'reciprocal containment' (G&G: 207–8).⁶

Diachronic identity: If files are individuated by ER relations (point iv) then converting a file to a new ER relation would involve distinct files over time, which would erroneously predict Frege cases (point v) where there is continuity in the thinker's conception of the object (G&G: 209–10, n15).

While neither problem is necessarily fatal for standard MFT, any solution would require adding or subtracting commitments in order to preserve the basic theory.

In light of these problems, G&G examine what the advocate of standard MFT is trying to explain, and what resources are really needed to do the work. First, they set out what they take to be the explanatory problem to which MFT

5. The term 'trading on identity' comes from Campbell (1987–88: 275–276). See also Recanati (2012: 47–49), Sainsbury (2002: 134–36), Millikan (2000: 140–44).

6. G&G also lay out versions of this worry concerning relational beliefs (e.g., the belief that A is taller than B), and the worry that attitudes other than belief can be involved in trading on identity.

is addressed. G&G remark, about the idea of co-filing meaning that the thinker can trade on identity, that

it is not yet clear what the modal force of ‘can’ is. Does it express a kind of psychological ability? A kind of rational permission? Something else? What, precisely, is the explanatory import of sameness of file in relation to trading on identity? . . . [W]e believe the imprecision captures a genuine ambivalence in standard expositions of the file framework. (G&G: 207)

They suggest that the core explanatory commitment of standard MFT is best seen as explaining the rational permissibility of a certain class of cognitive transitions, moves from one state to another, namely, those that trade on identity (G&G: 211).

With this in place, they argue that

(1) The rational permissibility of the transitions is ‘constituted’ by a relational representational feature—semantic coordination—that obtains between object representations in thought (*the semantic thesis*)⁷
and

(2) Semantic coordination relations obtain, in part, because of ERs that obtain between object representations and the objects they represent (*the metasemantic thesis*).

They further argue that

(3) Mental files play no role in semantic coordination, including how ERs ground semantic coordination.

And so they conclude that

(4) Files as mental particulars are explanatorily redundant for filing.

The two positive theses therefore require a revision to standard MFT. Files individuated by ERs are out. In their place are object representations, semantic coordination relations between them, and epistemic links between object

7. I put ‘constituted’ in scare quotes, as it is not clear what it means for the rational permissibility to be constituted by a semantic relation, their being quite different things. At one point, G&G use the term *partly constituted*, but then put the explanation of permission as a biconditional linking permission to coordination (see G&G: 212, the statement of *Coordination as Relational Representation*). The most charitable interpretation of what it means is that semantic coordination is explanatorily sufficient for rational permissibility. I dispute this in Section 4.

representations and objects. To prevent my presentation becoming unwieldy, I leave out further detail on how the positive theses are supposed to work and how G&G motivate them until the critical discussion below.

Before getting to that, it is important to recognize exactly the dialectical point G&G are pursuing: G&G's aim is show that the advocate of standard MFT would do better to drop the commitment to mental files. Standard MFT requires certain explanatory claims to be true, but if they are true, then files turn out to be a dispensable part of the theory. G&G are not advocating the mental file framework. Nor are they even committed to the existence of filing. What they advocate is the claim that standard MFT involves a set of commitments, (i–v) above, that can be reformulated without reference to files in a way that still provides the explanatory resources that files are supposed to provide. Call this the *revisionary alternative*.

It will help, in framing the discussion to follow, to put the issue in a slightly different albeit complementary way. Of the two options, G&G argue, we ought to choose the file-free revisionary alternative over file-positing standard MFT. Whether or not this is true when the question is limited to the two options considered so far, my aim here is to put a third option on the table. I will argue that an advocate of standard MFT would be better off *not* dropping the existence of files but rather shifting their understanding of what files are, and dropping the parts of standard MFT that G&G argue are well-enough captured by semantic coordination and its grounding in ERs. *Those* parts of standard MFT are unnecessary, not the files.

The 'standard' version of MFT is very much Recanati's indexical model. Whether or not it deserves the epithet is a moot point. More interesting is the question of the extent to which the indexical model goes beyond the basic idea of a mental file in explaining mental filing. This is sometimes missed, not least in Recanati's own discussion, which could be read as holding that the indexical model, with its central role for ERs and its attribution of representational properties to files, somehow merely makes explicit what is anyway implicit in the idea of a mental filing system. My aim here is to point out the advantages of a construal of the file idea which, while not incompatible with standard MFT/the indexical model as such, takes a more modest line.

The next section makes the case for this construal. The agenda of the sections thereafter will be to question the semantic thesis first of all, and then to examine the thinking behind the metasemantic thesis if the semantic thesis is rejected.

3. Files as Causal-Functional Entities

This section argues that a category of mental particulars that deserve to be called mental files will be part of a good explanation of filing. I acknowledge the two

phrases involved in this claim, ‘deserving to be called mental files’, and ‘a good explanation of filing’, need to be clarified; in the absence of an ability to say everything all at once, I will need to discharge the obligation to clarify both by the end of the section.

Recall that G&G hold that the explanatory aim of standard MFT is to explain the permissibility of inferences that trade on identity. G&G don’t go into what the question about permissibility involves, but it bears some analysis. What is the potential problem that would entail an inference being *impermissible*?

A common way to think about these inferences is to contrast them with what happens when a thinker makes an inference that has a belief about an identity as a premise. For example, Lucy wants to know whether all dogs like going out in the rain. She knows Honey the dog, and I tell her that Honey hates going out in the rain. So she concludes, on the basis of what she now knows, that there is at least one dog that does hate going out in the rain. We could represent Lucy’s cognitive transition as follows:

- (1) Honey is a dog.
 Honey hates going out in the rain.
 $\therefore \exists x (\text{Dog}(x) \ \& \ \text{Hates rain } (x))$.

Suppose instead that she knows of a dog called ‘H’. When I tell her that Honey hates going out in the rain, she would need to know also the identity ‘H is Honey’ to draw the same conclusion:

- (2) H is a dog.
 Honey hates going out in the rain.
 Honey is H.
 $\therefore \exists x (\text{Dog}(x) \ \& \ \text{Hates rain } (x))$.

Without the identity premise in (2), making the inference would put Lucy in danger of forming an inaccurate representation of things. The conclusion would be accidentally true. Moreover, the premises would manifestly only support the conclusion accidentally; it would be a manifestly invalid inference. So she would be entering into a state, taking a set of beliefs to support a conclusion that she knows is not supported, that is manifestly incoherent in an avoidable way.

As G&G point out, practical inferences also exhibit this pattern. Peter is in Birmingham, and looking at the departures board at New Street station, he forms a plan:

- (3) I want to go to London
 If I go to platform 1, I can go to London.
 \therefore I intend to go to platform 1.

Pierre is in Paris and wondering how to travel to Londres to find out what all the fuss is about. Arriving at Gare du Nord, he sees a train going to London but fails to realize that London is his desired destination. A helpful stranger informs him that Londres *is* London. Pierre can then form a plan:

- (4) I want to go to Londres
 If I go to platform 3, I can go to London.
 Londres is London.
 ∴ I intend to go to platform 3.

In the case of a practical inference, there is a danger of forming intentions that fail to realize the thinker's aims (or a failure to form intentions with the same result), and the point about avoidable incoherence applies here too.

Assessing a transition with these goals in mind—avoiding inaccuracy, achieving coherence—comes with two typical dimensions of assessment, in this case directed at the presupposition: Is it sufficiently linked to the identity of what the thoughts are about? Has the thinker taken appropriate care in making the inference? Both are inflected by context-dependent standards ('sufficiently', 'appropriate'), but in any context there will be a mechanism that will be assessed, the one that makes it the case that the thinker has the dispositions.

Note that neither dimension of evaluation is independent of whether the thinker in fact *has* the inferential dispositions. If there are no such dispositions, then the permissibility question is rendered pointless. To ask of someone who lacks it whether the disposition would be permissible despite their lacking it could only mean: would it be permissible, were they to come to have the disposition? The question wouldn't apply if the thinker *discovered* an identity, and so came to have an identity belief that they could call upon in inference, as their relying on an identity belief would not be a case of presupposed identity. Given a thinker's cognitive state, the permissibility of the inferences that presuppose identity is only in question, the conditions only need to be met, if in that state the thinker has the disposition to make inferences without an explicitly represented identity. Moreover, while the question of permissibility applies to making the inferences, the dispositions to make the inferences may be assessed along the same dimensions. The dispositions can be evaluated for permissibility even in the absence of the inferences actually being made. After all, it is the thinker's mental state that is up for evaluation, and this is not limited to its activity but also what activity it would undergo in certain circumstances, those which G&G call 'occasions for trading on identity' (G&G: 217).

The question of the causal basis of the distinctive inferential dispositions is therefore a question that is in an important sense prior to the permission question. The mental file theorist, who is concerned to account for the permissibility

of IPDs, is thereby also concerned with their causal basis, and has a distinctive account to offer: IPDs are based on mental files.⁸

Why might the mental file theorist feel moved to say this? Here is the answer I recommend: whatever the ultimate causal basis of the dispositions, at the level of functional description of the cognitive system, they will be caused by something with the functional profile of an entity with *file structure*.⁹ By file structure, I mean a structure consisting of an object, an open-ended and potentially changing collection of ways for things to be (of properties, or categories), and an association between the object and the collection such that a user of that entity would answer quantificational questions about the categories, such as ‘how many things fall into this collection of categories?’, with ‘at least one’ because of that association.¹⁰ The object is the file; the categories are picked out by the predicates a thinker represents some individual as satisfying.

This is an abstract way of describing how real-life, physical files work. Physical files have file structure: an object (a manila folder, for example), a collection of category representations (written on bits of paper, shown in diagrams or photographs, or recorded in other media), and an association between them (the latter being inside the former), such that the person looking in the file will take the information to bear on one object. In the analogous case of a mental file, there is an entity in the cognitive system (Lucy’s Honey-the-dog file, for example), some collection of category representations (‘is a dog’, ‘hates the rain’), and association is whatever it is in the thinker’s cognitive architecture that makes those representations enter into the appropriate cognitive transitions that presuppose identity. The appropriate transitions are those that depend on the functional role of the mental state linked to the associations of category representations with the file. A mental file is thus an entity at a certain level of functional description of the cognitive system; exactly how it is realized, and how it relates to other ways of describing the same system, is an open question.

8. G&G appear to be uninterested in what brings the IPDs about; at least I cannot locate a point where they acknowledge this as an explanatory question, much less provide an answer. Even if identifying permissibility as the explanatory target of MFT is correct as an interpretation of *Recanati’s* particular treatment of it, the argument here shows that the choice between permissibility or the psychology of the inferential dispositions as ‘the’ explanatory target is a false one. A mental file-theoretic explanation of permissibility is better if it includes an explanation of the psychology of filing too.

9. Millikan’s (2000) ‘Strawson markers’ are mental files, in this sense. Her more recent (2017) notion of ‘unicepts’ is a class of mental entities that includes entities with file structure as well, although she officially disclaims the terminology (2017: 49n4).

10. This is a necessarily circuitous way of putting it. The associations between files and predicates need not always have the force of attributing the property to the object; rather, they might have the force of desires or fears, graded credences or conditional beliefs, and so on.

(First obligation discharged: this is what I mean by 'deserving to be called a mental file': being a particular with file structure, so belonging to a functionally individuated kind that plays a specific causal role in the cognitive system.)

Why accept this suggestion? Well, what would have to be the case for it to be false? Bear in mind that the claim isn't that there needs to be any particular physical (or chemical, or biological, or neural) arrangement that realizes the file structure, any more than there need be any such arrangement that realizes the same structure in a filing cabinet, database, or file directory on a computer. For the claim to be false, therefore, it would have to be the case that there is no causally efficacious structure in the cognitive system that fulfils the function of putting IPDs in place. How could that be? There are only two options: either the function is not fulfilled, which entails there being no IPDs; or it is, and so there are IPDs, but they have no causal basis. I submit that neither of these options is better than accepting the claim that there are mental files, understood as entities with file structure that play the causal-functional role. One option denies a manifest fact about the mind, the other accepts the fact but denies it has any causal basis.

(Second obligation discharged: a good explanation is one that doesn't deny the data and doesn't leave them dangling either.)

Two comments before moving on:

First, to return to the comment on Recanati's exposition of the indexical model in the previous section, I want again to stress that causal-functional files are not *as such* mental representations. Being a mental file, therefore, is not *per se* to be an object representation on this construal, nor is it to play any role subsidiary to that. An object representation picks out a particular object, or purports to. The identity of the object to which the category representations apply is a purely formal question, as far as their association with a file is concerned: they apply to one particular object, but *which* object doesn't matter. A file contributes to a thinker's mental state purely by causally underwriting the sort of transitions the paradigm of which is conjoin-and-generalize, with semantically evaluable end states (a true belief, a satisfiable intention, etc.). This contrasts with standard MFT, which holds that they have both roles: files explain clustering *and* act as conceptual representations, to use Recanati's terms. Since mental files are not concepts, on the causal-functional way of understanding files, they are not subject to the containment puzzle: mental files are in no sense constituted by propositional attitudes, nor are the thinker's dispositions constitutive of predicates being associated with files.

Second, the issue driving the argument for causal-functional MFT isn't whether standard MFT *can* be addressed to the question of what explains the inferential dispositions. I have argued that *if* standard MFT is addressed to the question of permissibility of undergoing the distinctive cognitive transitions, then it must (i) include the presence of the dispositions in any account of the

explanation of permissibility, and therefore (ii) would be a better account were it to include a story about what makes those dispositions present that makes appeal to files.

4. Permissibility without Semantic Coordination

The previous section argued that adherents of standard MFT would do well to accept the existence of mental entities with file structure—that is, mental files—as the causal basis of the presence of IPDs. The presence of IPDs, and so their having a causal basis, is a necessary part of the story about their permissibility. Therefore, standard MFT would be better off positing files for that explanatory purpose. This section argues, on this basis, that semantic coordination is redundant, as against G&G who hold that semantic coordination is needed to explain the permissibility of transitions that trade on identity. The argument, in brief, is that the presence of IPDs, plus a further sort of disposition to be described, does the work that semantic coordination is supposed to do in G&G's revisionary alternative. The claim is not that there is no such thing as semantic coordination, on which I do not intend to take a stand. Rather, the claim is that semantic coordination is explanatorily irrelevant to filing. If a thinker has IPDs, the permissibility of those IPDs simply requires the thinker to take reference to be the same, and to do so in the absence of any easily accessible indication that this is a mistake. How so?

IPDs, based on files, are means a thinker has for extending what she knows, and forming intentions on the basis of what she knows, in order to achieve things she cares about and to realize values she holds. That is, she can undergo both theoretical and practical inferences which at some level must presuppose identity, the outcomes of which are used to pursue plans, to avoid bad situations, and to maintain good ones. The outcomes of her IPDs are, and are relevant to, matters of rational concern. She will therefore be motivated to *monitor* her IPDs, and if rational will do so.

Monitoring consists, in Lawlor's term, of file-maintenance dispositions (Lawlor 2001: 71–100).¹¹ What exactly these involve would require a lengthy discussion of how identity and identity mistakes might frustrate or promote one's ends, and how different ways of their being frustrated call for different ways of fixing the state of one's files. But speaking generally, the thinker's aim will be to achieve a state of coherence. That is, a thinker wants to avoid a state where

11. Lawlor's detailed analysis of the psychology of what she calls 'thinking with co-referential purport' is an excellent resource for understanding this crucial aspect of filing. Lawlor's account of the permissibility of IPDs differs from the one offered here.

something else she has reason to believe conflicts with the outcomes of the IPDs she has, because it indicates that the information associated with a file is not satisfied by a single individual. This would sometimes result in the need to 'prune' information (abandoning beliefs) or in the catastrophic case, to abandon a mental file and reconstitute the beliefs with a new set of files.

Since they presuppose identity, IPDs and the beliefs involved in those dispositions purport to track facts about an object. Monitoring is therefore responsive to whether this is done successfully or not. When there is no manifest incoherence, it confirms that they are successful. If a thinker's monitoring is sufficiently careful, then her inferences are safe, since they are produced by a cognitive mechanism that, due to her monitoring, would not easily produce them were they faulty. They are also responsible, since being disposed to monitor IPDs results in epistemically proper responses to manifest epistemic or practical conflicts which are in need of resolution. IPDs serve to put the identity presupposed into question, by purporting to track an object, and the monitoring dispositions serve to provide an answer, by confirming that tracking is successful, and therefore making it safe and responsible.

In short, once IPDs are put in place by the presence of a file, then a thinker with a concern to possess accurate information and coherent intentions, and the disposition to act on that concern, has everything she needs for them to be permissible. Call the combination of the causal-functional notion of files plus the permissibility of IPDs arising from monitoring dispositions *causal-functional MFT*.¹²

A useful way to frame the disagreement between this view and G&G's revisionary alternative is to ask how co-filing is understood on these views. On causal-functional MFT, it is a functional relationship between exercises of a thinker's capacities for categorizing particular things—tokenings of mental predicates. This relationship is based on the link between those mental predicates and a mental entity which brings about IPDs. Theories of categorization supply accounts of the cognitive architecture involved in these exercises; as to the nature of the entity that brings about this relationship between them, the causal-functional mental file theorist is committed only to there being something with that functional profile, not what does it or how. Moreover, co-filing does not equate to permissibility; it must occur within the context of a cognitive system that is disposed to monitor the existence of IPDs and responds to that disposition.¹³

12. It bears emphasizing that it is not the mere fact of co-filing that makes IPDs permissible. They need to be monitored also. This puts causal-functional MFT at odds with the claim made by Recanati, in the course of expounding standard MFT, that co-filing 'licenses' trading on identity all by itself; see Recanati (2012: 42; 2016: 11–12).

13. Lee (2022) takes co-filing to be a primitive feature of the MFT explanation of trading on identity, which is consistent with the causal-functional understanding of files, but also holds it to

In contrast, G&G hold that co-filing is best understood as (or is better replaced with) a representational feature of object representations. This feature has a metasemantics that involves functional facts, but is not reducible to those facts, since it requires interpretative facts as well. It therefore attributes to object representations a representational property in addition to their doing the job of representing objects. G&G canvass some possibilities about what this is, following Fine, but decline to commit, other than holding that it cannot be reduced to non-relational properties of object representations (G&G: 214; see Fine 2007).

Why prefer causal-functional MFT? We know that the causal-functional notion of co-filing needs to apply in a case of a thinker with IPDs, whether or not the representational notion applies. So the question is: once the causal-functional relationship is in place, and the further monitoring that a rational thinker would anyway need to carry out is also in place, what explanatory work is semantic coordination needed to do? None. All the semantic properties needed are there anyway, since we have object representations and the category representations that determine truth/satisfiability conditions for the thoughts involved, inferential dispositions to tie them together with a presupposition of identity, and the monitoring dispositions to complete their normative standing.

The point can be put like this: We know from the contrast between inferences that need identity premises and those that do not that some inferences would fail to be safe and responsible even though the thoughts involved are co-referential. That is, co-reference is insufficient for permissibility. Instead, whatever explains permissibility needs to be something that links the cognitive dispositions a thinker has to whether her thoughts are co-referential, such that she would not easily have the dispositions were they not co-referential. IPDs plus monitoring dispositions achieve this. Consequently, any semantic relationship stronger than co-reference is not needed for permissibility. Semantic coordination is such a semantic relationship. Therefore, semantic coordination is not needed for permissibility.¹⁴

Could it be explanatorily sufficient? That is, might the explanation of permissibility be over-determined? That depends on the relationship between the functional set-up necessary for permissibility and the metasemantic basis of semantic

be essential that files are concepts, which is not. Whether this combination is coherent is an interesting question, though not one I can address here.

14. This point applies to anti-realist neo-Fregean appeals to sense as well, cf. G&G: 213. None of the above speaks against a role for semantic coordination in thought in some other explanatory capacity, nor does it speak against semantic coordination in language, at least not without further argument. It may be that the interpretation of noun phrases and variables involves semantic properties like coordination, semantically required co-reference, co-reference *de jure*, etc. The contrasting point about IPDs would then be that there is no analogous point of interpretation, as the presence of the IPDs pre-empts the interpretative question to which semantic coordination is supposed to be the answer.

coordination. Either (i) semantic coordination obtains because and whenever the functional set-up obtains, or (ii) it does not. If (ii) is the case, then it follows that semantic coordination is explanatorily insufficient for permissibility because logically insufficient. If (i) is the case, then semantic coordination is dependent on the functional set-up providing the right metasemantic basis. In G&G's discussion, the metasemantic basis involves the thinker being rationally interpretable in a certain way, as having rational inferential dispositions. But adding semantic coordination doesn't *improve* the interpretation—it doesn't make the thinker more rational. More generally, adding the metasemantic fact that the set-up is sufficient for an extra representational feature makes no difference to whether the IPDs are permissible. What constitutes their being permissible is the fact that they are safe and responsible, and what makes them safe and responsible is their involvement in the right functional set-up. Semantic coordination therefore runs in parallel to the permissibility of IPDs, at best.

G&G could object here that I have not countenanced their argument that permissibility has to be explained by a representational feature. That argument runs as follows:

Suppose someone was to make a *single*-premise generalization of a kind similar to the *multi*-premise conjoin-and-generalize inference that is the paradigm of identity presupposition. For example, Smith infers from her belief that Twain is an author that someone is an author. This inference is permissible, and G&G 'take it to be obvious that the relevant permission is constituted by the representational features of Smith's belief: it is because the belief has the content that it does that Smith has the relevant permission' (G&G: 211). Now suppose that Smith also believes that Twain is from Connecticut. Jones believes that Twain is an author, but believes that Clemens is from Connecticut. Even if Twain is Clemens, G&G claim that

file-theorists will admit that Smith and Jones are in different representational states. The difference between them is not merely at some 'lower level' of explanation (functional, computational, or whatever). The difference between them is that, in Smith's case, the coreference of the two attitudes about Clemens/Twain is representationally encoded. (G&G: 211)

Why think that?

To deny this would be to hold that the *kind* of explanation that we give when we say that Smith *can* conclude that someone is an author because she believes that Twain is an author is of a radically different sort than when we say that Smith *can* conclude that some Connecticutian is an author because she believes that Twain is an author and that Twain is from Connecticut. (G&G: 211)

This is, they say, ‘bizarre on its face,’ since single-premise inferences are licensed by their representational content, and there is no reason to think that multi-premise inferences are any different.

But is it true that there is no reason to think that the permissibility of multi-premise inferences might differ from the single-premise case? Clearly there *is* a difference between the two sorts of inference, a difference which is directly relevant to the question of permissibility that we are discussing. The question about permissibility in the multi-premise case is: is the thinker correct (responsible, safe) in taking her object representations to be co-referential? That question doesn’t apply in the single-premise case, as there is no question of co-reference. So it is not obviously right to insist that the facts that bear on permissibility must be the same, since the question of permissibility cannot be the same question in the two cases. It would need to be shown that there is something that suffices for the permissibility of the single-premise inference which *must* also apply to the permissibility of the multi-premise inference, specifically with respect to the question about taking object representations to be co-referential. The analogy itself cannot show this.

G&G might want to appeal to a more general principle, which the single/multi-premise inference comparison is perhaps meant only to illustrate. Inferences are cognitive processes that respond to the content of attitudes. Content depends on representational features. So an explanation of a property of inference, as a content-sensitive transition, must also depend on representational features. It might be concluded from this that an explanation of the permissibility of inference that does not appeal to a representational feature additional to reference is ruled out. This could be what G&G mean when they say that a non-representational explanation would be of ‘a radically different sort’ to the one that applies in the single-premise case.

The general principle might well be correct, but it does not support the conclusion. It would be odd, bizarre even, if permissibility of an IPD had *nothing* to do with representational features of the attitudes involved. But it does not follow that permissibility (or rather, the conditions thoughts need to meet for there to be an explanation of permissibility) of IPDs must be *constituted by* representational features. In other words, it does not follow, from the agreed facts that inferences are responses to content and that content depends on representational features, that the aspect of those transitions that explain their permissibility—whatever that involves—must be representational facts. Rather, the explanation must involve features of the attitudes that are *suitably related* to representational facts. This can include not just the contents of object representations themselves but also how a thinker responds to them and how she manages her response.¹⁵

15. Heck (2012: 156–58, 163–72) makes essentially the same point in reply to an objection to their formal relationism, a view addressed to psychological explanation, not to normative

What exactly this means depends on what the question of permissibility is, what aspect of the inference may be in question and for what reason. In this case, it concerns undergoing a transition that depends on object representations picking out the same object: if they do not do this, the state to which the thinker transitions is unsupported by what she knows or believes. So the explanation must relate permissibility to object representations; it must be the case that what makes for permissibility depends on the nature of the object representations *qua* object representations. The causal-functional explanation fulfils this remit. The monitoring dispositions only apply because of the representational nature of the object representations and category representations involved. Without that, they would have no basis or direction. Similarly, it is because of the representational content of attitudes that IPDs in combination with the right sort of monitoring dispositions produce the positive normative status.¹⁶

The claim here is that the explanation of the permissibility of IPDs does not include semantic coordination, or any additional representational features than those already present in having the IPDs, irreducibly relational or otherwise. The analogy with single-premise inferences does not succeed in establishing the need for an additional representational feature to figure in the explanation, and the general principle linking the grounds of permissibility of inference to representational properties does not undermine the causal-functional explanation. To that extent, the redundancy of semantic coordination for filing stands.

5. ERs and Tracking

The previous section showed that semantic coordination is not relevant to explaining the permissibility of IPDs. This means that the truth or otherwise of the metasemantic thesis G&G put forward is irrelevant to the present discussion. However, it does raise a question concerning the place of ERs in causal-functional MFT, or otherwise put, concerning the relationship between how a thinker forms her beliefs and the rationality of her understanding of their inferential relationships given the causal-functional account.

Recall that ERs are epistemic links to objects. On standard MFT, the mind bases its mental files on the ERs it bears to objects, and a file is individuated by which ER it is based on. G&G reject the individuation claim, since they reject

evaluation. Formal relationism is compatible with the position I argue for regarding the evaluation of IPDs.

16. One might want further to suggest that it is these dispositions which serve to direct attitudes towards a particular object, and that this somehow grounds the representational facts—see Lawlor's (2001) notion of maintaining a cognitive link, and Dickie's (2015) idea of cognitive focus. No such claim is needed for the causal-functional explanation of permissibility, however.

the existence of files, but they accept a metasemantic role for ERs. The metasemantic basis that G&G adduce for semantic coordination has a forward- and a backward-looking aspect. The forward-looking aspect is trading on identity, the dispositional marker of filing. The backward-looking aspect is reliance on ERs: object representations being produced by responses to information received about an object, where it is the same object. To motivate this idea, G&G ask us to consider a creature who is sometimes disposed to engage in inferences that trade on identity, but where this disposition is unrelated to any tracking ability, because it puts incoming information into files at random. In such a case, G&G claim, attributing rationality to this creature would ‘sever the constitutive connections between representation, rationality, and non-accidental cognitive success’ (G&G: 220).

G&G take from this the claim that a process that reliably delivers information about a single object is required for rational filing. This point might be thought to survive the redundancy of semantic coordination. That is, it might seem that ERs must have a role in rational filing even if the metasemantic role G&G give to them is irrelevant to permissibility. Like the revisionary alternative, causal-functional MFT also rejects the individuating role for ERs, but is neutral on the metasemantic role G&G give them. How, then, does the way a thinker manages its incoming information bear on the permissibility of IPDs, according to causal-functional MFT? The answer is that it contributes to how the monitoring dispositions work and how they need to work to achieve the good normative standing of IPDs. But an accurate view needs to grant a primary role to the presence of IPD-grounding files, rather than to ERs.

Let’s call the process of using ERs to gain information about an object over time ‘tracking’.¹⁷ Tracking involves identity, since it is a relationship between a thinker and an object which requires the thinker to recognize one and the same object as the source of information. How exactly this works depends on how ERs are counted. If it is a single ER that is used to gain information, then tracking means recognizing it as the same ER; if it is multiple ERs, then these need to be recognized as bearing on the same individual despite their being different. However ERs are counted, they depend on the identity of the object in question, and so the recognition of sameness or difference of ER means interpreting them as relating the thinker to the same object.

The identity involved in tracking is therefore *substantive*: it is a particular object which is the source and subject of the information gained. ERs are, in this respect, similar to monitoring dispositions, but different from IPDs. For IPDs, the exact identification of the individual is not important; it only matters that

17. This is related to but not exactly the sense of ‘tracking’ as introduced by Evans (1985); it does accord with Recanati’s ‘tracking relations’ (2012: 72–75).

it is the same, whichever one it is. Identity is *formal*, we might say, as far as the IPDs are concerned. In contrast, monitoring involves a substantive identity. How a thinker monitors her IPDs is responsive to how the thinker conceives of the individual which the thoughts concern. Whether there is incoherence and how to resolve it, which is what monitoring needs to monitor, is determined by the thinker's conception of the object. A conception of an object depends on the place of the corresponding mental file in the broader conception of the world that a thinker's mental filing systems articulates: which objects there are, what kinds of objects there are, how they relate to each other, how having one property might bear on having another property, and how all this relates to the thinker in her self-conception. Since tracking depends on recognition, the same must also be true of tracking. The way in which a thinker responds to an ER as the same, or to multiple ERs as bearing on the same, depends on the content of the file in which it is going to be placed, and on other files on which it might have an influence.¹⁸

This point is well made in Millikan's examination of what it takes for a thinker to keep track of an object.¹⁹ She invites us to imagine keeping track of a person (called 'Kate') at a party:

For a brief moment—not much longer, suppose, than a saccade—you divert your eyes to the face of a friend, but immediately pick up Kate's face again. . . . Looking at Kate and hearing her voice, you perceive these as having the same source, as locating the same person. Now Kate passes for a moment into another room, but you continue to hear her voice—though of course there are spaces between the words—and she soon emerges again. . . . Now suppose that Kate looks and sounds familiar also an hour later and then a day later when you meet her again. . . . Probably you would not have recognized her, however, had you met her in Singapore—in some radically disjoint context. . . . Further now, suppose that Kate's name has become familiar, and as more time goes by you often pick up information about her from friends. Again, you usually know which 'Kate' they are talking about from the context, from anticipating

18. This point can be obscured by the myth of a file that is empty of all but one or two pieces of information. While such a file might be possible for, say, beliefs about abstract objects, any conception of a concrete individual will automatically be rich, if not in outright belief or knowledge then in hypotheses, guesses, conditional credences, and so on.

19. Millikan (1997; 2000) is frequently cited in discussions of mental files, though her main insights about what she calls 'sameness tracking' are rarely acknowledged, much less taken up. Her concern is not with the normative dimensions of tracking, which may well be a deficiency of her view, but her point about the holistic and embedded opportunism involved in tracking can easily be combined with the explanation of permissibility offered above.

her possible projectories, and the possible projectories of various kinds of information emanating from her. (2000: 154–55)

As the case illustrates, identity is tracked through the exploitation of multiple sources of information, and any adequate way of tracking an object will necessarily involve at least the potential use of a range of methods. The fact that the object is part of a complex causal network distributed in time and space, and that the thinker is also so situated, makes this inevitable. The crucial point to note is that which methods are used, and which will work, depends on the kind of object that is being tracked, the nature of the environment it is moving through, and how the tracking thinker understands both of these things in relation to her own cognitive and physical capacities.

Millikan concludes that one cannot individuate ways of tracking objects, as a basis for how identity is represented in thought, by somehow using sameness and difference in ways of getting information from objects to explain what it is rational to think (2000: 155–58). In that spirit, causal functional MFT allows us to see that ERs are not of primary importance for the success of tracking, and we do not need a way of counting them as the same or different to understand the permissibility of IPDs. The important question for filing is: given an information link to some individual, which individual is it? It is the individuation of objects thought about that matters, not (*contra* Recanati) the individuation of information links we bear to them. Epicyclical accounts of the right way to individuate ERs are misguided because we track objects, not ERs—we *use* ERs to do this, but we do so by *interpreting them* in terms of our ‘cognitive map’ of the world, which is sustained by and based upon our mental filing systems.

The thinker tracking an object must have a conception of it, and of how it relates to other objects (including, of course, the thinker herself), and so how the ERs she does and might come to stand in, however individuated, deliver information. The epistemology of empirical identities therefore depends on having a view of the world that is already informed by the structuring of information that the file system, with its attendant monitoring, creates. ERs cannot by themselves *guarantee* identity, but rather *indicate* identity by occurring in a context and against the background of a constantly updated conception of the world through which their outputs can be interpreted as picking out a particular object.

To go back to G&G’s creature, it is possible to confuse the creature’s carelessness in sorting information—associating information with files regardless of its source—with the absence of monitoring dispositions. If the creature lacks monitoring dispositions, then clearly its filing system will not have the right epistemic credentials. But failing to sort incoming information is not the same as failing to monitor the effects of the resulting information bundles. If the creature does

that, then it may still have permissible IPDs, though the range of its permissible IPDs is likely to be extremely limited. Fortunately, creatures with our cognitive dispositions and capacities, living in a world populated by relatively stable objects with relatively stable collections of properties, do not face this problem incorrigibly.

ERs are relevant to the permissibility of filing, then, but the role they play is secondary to the other factors that bear on permissibility. Monitoring dispositions may of course include checks on the sources of information and the validity of the identifications they (seemed to) support. However, there is no simple connection between the epistemic links through which information is gained and the rationality of the thinker's understanding of that information. There must be something that accounts for how ERs are understood as bearing on the same thing that is in place *prior* to those ERs being exploited. So we need files to supply the basis of IPDs. This underscores the point that we cannot do without files in our account of filing.

When putting the case for the metasemantic thesis, G&G ask: 'How could we see the disposition to trade on identity as rationally relevant if it is not at least the typical downstream effect of a process that reliably delivered information about a single object?' (G&G: 219). But the relevant process does not just concern what happens 'upstream', but further downstream too. Or, better, think of it as a highly open-ended, inferentially articulated set of cognitive tasks, dependent at a basic level on the existence of a mental filing system and the congruence of that system with an environment. Files without ERs may be empty, but ERs without files are blind.

6. Conclusion

G&G argue that standard mental file theory unnecessarily burdens itself with positing mental files. While I concur that files as mental representations are unnecessary, and that ERs do not in any sense individuate mental files, I have argued that there is nevertheless a need for mental files and so G&G's recommended revision fails. The scope of mental file theory needs to include a causal basis for the presence of dispositions to engage in a particular sort of cognitive transition. Once files are acknowledged as the basis for these dispositions, certain consequences follow: there is an essential role for files as mental particulars; semantic coordination is redundant; ERs play a secondary role in permissible filing. The point is not to haggle over how to analyse the notion 'mental file', but to ask: what is the best explanation of filing? We have three candidates on the table. I have argued that causal-functional mental file theory offers the best explanatory resources out of the three.

The prominence of Recanati's indexical model, and the detail and skill with which it has been articulated and defended, has perhaps had the unfortunate result of limiting how mental files have been understood. G&G move away from that model, but in the wrong direction. Positing files as particulars with a causal-functional role is a more promising way of using the idea of a mental filing system, and worth exploring further.

Acknowledgments

Thanks to two anonymous reviewers for this journal for detailed comments which greatly improved an earlier draft of this paper.

References

- Campbell, John (1987–88). Is Sense Transparent? *Proceedings of the Aristotelian Society*, 88, 273–92.
- Clarke, Henry (2018). Frege Puzzles and Mental Files. *Australasian Journal of Philosophy*, 92(2), 351–66.
- Dickie, Imogen (2015). *Fixing Reference*. Oxford University Press.
- Evans, Gareth (1985). Understanding Demonstratives. In *Collected Papers* (291–321). Oxford University Press.
- Fine, Kit (2007). *Semantic Relationism*. Blackwell.
- Goodman, Rachel and Aidan Gray (2022). Mental Filing. *Noûs*, 56(1), 204–226.
- Heck, Richard Kimberly (2012). Solving Frege's Puzzle. *Journal of Philosophy*, 109(1/2), 132–74. (Originally published under the name 'Richard G. Heck, Jr.')
- Lawlor, Krista (2001). *New Thoughts about Old Things: Cognitive Policies as the Ground of Singular Thought*. Routledge.
- Lee, Pong Shil (2022). Co-filing and De Jure Co-Referential Thought in the Mental Files Framework. *Erkenntnis*, 87(1), 309–345.
- Millikan, Ruth Garrett (1997). Images of Identity: In Search of Modes of Presentation. *Mind*, 106(423), 499–519.
- Millikan, Ruth Garrett (2000). *On Clear and Confused Ideas*. Cambridge University Press.
- Millikan, Ruth Garrett (2017). *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford University Press.
- Murez, Michael and Francois Recanati (2016). Mental Files: An Introduction. *Review of Philosophy and Psychology*, 7(2), 265–81.
- Recanati, François (2012). *Mental Files*. Oxford University Press.
- Recanati, François (2016). *Mental Files in Flux*. Oxford University Press.
- Sainsbury, Mark (2002). Fregean Sense. In *Departing from Frege* (125–36). Routledge.