Mental Filing, Continued…

Rachel Goodman
Aidan Gray
University of Illinois, Chicago

[Forthcoming in *Synthese*. Please cite published version]

According to traditional versions of the mental file theory, we should posit *mental files*—that is, mental representations with containment structure—to explain both rational relations between the attitudes, and the persistence of the attitudes across time. However, Goodman and Gray (2022) offer a revisionary interpretation of the file framework, according to which its explanatory commitments are better presented by positing *mental filing*, as a process, but not *mental files*, as mental representations with file structure. Goodman and Gray focus on a certain class of synchronic explanations, but files have also been thought to play an essential diachronic role in the maintenance and update of object-directed thought. This paper has two aims. First, we clarify the aims and commitments of Goodman and Gray (2022). Second, we extend their mental filing view to show how it can account for continued belief, change of mind and persistence of the attitudes in general.

Our cognitive relation to objects has an essentially diachronic dimension. We are not only able to think about the objects we encounter *when* we encounter them, but also able to form beliefs (and desires, and other attitudes) about them, to continue to believe (desire, etc.) about them, to change our minds about them, and so forth. And, it is both natural, and arguably required, to think that this involves forming representations of objects and collecting information about them, and maintaining or updating these representations and information over time, such that they are available for use in thought, inference, action, etc.

This raises a set of questions at the intersection of theories of reference, theories of the attitudes, and theories of cognitive architecture: how must object-directed representations and object-sourced information be arranged, maintained and marshaled, such that the subject can be said to continue to hold the same attitudes across time? What, in a representation of an object, must *persist* across time such that the subject can be described as having *changed her mind* about it? What is the rational basis for information collected at different times being deployed together in inference or action? What kind of *content* must attitudes about objects have, and how must this content be functionally realised, given the diachronic dimension of our cognitive relation to objects?

In philosophy, the *mental files* framework is a systematic attempt to theorise about attitude content, mental reference and cognitive architecture. In fact, one of the central explanatory roles claimed for mental files has been the idea that positing them puts us in a position to explain the diachronic profile of referential attitudes. In this paper, however, we take up and extend our own recent attempt (in Goodman and Gray (2020)) to offer a revisionary interpretation of the file framework. On this revisionary interpretation, the central, explanatory commitments of the file picture are better presented by positing *mental filing,* as a process, but not *mental files*, as mental representations with file structure. The current paper has two, related goals. First, we clarify the aims and commitments of Goodman and Gray (2020) (from hereon in, we refer to that work/its authors as 'G&G', in order to distinguish the content of the previous work from our current commentary on and extension of it). Clarifying G&G will put us in a position to achieve our second, and central, goal: to extend the *mental filing* interpretation of the file picture to take into account the role that files have been claimed to play in the formation, maintenance and update over time of object-directed propositional attitudes.

1. Explanatory Roles for Files

Before turning to Goodman and Gray (2020) some ground-clearing.

First, our discussion here is focussed on (what we think of as) the *philosopher's* notion of a mental file.[1] This involves the use of 'file-talk' to give an account of *propositional attitudes*. In particular, files appear in the philosophical literature as part of attempts to give a broadly Fregean account of the attitudes: that is, to give an account of the attitudes according to which more than just referential content is needed to explain their 'cognitive significance'. Files, or file-like structures, have also been posited in vision science and linguistics, most often to give accounts of perceptual states of various kinds and of reference and quantification in natural language.[2] Though some of our discussion of the notion of a file as it is used to theorise propositional attitudes might have application to literatures concerned with these perceptual and linguistic phenomena, we leave that question for future work.

Second, we follow G&G in taking Recanati's work as our main foil. This is for two reasons. First, he has done the most to flesh out the philosopher's notion of a mental file (that is, the use of files to theorise the propositional attitudes). Second, we have sympathy for many of the moves that he makes in doing this. We don't see the project as arguing against Recanati but, rather, as a kind of revisionary Recanati interpretation. Our question is: how can we take the things that seem to matter to Recanati and express them in a way that clearly displays the explanatory moving pieces? Our central claim is that the best way to do this is to write files, qua mental particulars, out of the canonical statement of the view entirely.

Finally, our project is guided by two related sets of questions. The first is foregrounded in G&G (2020):

[1] For an overview of the philosopher's notion, including a discussion of its status *qua* metaphor or theoretical posit, see (Goodman, 2024).

[2] For files in linguistics, see (Kamp, 1981), (Heim, 1982). For the relationship between that notion and the philosopher's notion see (Maier, 2016) . For files in psychology, see (Kahneman and Treisman, 1984), (Kahneman, Treisman and Gibbs, 1992), (Scholl, 2001). For the relation between that notion and the philosopher's notion see (Murez, Smorchkova & Strickland, 2020).

> Should we take talk of files as a *helpful metaphor* or as a *theoretical posit*? That is, is talk of files ultimately just a convenient way to introduce the commitments of the best version of the theory? Or is it an ineliminable theoretical posit of the best version of the file theory?

The second question is less explicitly thematized in G&G, but is crucial for seeing the motivation for our revisionary interpretation of the file framework:

> At what *explanatory* level do the various mechanisms posited by file theorists live? Which of the file theorist's claims are about the *content* of the attitudes? Which are about the *metasemantics* of the attitudes? Which are claims about the *functional implementation* of the attitudes?

As we see the matter, all of these questions are closely connected. It can be hard to pin down exactly what mental files are meant to explain, and how they're meant to explain it. In our view, this is partly because file theories seem to appeal to files at different levels of explanation, and therefore to bridge different explanatory projects. And, this is not always made explicit.

Given the importance, on our view, of distinguishing the explanatory domains of content, metasemantics, and functional explanation, our first task will be to say a little about how we understand them, and their relation to one another. But we note an important caveat up front: clearly these three domains are interrelated, and different people might look at the same phenomenon and assign it to a different level of explanation. We take it, though, that everyone in this area is operating with either an explicit or implicit sense of what distinguishes these levels of explanation. We think that part of the work required to properly understand the file-picture involves seeing how it divides the explanatory work between these levels of explanation, and asking, 'at which level/s, if any, do files do explanatory work?'.

Content Explanations

At least some proponents of the philosopher's notion of a file hold that files are part of a theory of the *content* of the attitudes. Recanati, for example, holds that files are Fregean senses (Recanati, 2015, p.12, p.71).

Given this, one would like to begin with some idea of what it is to hold that something is a feature of attitude content. But this is difficult, in that people operate with different, and often less than fully explicit, conceptions of what content is and of when some posited psychological feature counts as a feature of content.

On one conception, content is understood broadly truth-conditionally. On this conception, content is ultimately understood in terms of truth-with-respect-to-a-point-of-evaluation, be it a world, situation, context, or tuple of such things. Clearly, the claim that files are features of content does not employ this conception of content. That two attitudes are associated with different files does not imply that there is a truth-conditional difference between them. This is part of what Recanati is trying to capture in saying, for example, that there can be distinct files that contain all of the same predicates (2012, p.40). So the file-theorist must be working with a different conception of content.

We think the file-theorist's conception of content can be captured, roughly, with CONTENT:

> **CONTENT:** The content properties of an attitude are those properties that play a role in rational evaluation and rationalizing psychological explanation.

We take something like this idea to be in the background of most Fregean theorizing about the attitudes.[3] Of course, it is only substantive to the extent that we have a grip on the sorts of evaluations and explanations that count as rationalizing. Though offering a precise characterization of rationalizing explanation is not something we can do here, it suffices, for our purposes, to say that rationalizing explanations are person-level explanations that explain an agent's inferences and actions in a way that captures the reasons they have to make those inferences or perform those actions.

For example, file theorists will often say that, if a subject has attitudes that are associated with the same mental file, they thereby have attitudes whose content rationally licenses inferences that 'trade on identity' (Campbell, 1987). We take this to be an expression of the commitment concerning the explanatory role of content expressed by CONTENT. Similarly, file-theorists will explain the rational permissibility of holding attitudes whose truth-conditional content is incompatible in terms of those attitudes being associated with distinct mental files. We take this to be an expression of the same commitment.

Summing up: a central explanatory role for files is that files are meant to explain the rational status of inferences that trade on identity, and the possibility of Frege cases. The claim that files are features of content, along with a conception of content that links content to rational explanation, is the way they do this.

Metasemantic Explanations

Metasemantic questions are questions about what grounds content properties. So, for example, we might ask: in virtue of what does a subject have an attitude that refers to some object *o*? This is a metasemantic question.[4] In answering it, one will usually appeal to facts that are not themselves content facts.

---

[3] See, for example, (Loar, 1988), (Devitt, 1989), (Fodor, 1995), (Heck, 2002, 2012, 2014), (Almotahari & Gray, 2020)

[4] One might fairly wonder (as a referee did) whether, once we accept CONTENT, and plump for Fregeanism, we can really think of *reference* as a feature of content. After all, Fregeans hold that rationalizing explanation appeals to *senses*, not directly to referents. Doesn't that mean that the Fregean, if she accepts CONTENT, cannot think of reference as a feature of content?

This is reasonable enough. We think it is mostly harmless to talk as if file-theorists countenance reference as an aspect of content because they appeal to *non-descriptive* senses. Individuating such senses, for the file-theorist, involves specifying which object is at the other end of the ER relation that governs the file. So, strictly speaking, the relevant metasemantic question here is a question about non-descriptive sense. But an answer to that question must appeal to facts about reference determination. So we can think of reference-determination as an aspect of the metasemantic question about non-descriptive senses.

When it comes to the relationist—whom we introduce shortly and also understand to be committed to CONTENT—there is no barrier to thinking of reference as a feature of content according to

Parts of the file picture are clearly aimed at metasemantic questions. For example, Recanati holds that the reference of an attitude associated with a file is determined by the epistemically-rewarding (ER) relation that governs that file (2012, p.35).

Of course, what the meta-semantic questions are depends on what the content facts are. On G&G's preferred reconstruction of the file theory (which we spell out in Sections 3 and 4), file-theorists posit a particular kind of content fact (*coordination*) along with a particular metasemantic story about it.

Functional Explanations

Functional questions about the attitudes are questions about how they are implemented in causal information processing networks. Functional explanations in this domain typically posit some kind of underlying symbolic structure, along with information-processing mechanisms defined over that structure.

Clearly, claims at the functional level are interrelated with content claims and metasemantic claims. Plausibly, the functional properties of an attitude state constrain what content properties it can have. On some views, for example, if the functional implementation of an attitude doesn't put that attitude into the right sort of causal rapport with Xs, under the appropriate conditions, then the attitude cannot be *about* Xs. But it's important to note both *i*) that claims about the functional implementation of attitudes are not, *as such*, claims about attitude content, and *ii*) that there can be functional structure that is not reflected at the level of attitude content (as conceived according to CONTENT).

For example, prototype theories of concepts posit a certain structure for concepts-qua-representations, along with certain processes which are defined over that structure. But nothing about *rationalizing explanation* follows merely from positing prototype structure for concepts, along with, say, access mechanisms that appeal to that structure. It might be that these are facts about the implementation of attitudes that don't have any characteristic effect at the level of rational explanation or evaluation. So, for example, a difference in BIRD-prototype might explain a difference in reaction-time between two subjects in identifying a chicken as a bird. But this, by itself, need not imply any difference in the *content* (in our preferred sense) of their attitudes about birds. This difference might not make a difference to what each subject is rationally licensed to infer about birds, or to what bird-involving actions their attitudes give them a reason to perform.

We take the claim (often made by philosophical proponents of the notion of a mental file) that files are *vehicles* for thought, and the claim that files are mental particulars that *contain* token predicates, to take place at the level of functional explanation. But, as we will suggest shortly, it is not in fact clear how to understand them as playing a genuine/legitimate role at that level.[5]

---

their ideology. The 'second layer' of content that relationists appeal to is *coordination relations*. And these are relations *on* referential contents. So relationists are committed to reference as a feature of content.

[5] It is worth stressing that we take philosophical proponents of the notion of a mental file to be making claims about the vehicular structure of the attitudes and that, in specifying our interest in the 'philosopher's notion of a file', we do not assume this notion is intended as part of an a priori or 'non-empirical' theory. Rather, we ask: are claims made by file-theorists about the vehicular structure of the attitudes explanatory? Or, are they superfluous to the theory properly understood?

2. Goodman & Gray 2020

Our second task will be to rehearse and clarify the approach in G&G (2020), before applying it to the diachronic role for files.

The guiding question for G&G was: does the best version of the file theory treat talk of files as a helpful but inessential metaphor, or as a theoretical posit, which does genuine explanatory work? The strategy for addressing this question was to think about each theoretical role discussed above—content, metasemantic, and functional— and ask what explanatory work, if any, files were supposed to do at that level.

More narrowly, the particular focus of G&G (2020) was on the use of files to give an account of phenomena associated with synchronic versions of Frege's puzzle. That is, the question was: do files have a semantic, metasemantic or functional role to play in explaining the rational status of trading on identity and Frege cases? And the answer was: no. In their view, the explanations offered by the file theory can be given without positing files.

To understand how this goes, we should begin by noting that G&G are interested in the *justificatory status* or *rational permissibility* of inferences that trade on identity. Thus, they took the explanatory target to be a set of *synchronic facts*: the rational facts about a body of attitudes at a given time.

Their first claim was that files and file-structure are not an essential part of the file-theorist's explanation of the rational permissibility of inferences that trade on identity *at the level of content*. This will come as a surprise to anyone who has read file-theorists, particularly Recanati. He and other file theorists will often be found claiming that a subject is licensed to trade on the identity of the reference of two pieces of information if and only if they are contained in the same file. And, he claims that files are to be construed as non-descriptive modes of presentation (2012, ch. 3). Thus, the file picture seems to posit files *as Fregean contents*, whose role is to account for a subject's license to trade on identity. G&G's claim, however, was that file talk as it appears here–that is, as part of a content-level story about the rational permissibility of inferences that trade on identity–is, in fact, dispensable.

To explain this, we must introduce a technical term: 'coordination'. In G&G's use, *coordination* is a relational feature of attitude content. So, for example, take two subjects: A and B. Both A and B have beliefs with the referential content that George Eliot is an author and that George Eliot is from Nuneaton. But A would express both of those beliefs with the name 'George Eliot', whereas B would express one belief with the name 'George Eliot' and one with the name 'Mary Ann Evans'. The file theorist holds that A and B are in attitude states with *different content*. After all, they are in states that license different inferences (A but not B is entitled to infer that some author is from Nuneaton). *Coordination* is just a name for the representational relation that holds between A's two attitudes about Eliot, but not B's two attitudes about Eliot.

On G&G's interpretation of the file-framework, when Recanati, for example, says that a subject can trade on the identity of information when it is contained in the same file, the ultimate explanatory import of this claim is really a thesis about *coordination*. It is the claim that coordination is an *irreducibly relational* representational feature. That is, *coordination* is a representational relation that holds between two (or more) token attitudes, which is not determined by any representational feature that each of those token attitudes possesses considered *independently*, outside of their relation with one another.

This might seem like a bit of a leap, but here is the idea. The background against which Recanati makes his claim that a subject is licensed to trade on the identity of the reference of information if and only if it is co-filed is what we might call, the *traditional Fregean approach to coordination*. For the traditional Fregean, senses are descriptions and coordination is sameness of sense. This is to say, for example, that two attitudes about Eliot are coordinated only if they are associated with the same descriptive presentation of her (for example: as *the author of Middlemarch*). One of the functions of Recanati's claims that files are mental particulars, which are not individuated by their contained predicates, and that the co-filing of information explains the license to trade on the identity of its reference, is to insist that *coordination* is not explained by a 'match' of descriptive information. That is, the point of saying *coordination* is explained by the filing of information in the same file is to reject the traditional Fregean story about coordination.[6]

G&G's key claim is, essentially, that we can reject the traditional Fregean account of coordination without giving files themselves an explanatory role to play at the level of content. We can simply say that coordination is a representational feature in its own right. It need not be explained by the matching of some representational property, be it descriptive or otherwise. This is what G&G called *relationism* about coordination.[7] On this way of thinking about things, the temptation to say that files are *senses*, and that coordination is *sameness of sense*, is just an unnecessary holdover from the traditional Fregean picture. Coordination need not be explained by the sameness of *any* representational property. Therefore, insofar as the explanation of *coordination* is the central role of files at the level of content, we have no essential need for files here.

Though it is less explicitly thematized, G&G's second claim is that existing philosophical versions of the file-theory don't give files a genuine *functional, implementational* role to play. That is, on existing versions of the file-theory, files are not essentially involved in a functional, implementational explanation of the rational permissibility of inferences that trade on identity.

The general idea behind this claim is this: if we posit files as part of the functional/implementational story of the attitudes, we had better posit psychological processes, or mechanisms, that are sensitive to *file-structure*. But what one finds when one looks at existing philosophical file theorists is, instead, that they tend to posit representations with file-structure, and then tend to do their best to avoid any functional commitments that are connected to this structure.

Here we should pause to be fully explicit about something: G&G take the file theorist *seriously* when she says that files are mental particulars that are collections of the monadic predicates to which a subject is doxastically committed. This is a specific claim about the implementation of the attitudes. It is not simply a general commitment to the representational theory of mind, but rather a commitment to a particular claim about vehicle structure.

To illustrate this commitment, we can contrast it with some potential alternatives. One could hold that object representations, instead of being collections of monadic predicates to which the subject is doxastically committed, are collections of monadic predicates which capture what a subject *desires* about the referent of the file. This would have some of the attractive

[6] See, especially, the discussion in (Recanati, 2012, Part IV) and (Recanati, 2015, Part I).

[7] On relationism, see (Taschek, 1995), (Fine, 2007), (Pinillos, 2011), (Heck, 2012), (Pryor, 2016), (Gray, 2017), (Goodman & Gray, 2020).

features of positing files (for example, it would secure the possibility that a subject has distinct modes of presentation on the same object while believing all of the same things about it under each mode of presentation). But it would raise immediate questions about how to understand the implementation of attitudes other than desire (just as there are questions for the file-theorist about the implementation of attitudes other than belief (see G&G p.221)).

Or one could hold that object representations, instead of being collections of monadic predicates to which the subject is doxastically committed, are collections of *dyadic* predicates to which the subject is doxastically committed. This would also have many of the attractive features of the file approach. But it would immediately raise the question of how *non* dyadic belief is implemented (just as there are questions for the file-theorist about the implementation of non-monadic belief, as we will rehearse shortly).[8]

We assume that nobody would be particularly attracted to either of these alternative accounts. And this would, at least in part, be because there is really no reason to think that monadic desire, or dyadic belief, have any privileged implementational status. G&G's claim is that the same goes for monadic belief. Or, more carefully, that file theorists have not made any attempt to *convince* us that monadic belief is implementationally privileged.

This dialectic comes out, though perhaps in a slightly compressed way, in G&G's discussion of the 'containment puzzles' for the file approach. One such puzzle is about relational belief: if beliefs and other attitudes are implemented by files containing predicates, then how are relational attitudes, like the belief that *a* loves *b*, implemented? Does this involve the predicate 'x loves *b*' being stored in a file that refers to *a*, the predicate '*a* loves x' being stored in a file that refers to *b*, or both?

Suppose that the file theorist takes the third option and claims that relational beliefs are 'redundantly' encoded in the files for both relata (whereas monadic attitudes are non-redundantly stored only in the file for their referent). If this is a well-motivated functional, implementational claim, shouldn't it be visible somehow in the processes involved in forming, maintaining, or using those attitudes? Might we not expect, for example, that certain interventions might destroy one, but not both, of the redundantly stored predicates? Or, that certain processes will work differently if they are operating on redundantly stored beliefs as opposed to non-redundantly stored ones?[9]

The point is *not* that these expectations are not *in fact* fulfilled in the case of minds like ours. For all we know, they are. The point is rather that philosophical file–theorists tend not, in fact, to be motivated by appeal to them. They make no attempt to identify phenomena such as

---

[8] It is worth mentioning that these questions could be answered, albeit in a non-elegant way. Suppose that S has a dyadic file on the ordered pair of objects <*a*,*b*>. This file consists in a set of dyadic predicates which the subject believes apply to that pair ('*x* loves *y*', '*x* is taller than *y*', etc). How could the monadic belief that *a* is French be implemented? It could simply be the presence of the dyadic predicate `x is French and *y* = y' in the <*a*,*b*> file. Here the nature of dyadic files forces us to give a more elaborate account of monadic belief. But the same thing is true of (traditional) files; their nature forces us to give a more elaborate account of dyadic (and poly-adic) belief.

[9] We have been following file-theorists in acting as if a subject's cognitive relation to properties can be fully captured by the predicates that occur in object-files. But, plausibly, subjects will keep track of properties too. And an analog of trading on identity occurs here (consider: John is French, Sally is French ∴ There is a something that John and Sally have in common). So the file theorist should posit files on properties as well. But now the reduplication problem gets worse. The belief that *a*R*b* should be redundantly stored in the *a* file, the *b* file, *and* the *R* file.

these, which would properly motivate particular functional, implementational claims (like the claim that relational beliefs are redundantly stored, or, instead, stored in only one of the files referring to the relata).[10] All of this suggests that, ultimately, they do not see the kind of file-structure that they posit–that is, structure which would force a *choice* between the options above concerning the implementation of relational belief–as playing a functional, explanatory role.

G&G's (2020) claim about the containment puzzles is that, although file-theorists might be able to make various fixes, or give various interpretations of the containment claim, which solve these puzzles, if one doesn't posit file-structure to begin with, the puzzles don't arise.[11] From our perspective, the more *general* point, not explicitly made by G&G, is that positing files as part of one's functional, implementational story about the attitudes ought to have various downstream effects, visible in relation to the psychological processes that act on files. The containment puzzles are simply one way of pointing out that file-theorists are not actually in the game of characterizing such effects. Their general approach is usually, instead, to weaken or reinterpret the claim of containment so as to dissolve the puzzles.[12]

We will emphasize again: the target of this discussion of the functional, implementational explanation of the attitudes is *file-structure, as such*. That is, the target is *not* the representational theory of mind, or the general idea that our attitudes about objects are implemented in mental vehicles. G&G themselves appeal to the existence of object-representations in their 'mental filing' story. This is entirely consistent with their rejection of files as mental particulars. A file theorist who posits mental files as mental vehicles thereby posits specific structures: collections of monadic predicates to which a subject is doxastically committed. The point of the discussion above, and of G&G's discussion, is to take them seriously in this, and to ask whether those structures are doing any explanatory work.[13]

Finally, having claimed that mental files do not have a semantic or functional, implementational explanatory role, G&G ask: does file-structure have a *metasemantic* role to play? Suppose we accept G&G's claim that the file theorist should hold that coordination is irreducibly relational at the level of content. We might still ask: what *grounds* the content-level

---

[10] In a footnote, in response G&G's rehearsal of this worry, Tałasiewicz (2022) floats the possibility that the file theorist could allow that relational belief is only implemented in one of the relata files, that this asymmetry could correspond to how 'important' the relevant belief is in the subject's conception of each relatum. This is an interesting idea. If anything like this were true, it would, perhaps, motivate file-structure. But note that Tałasiewicz does nothing to support this hypothesis or spell out this notion of importance (what would establish it?). And the discussion is consigned to a footnote. It seems fair, then, to call it post-hoc speculation, rather than essential motivation for the files-as-particulars view. Still, we should say: this is precisely the sort of thing that file theorists *should* explore if they want to justify file-structure.

[11] The broad idea that file-theorists have over-committed with respect to the structure of object representations is not new to G&G. They cite Pryor (2016) as inspiration for their discussion.

[12] See the discussion on pp. 4-5 in (Goodman & Gray, 2020). See also (Tałasiewicz, 2022).

[13] Thus, some apparent defenses of files-qua-particulars, in response to G&G's criticisms, are not obviously defenses of that view at all. Clarke (2021) takes himself to be defending the files-qua-particulars view from G&G's criticism. But, he sometimes seems to proceed as if that view is simply equivalent to the claim that there is *some functional organization or other* that implements a system of mental filing (see pp. 666-667, in particular footnote 9). On the other hand, he often talks as if files are to be functionally characterized as collections of monadic predicates to which a subject is doxastically committed. But he does not, as far as we can tell, make any functional claims which make use of that structure in particular.

facts about coordination? Perhaps mental particulars with file-structure are needed to play this grounding role.

G&G respond to this suggestion, by offering their 'mental filing' proposal about the grounding for coordination. So, our next task is to rehearse and clarify this proposal.

3. Mental Filing

G&G offer a Recanati-inspired story about the grounding for coordination in thought. The story posits *filing,* as a process, but not *files*, as mental particulars. In overview, the account is that object-representations (which do not, themselves, have containment structure) carry coordinated content in virtue of occupying a certain functional role with respect to (*i*) each other, and (*ii*) the information-management processes that are associated with them.

With respect to (*i*), G&G hold that, when object representations are *coordinated*, this is partly in virtue of the fact that the subject is disposed to 'trade on' the identity of their referents, under appropriate circumstances*.* This means that, *if* an occasion arose in which trading on identity would be useful for the agent's epistemic, practical or inferential purposes, then she would, all things equal, do so. For example, imagine that I am currently investigating the question of whether any birds have both yellow bills and graduated tails. Imagine, also, that I have a belief with the referential content 'that bird has a yellow bill' and another with the referential content 'that bird has a graduated tail', but that I *do not* trade on the identity of the reference of these beliefs to draw the conclusion that there *are* birds with yellow bills and graduated tails (thereby answering the question of my investigation). This would be evidence that my attitudes are *not* coordinated. If they *were* coordinated, then, ceteris paribus, I'd make the inference.

With respect to (*ii*), G&G claim that two representations *a* and *b* are coordinated *only if* there exists an ER relation, R, associated with them, such that any representation, *c*, which is a deliverance of R, would be such that the subject would be disposed to trade on the identity of the reference of *c,* with *a* and *b*, under appropriate circumstances. For simplicity, we'll sometimes put this by saying that *a* and *b* must be *governed* by an ER relation. Since G&G say they conceive of ER relations, in broad terms, as relations to objects that deliver information about them, this condition essentially states that, where there are *coordination relations*, there must be information delivering processes in place, which *produce* the disposition to trade on identity cited in (*i*). Thus, (*ii*) reflects the view that the activity of collecting, sorting and updating information plays an essential role in grounding the rational status of inferences that trade on identity–that is, in grounding the existence of coordination relations.

The claim that coordination is partly grounded by the existence of ER relations that produce dispositions to trade on identity is motivated by the idea that the *rational credentials* of inferences that trade on identity rely on our abilities to reliably track objects. To be convinced of this, imagine a disposition to trade on identity that was unrelated to a tracking ability. Imagine, that is, that information comes in from different sources and is sorted and attached to object-representations randomly, such that the creature in question is then disposed to trade on the identity of reference in randomly generated ways. G&G follow Recanati in holding that this creature's inferences would lack positive rational status. A disposition to trade on identity only has rational credentials when it is the typical downstream effect of a process that reliably sorts information, so that information from the same object is treated as such. In other words, like

Recanati, G&G adopt a conception of rationality that makes a connection between rationality and non-accidental cognitive success.

Note that, with condition (*ii*), G&G make ER relations part of the grounding for *coordination* but aim to do this in a way that allows that two object representations can be coordinated without having been the upshot of a *single* ER relation (however one is individuating ER relations). What's required for two object representations to be *coordinated* at a time is that they are associated, at that time, with *some* ER relation that can produce representations such that the subject is disposed to trade on the identity of their reference. Thus, the *mental filing* picture is meant to allow for what people who posit files have in mind when they posit conversion, incremental conversion, temporally coarse-grained files, etc.

G&G's metasemantic story about coordination does not make reference to mental files. Essentially, they view files as a kind of 'middleman', which can be cut out of the picture: they are not needed and don't do independent explanatory work. *Recanati* posits ER relations that govern files, which then explain dispositions to infer in certain ways. He thereby claims that ER relations ground the rational status of inferences *via files*. G&G cut out the files: they propose a view on which ER relations generate certain inferential dispositions and thereby ground their rational status.

Both G&G and Recanati appeal to ER relations, but neither says very much about what ER relations are. Without giving a full account of what an ER relation is, we would like to make some assumptions about them explicit, since we take them to be part of the mental filing picture. This will help us to clarify the commitments of the mental filing approach. The assumptions we outline will also be relevant when we extend the picture in Section 6.

We will assume three things: *a)* ER relations hold in virtue of causal processes (or dispositions to undergo causal processes) that generate *predications*; *b*) there is a canonical individuation of such processes; and *c*) when an ER relation generates a *predication*, there is typically a fact of the matter about the object from which the information employed by the predication derives.

Our next step is to explain these assumptions, along with some of their implications.

With respect to (*a*), Recanati holds that the ER relations that govern a file deliver predicates that are added to the file. We want to take this idea, and strip it of any commitments about vehicle *structure*. ER relations generate *predications*: token representations of an object bearing a property.[14] So for example, if a subject stands in a perceptual ER relation to a bird, *b*, flying across the sky, this ER relation may result in representations of *b*'s being red, *b*'s being large, and so forth. This is, in principle, consistent with many different claims about the *structure* of the vehicles that carry these contents.

With respect to (*b*), G&G follow Recanati in assuming that ER relations can be individuated. For example, they claim that two representations of the same object are coordinated at a time only if there is a single ER relation that *governs* both of them at that time. We don't pretend that it will be straightforward to characterize the individuation conditions for ER relations, but we follow Recanati in assuming they can be individuated. For example, we

[14] Here and throughout the paper, the noun 'predication' is used to refer to representations that are characterised by their representational *content*, not their form. Forms of the verb, 'to predicate', are used throughout the paper, to characterise the content of predications. We sometimes call this content, 'predicational content'.

assume that, if a subject is unknowingly, perceptually attending to the same object *twice at the same time* (perhaps because of a well-placed mirror), then she stands in two distinct ER relations to that object. It may be that, ultimately, there is some way to rid the mental filing view of this assumption, but it will play a role in our way of extending the view in Section 6.

With respect to (*c*), we assume that *predications* can 'employ' or be 'based on' *information*. For example, a predication representing that an object, *o,* is F might be based on perceptual information, or on testimonial information. We do not offer an account of information (though, neither does the traditional file theorist) but we assume a broadly causal account. And, recall, (*c*) is the assumption that there is a fact of the matter about the source of any piece of information that a predication may be based on or employ (e.g., it might be a fact that its causal source is *a*, rather than *b*, or no object at all). By making this assumption, we allow (in principle) for a *distinction* between the object represented by some predication, and the object that is the source of the information on which the predication is based (Lawlor, 2001. p.62). For example, when a perceptual ER relation generates a predication to the effect that an object *o* has a red tail, the perceptual information this predication is based on will (usually) have some object as its causal source. This object is usually the object represented to have a red tail—in this case, *o*—but it need not always be.

For example, imagine that S is tracking *b* amongst a flock of birds. She is, mostly, successful at tracking *b* as it darts back and forth—that is, most of the predications generated by her perceptual tracking relation (that *b* has white wings, has a large beak, etc.) are based on information whose causal source is *b*. But, for a moment, unbeknownst to her, she visually attends to a distinct bird in the flock. It is therefore the redness of *this* bird's tail that is the causal source of the information employed in her predication that *b* has a red tail. In this kind of case, we want to say, the coordinated body of predications generated by the perceptual tracking ER relation in fact contains information causally derived from different objects (from two distinct birds). However, all of these coordinated predications represent *b* as being this way or that. (*c*) is the assumption that allows the mental filing theorist to say this. It thereby allows her to accommodate the sort of thing that file-theorists want to say about cases in which files end up containing some wrongly sourced information, without this entailing referential failure, ambiguity or indeterminacy.

Notice, also, that this distinction—between the object that is the causal source of information employed by a predication and the object represented in that predication—allows us to characterize one sense in which ER relations are *epistemically rewarding*: they *typically* deliver information that is causally sourced in the *same object.* A perceptual tracking relation is epistemically rewarding because it typically (though not infallibly) supplies information from a single thing.[15] This sense in which ER relations are epistemically rewarding is also what makes them suitable to provide the grounding for *coordination*. To mark this, we will sometimes say that it is 'no accident' when deliverances of a single ER relation carry information from the same object. In our view, it is because of this that the objectual components of the predications generated by a single ER relation are *coordinated.*

[15] It's fair to ask about the sense in which ER relations like perceptual tracking 'typically' supply information from a single thing. For example, one might offer a teleo-functional elaboration of this idea, or a non-reductive epistemic elaboration. We don't take a stance on this here. Rather, our point is that *some* such idea is in the background of the file-theorist's talk of ER relations.

It should be stressed, again, that G&G in no way deny the existence of mental particulars. In fact, they explicitly allow that coordination relations may hold between object-representations (p.212, 214 amongst others). However, their picture does without mental particulars *with file-structure.* They posit coordination relations between object representations, grounded in information-marshaling processes and inferential dispositions. What they don't posit is the containment of predicates in files.

In Section 2, we claimed that, although it is natural to interpret the file-theorist's claim that files are mental particulars as a claim about functional structure, she does not in fact specify any functional explanatory work for file-structure *in particular* to do. In this sense, file-theorists are *over-committal* about the functional structure they posit, given their central explanatory commitments. With their metasemantics for coordination, G&G offer a way to take the broad goals of the file-theory and package them so as to posit only as much functional structure as is required by the explanations that the theory aims to offer. G&G themselves are non-committal about the nature of what they call 'object representations'. They claim only enough structure for them as is required to satisfy the explanatory ambitions of the file-approach: that they carry referential content, and that they can stand in certain functional relations to ER relations and to each other.

4. A Role for Diachronic Files?

G&G claim to explain the rational relations between beliefs that hold at a time, and the rational permissibility of trading on identity, using the central explanatory tools of the file theory, without positing mental files. However, this leaves open that there are *other* explanatory roles that justify positing mental files. Here, we would like to consider one such possible role, often posited for mental files, and to ask if a 'mental filing without mental files' view can handle this role. Our answer will be that it can.

One of the earlier philosophical applications of the file-picture is offered by John Perry (1980), as part of his account of *continued belief*. Perry is concerned to give an account of what we think of as an ordinary notion of 'continuing to believe the same thing', and posits files as part of this account.

Here is the sort of case that Perry focuses on. Imagine that I am at a party, with many people in attendance whom I have never met. I'm intrigued by the guests and I spend the evening people-watching. At t_1, I believe what I would express with, "That man by the bar regaling people with stories is French". At t_2, I believe what I would express with, "That man on the balcony looking tired and emotional is French". This *could* be a case of me continuing to believe the same thing at t_2 that I believed at t_1. The question is: what would have to be true of me for this to be a case of continued belief?

Clearly, the fact that the thought I expressed at t_1 and the one expressed at t_2 are about the same man is *not sufficient.* I may have lost track of the relevant man during the evening (perhaps he has changed his outfit, or is looking much less glamorous and more disheveled at the end of the evening than he was at the beginning). It might therefore be, so far as my *cognition* is concerned, just as if my belief at t_1 and my belief at t_2 were about two different men. However, Perry also thinks that sameness of any other *semantic* feature of my attitudes at t_1 and t_2 is *not necessary*. For the case to be one of continued belief, it is not as if I need to have some uniquely identifying descriptive take on the man who was at the bar and then on the

balcony, which persists from t_1 to t_2. Similarly, I don't need to have any particular egocentric relation to the man—expressible with a single indexical—that persists from t_1 to t_2.

According to Perry, trouble comes when we try to give an account of continued belief in terms of a relation between the *semantic* facts about my attitude at t_1 and the *semantic facts* about my attitude at t_2. Instead, he thinks we should focus on the causal/psychological process that connects my belief state at t_1 to my belief state at t_2. Perry calls the kind of connection via a *causal/psychological process* that would make for continued belief, 'internal identity'. However, he also claims that internal identity involves the maintenance and development of a single *mental file* over time. Thus, he gives an account of continued belief that makes reference to sameness of file over time: if, in our case above, the attitude at t_2 is an expression of the same *file* as the attitude at t_1, regardless of what predicates are in the file at the two times (apart from 'is French', which must presumably remain in it throughout), we have a case of continued belief.

We would like to ask: does theorising continued belief justify positing files as mental particulars? Or, can essentially the same theoretical work that Perry envisages be handled within the *mental filing* picture?

Given the venue in which this paper appears, we want to note, before proceeding, that there is an analogy between questions about continued belief and questions about what it takes for a representation at time t_2 to be a memory of an experienced event at time t_1.[16] The analogy is worth pursuing in detail but it has limits. Firstly, though the latter questions are about episodic memory, to the extent that the former questions concern memory, they seem to be about semantic memory. Secondly, as we will emphasise in Sections 5 and 6 of the paper, continued belief is itself but one instance of a broader phenomenon (which should be seen as the topic of this paper): *coordination* between attitudes across time. Questions about diachronic rational relations between, and update of, the attitudes arise even when there is no shared or preserved predicational content across time. And, file theorists (including, arguably, Perry, though he frames his discussion in terms of continued belief) are interested in the rational and psychological relations that can hold between a coordinated body of object-representations at t_1 and a coordinated body of object representations at t_2. This is a wider concern than the question about maintaining a particular belief from t_1 to t_2, though it encompasses that question. In part for these reasons, in part to avoid distraction from the focused argument of the paper, and in part because we think the analogy between questions about continued belief and questions about episodic memory deserves more detailed and careful examination than we are in a position to give it, we will confine our further (preliminary) remarks about it to footnotes.[17] However, we hope that the appearance of this paper in a topical collection on reference and

[16] For example, a referee's comments suggest an analogy between the view that continued belief requires sameness of file across time and the view that episodic memory requires a 'trace' of the remembered event. However, as we see it, these views are not analogous. Proponents of traces—even the apparently most strict ones—do not require that a token representation persists from t_1 to t_2 but rather allow that it might be a sequence of representations that runs from t_1 to t_2 (see, for example, Martin and Deutscher's talk of 'a succession of states' (1966, pg. 166)). This succession of states is even sometimes called a 'process' (for example, in (Werning, 2020, pg. 304). So even the strictest version of the causal theory of memory does not provide analogical support for diachronic files. At best, it provides support for continued mental *filing* (see Section 6), since there is a shared appeal– by mental filing theorists and causal theorists of memory— to a causal chain of mental representations linking a representation at t_1 to a representation at t_2.

[17] See fn. 16, 19, and 20.

remembering will stimulate discussion of what we might learn by bringing these two literatures into closer dialogue.

With that said, we return to our question of whether theorising continued belief justifies positing files as mental particulars, or whether it can be theorised from within the *mental filing* picture. To get some clarity about this question, we want to start by distinguishing three distinct lessons that Perry takes from his discussion.

The first of these lessons is a negative claim: that continued belief is not a matter of any match between truth-conditionally specifiable content properties at distinct times. With respect to this claim, we should note that Perry thinks that all content properties are truth-conditionally specifiable. This means that he *also* asserts the *stronger* negative claim that continued belief is not constituted by sameness of *any* kind of content across time. In contrast, neo-Fregean file-theorists like Recanati don't accept the picture of content that motivates this claim (and neither do we). This means that they (and we) won't accept the stronger negative claim (more on this in Sections 5 & 6). But we are perfectly happy to accept the weaker negative claim. Perry's examples do seem to show that much.

The second lesson that Perry takes from examples like the one above is a minimal positive claim: that continued belief involves a causal/psychological process that connects a belief state at an earlier time to a belief state at a later time. We call this positive claim 'minimal', because it treats continued belief as partly a matter of causal processes on mental representations, but does not involve any commitment about the structure of those representations. This means that the mental filing theorist can also accept this claim.

However, the third lesson that Perry draws is a maximal positive claim: that continued belief involves causal/psychological processes *that modify persisting mental files*. We call this positive claim 'maximal' because it holds not *only* that continuing to believe involves a particular causal/psychological *process*, but also that this process involves the manipulation of file-like mental particulars. For us, this will be the sticking point. The mental filing theorist will have to reject it. But, in truth, we do not actually see an argument for the maximal positive claim in Perry's discussion.

If we accept that continuing to believe is a matter of a certain process defined over mental representations (and we, on behalf of the mental filing theorist, are happy to accept this), the question becomes: what additional features of cases like the one above are supposed to motivate file-structure in those representations? We take file structure to have specific features: it involves the idea that the propositional attitudes are implemented functionally by containment of predicates in files, and so it privileges monadic belief in its way of theorising the attitudes. But, we don't see what this structure has to do with Perry's argument about continued belief. Instead, we think that the file-free *mental-filing* picture can handle the notion of continued belief in terms of causal/psychological processes that act over object representations.

Furthermore, we see this to be a suggestion very much in line with the spirit, if not the letter, of Perry's view. For, though he does posit files as part of his story about continued belief, Perry also suggests that the cash value of this is to be found in his claim about the causal/psychological connection required between one's belief states at different times:

> '…What is essential in these metaphors and analogies is a path from the production of texts at one time back to the original perception of (or other introduction to) the source at

> an earlier time. This path in the mind plays the role of an object in the world. So our notion of internal identity, and so ultimately of believing the same thing, depends on the identity of the internal causal path or chain.'—Ibid, p. 88

There is, in our view, little daylight between Perry's view and the *mental filing* theorist's claim that mental files should be viewed more as helpful metaphor, than as essential theoretical posit.

5. Diachronic Files in Recanati's Frege/Evans-inspired Framework?

Before explaining in slightly more detail how the mental filing view will handle the putative diachronic role for mental files, we want to pause, and turn from Perry back to Recanati. We do this to consider a potential motivation for diachronic files, which can also be framed in terms of questions about continued belief, but tends to be discussed instead under the banner of *Fregean Cognitive Dynamics*.

As we understand him, Recanati would agree with Perry's claim that continued belief does not consist in a match between truth-conditionally specifiable content of belief states at different times. However, he rejects Perry's restriction of content to truth-conditionally specifiable content. Instead, he takes a Frege/Evans-inspired approach to theorising the propositional attitudes, both synchronically and diachronically. What we'd like to do now is reconstruct, in our own terms, the Frege/Evans inspired approach to diachronic individuation of the attitudes.

Recall that we started, in Section 1, with the big-picture idea expressed by CONTENT: that the content properties of attitudes are the properties that play a role in rationalizing explanations and evaluations. If we accept this, we want to suggest, we can probe the diachronic individuation of attitude *content* by looking at the structure of diachronic norms of rationality.

It is not completely obvious that there are any essentially diachronic norms of rationality. If there are not, this would imply, given CONTENT, that diachronic issues do not introduce any motivation for adjusting our picture of content. There would still be, perhaps, questions about the ordinary notion of *continuing to believe the same thing*. But it would not be obvious why that notion should be given a content-level gloss (recall that Perry sees no reason for this). Diachronic issues, in that case, might be confined to the levels of functional and metasemantic explanation.

We will set this possibility aside and assume that there are essentially diachronic rational norms. Suppose, for example, that there is something like a diachronic norm against capricious changes of mind. Next, suppose that we try to frame it merely in terms of referential content. It would look something like REFERENTIAL DIACHRONIC NORM:

> **REFERENTIAL DIACHRONIC NORM**: If S has a belief with the referential content that <*o is F*> at t_1, and doesn't gain relevant information between t_1 and t_2, then S should not have a belief with the referential content that <*o is not F*> at t_2.

This is clearly too simplistic as a norm. But the relevant thing, for our purposes, is that, if anything in the vicinity of a norm like this is wanted, referential content is too coarse-grained to capture it. To see this, note the following. If I believe what I would express with "George Eliot is

an author" and "Mary Ann Evans is not an author" at t_1, and I believe the *same things* at time t_2, I would count as violating REFERENTIAL DIACHRONIC NORM in virtue of my 'George Eliot' belief at t_1 and my 'Mary Anne Evans' belief at t_2. But, continuing to believe that George Eliot is an author and that Mary Ann Evans is not an author from t_1 to t_2 is not a capricious change of mind (it's not a change of mind at all!). So, the norm expressed in terms of referential content gives us the wrong results. To capture a plausible version of the wanted norm, we need to appeal to some finer-grained conception of content.

We take this to be the sort of motivation that Fregeans like Recanati have for positing diachronically persisting files, construed as dynamic senses. That is, Recanati wants to be able to say that I could count as changing my mind from t_1 to t_2 only if my thought at the two times involves the same *non-descriptive sense* (constituted by a *file*). From their perspective, there is a role for diachronic sameness of fine-grained content, and diachronic files play this role.

To start off with, the view that diachronic files account for continued belief and change of mind is going to raise challenges generated by fusion and fission cases. These will not arise if we don't posit files as continuants. Recanati of course knows this, and this is why he is in fact somewhat ambivalent about positing files as diachronic entities.[18]

Furthermore, just as there are apparently awkward consequences of the containment picture, which file-theorists must massage away in the synchronic domain (the 'containment puzzles'), there will be new, awkward consequences introduced by a diachronic application of the picture. For example, G&G note that insofar as trading on identity is licensed between beliefs and attitudes of other types (e.g. desires), and because *containment* is meant to be a kind of doxastic commitment, the file theorist will have to weaken their claim that trading on identity requires co-containment (pp.207-208) This is because no predicates in a subject's file correspond to an agent's desires. An analogous issue will arise diachronically. We need to understand *continuing to desire the same thing* just as much as we need to understand continuing to believe the same thing. And for the same reason just mentioned—that file-containment is supposed to be a kind of doxastic commitment—we will not be able to understand continuing to desire the same thing in terms of containment in the same diachronically-persisting file.

But, putting these issues aside, we want to ask the same question G&G asked about the file-theorist's explanation of the rational facts about a body of attitudes at a given time: that is, is there any *motivation* for positing files here? It is certainly *natural* to appeal to files in this context if one has *already* posited them to make sense of the synchronic phenomena. But if G&G are right that the synchronic phenomena don't motivate positing file-structure, is there any *new* motivation here?

6. The 'Mental Filing' Story, Continued…

Our answer is that there is not. And, from our perspective, the dialectic that establishes this will be analogous to one pursued by G&G in their discussion of the possible synchronic role for files. We will ask, of the different possible levels at which files *might* do explanatory work, whether they in fact do work at that level.

[18] See (Recanati, 2015, 2016, 2021). For related discussion, see (Prosser, 2019).

At the level of *content*, we claim, what is required to capture diachronic rational norms is just that *coordination* is diachronically applicable–that is, that representations of the same object *at different times* can stand in a representational relation that constitutes there being 'rational relations' between them. This means that, when we come to express our norm against capricious changes of mind, for example, it would look something like the following COORDINATION-BASED DIACHRONIC NORM.

> **COORDINATION-BASED DIACHRONIC NORM**: If S has a belief with the referential content that <*o is F*> at t_1, and doesn't gain relevant information between t_1 and t_2, then, at t_2, S shouldn't have a belief with the referential content that <*o is not F*>, in which the representation of *o* is coordinated with the belief at t_1.

Notice, this norm does not treat our agent from Section 5, who maintains two beliefs from t_1 to t_2, which she would express, respectively, with "George Eliot is an author" and "Mary Ann Evans is not an author", as having changed her mind. And it makes no reference to files.

As with G&G's claim about synchronic coordination, our claim about *diachronic coordination* is that we do not need to ground it in *sameness of sense*. COORDINATION-BASED DIACHRONIC NORM introduces a diachronic role for fine-grained content, in the form of *diachronic coordination relations*, but it does not introduce diachronic *sense*, or diachronic files. We see the desire to posit non-descriptive senses, and to thus explain coordination as *sameness of sense*, as a holdover from the traditional Fregean picture. And, we do not see any reasons offered by the file-theorist to think that diachronic coordination is constituted by a *match* in content properties, rather than by irreducibly *relational* content. In this sense, any 'file talk' that may appear in the file-theorist's content-level story about continued belief or diachronic rational relations between attitudes, seems dispensable.

Our next question is about the functional implementation of diachronic coordination: in particular, does it involve files as diachronically persisting mental particulars? Here, the question should be the same one we pressed in our discussion of G&G, in Section 2. The theorist who posits mental particulars with file structure is making a substantive functional/implementational claim: that object representations are structured as clusters of monadic predicates (they have containment structure). To be legitimate as a functional, implementational claim, this must be reflected in some downstream effects. If it is not, then it is superfluous, and the file-theorist has needlessly overcommitted.[19] With respect to *diachronic coordination*, we see no new reasons offered by file-theorists to think that object representations with containment structure are required. There *may* be a need, in our functional, implementational story, for object representations to serve as the relata of *diachronic coordination relations*. But this is a different claim, and the *mental filing* picture does not rule this out.

[19] There is an analogy here between the mental filing theorist's claim that there is no explanatory work done within the file theory by mental particulars with file structure, and Robins (2016, p. 3000) claim that Martin and Deutscher over-commit concerning the structure traces in developing their causal theory of memory. Martin and Deutscher (1966) claim that memory traces are "structural analogues" of remembered events, but Robins argues that this a stronger claim than they need: given what they are trying to achieve, it is enough for them simply to say that the traces are mental representations of remembered events (whatever the structure of such representations turns out to be).

Our final question is about the metasemantics for diachronic coordination: does the file-theorist's story about this really require her to posit files? And again, our strategy here mirrors G&G's approach to the synchronic case. Recall that, instead of mental files, G&G's metasemantics for synchronic coordination posits object representations, connected by dispositions to trade on identity, governed by ER relations. They called this a mental filing story. To extend this story diachronically, we need to talk about how this two-aspect, functional structure develops across time. Our slogan here will be this: the diachronic metasemantics for coordination will involve *continued mental filing* (but no diachronic files).

As with G&G's original mental filing story, we do not take ourselves to be disagreeing with file theorists, exactly. In fact, the story below is largely inspired by the sorts of things that file theorists say about file-maintenance over time (particularly, Perry' s (1980, 1997) talk of internal identity and information-games, Lawlor's (2001) talk of policies that reliably maintain intentional relations to objects, and Recanati's (2012) appeal to ER relations). Essentially, we will take the parts of those approaches that we take to be explanatory, and re-write the theory without mentioning files.

To see how this will go, it helps to pause and compare Perry with the *mental filing* theorist. The *mental filing* theorist agrees with Perry that continued belief cannot be understood in terms of a match in truth-conditionally specifiable content between an attitude at an earlier time and one at a later time. However, following *Recanati*, she understands continued belief (and continuation of other kinds of attitude) to involve fine-grained content facts. For her, the content-theoretic explanation of continued belief involves specifically relational content–that is, coordination relations that hold between a belief at an earlier time and a belief at a later time (recall, she sees Recanati's claim that continued belief involves the continuation of files, construed as Fregean senses, as an unexplanatory hold-over from the traditional, descriptive Fregean picture that he is rejecting). In contrast, Perry eschews a content-theoretic explanation of continued belief and understands continuation of belief only in terms of belief states at different times being connected by the right kind of causal process. Despite this difference, the mental filing theorist thinks that Perry is right to appeal to the idea of a causal process in explaining continued beliefs (and continued attitudes in general). She accepts what we called Perry's *minimal positive claim* (in Section 4), but thinks that the kind of causal processes Perry posits between attitude states at earlier and later times play a *metasemantic explanatory role*: they provide a *grounding* story for *diachronic coordination*. She differs from *both* Perry and Recanati in thinking the entire story of continued attitudes can be given without representations with file-structure.

What we would like, in order to convincingly extend the *mental filing* view to explain diachronic as well as synchronic *coordination,* is a story about what *kind* of causal process counts as one of the right kind to ground diachronic coordination (and, thereby, continuation of attitudes). Imagine a collection of (synchronically) coordinated object representations, C1, at t_1, and a collection of (synchronically) coordinated object representations, C2, at t_2. Our question is: what kind of process would have to connect C2 to C1 in order to make it the case that the members of C2 are *diachronically coordinated* with the members of C1?

In brief, our answer is this: for diachronic coordination to hold between the object representations in C1 and C2, they must be connected by a process that *makes it no accident* if the information used to form the representations in C1 and the information used to form the

representations in C2 is causally sourced from the *same object* (that is, a process that typically, though not infallibly, supplies information from a single object). We'll call processes of this kind *object-directed update processes*.[20]

An *object-directed update process* takes a body of coordinated object representations at one time and generates a new body of coordinated object representations at a later time. What makes it *object-directed* is the fact that, typically, the information employed to form the earlier body of representations and the information employed to form the later body of representations is causally sourced from the *same object*.

To make this idea more concrete, we can give some examples, which illustrate both the different *varieties* of object-directed update process and also what they all have in *common*.

A first kind of object-directed update process employs *incoming information*–that is, it involves one or more ER relations. For example, C1 might be a collection of perception-based thoughts that is *governed* by a visual tracking ER relation. And C2 might be the result of adding new information supplied by that same visual tracking relation. Given the way that visual tracking works, it will be no accident if the information in C2 derives from the same object as the information in C1. So, this will count as an object-directed update process, and the representations in C1 and those in C2 will be *diachronically coordinated*.

Object-directed update processes that employ incoming information need not involve only *one* ER relation. For C1 and C2 to be diachronically coordinated, the representations in C2 need not be generated solely by an ER relation that *governed* C1. It might rather be that a new ER relation has been *added*, or has *replaced* the one that governed C1. For example, imagine that between t_1 and t_2 I learn the name of an object that I have been visually attending to, and continue to visually attend to. In this case, C2 might contain information that came in from the perceptual channel as well as information that came in through a testimonial channel. In this case, each of the ER relations is associated with an object-directed update process. And we can think of the composite process—which involves adding a new ER relation and acquiring information through each ER relation—as, itself, an object-directed update process. Why would the process that combines the two ER relations in this way count as object directed? We take it that adding a testimonial ER relation to a perceptual ER relation will involve (and *require*) sensitivity to whether the testimonial information and the perceptual information are indeed causally sourced in the same object. Therefore, it won't be an accident if the body of representations combining information from the two channels at t_2, (C2), employs information causally sourced in the same object as the body of representations that employed information from the perceptual source at t_1, (C1).

A second kind of object-directed update process involves what might be thought of as 'maintenance', but no incoming information. One case of this kind is a 'null case', in which there is no difference between what is predicated of the object represented by a coordinated body of object representations at t_1 and what is predicated of the object represented by a coordinated

[20] There is another analogy here between our mental filing view and the causal theory of memory, in that both appeal to the reliability of causal process. For example, Werning notes that if episodic memory is to sometimes justify beliefs about remembered events, the processes that form and maintain memory traces must reliably create true representations of remembered events (2020, p. 305). Note, though, that the kind of reliability we appeal to is not about the truth of any representation, but rather about a shared source of information.

body of object representations at t_2. Since no incoming information is employed in the 'null case', it is clear that the information employed to form the body of representations at t_2 will have the same causal source as the information employed to form the body of representations at t_1, and that this will be no accident. However, it's important to properly understand the sense in which no change has occurred. We do *not* presuppose, for example, that the null case involves anything static at the level of vehicles. For all we know, the subject's continuing to predicate the same properties of the same object might involve continuous *alterations* at the level of vehicles. Whatever the implementation of the maintenance process in this case, what's important is that the update process makes it no accident if the causal source of the information employed to form the earlier body of representations and the later body of representations is a *single object*.

Another species of 'maintenance' update process *does* involve a change in what is predicated of the object represented by C2 compared to what is predicated of the object represented by C1. But this change comes by way of *inference*, not incoming information. For example, imagine that C1 is a body of representations that contains a predication that *o* has a yellow bill, and a predication that *o* has a graduated tail. And imagine that C2 also contains a predication that *o* is a yellow-billed cuckoo, added on the basis of an inference from my knowledge that only yellow-billed cuckoos have both yellow bills and graduated tails. Whatever vehicular implementation this update process acts over, the point is, again, that it will count as an *object-directed* for the same reason that any object-directed update process does: it is no accident if the information employed to form the representations in C1 and C2 are causally sourced from the same object.

Above we introduced a basic idea: the representations in C2 are coordinated with those in C1 when C2 is the result of an *object-directed update process* on C1. Diachronic coordination is the reflection, at the level of content, of a rational relation between attitudes at different times grounded in a causal process that makes it no accident if the information involved in forming those attitudes is causally sourced in the same object.

Note that this approach, like the file approach, does not require that there is any overlap between C1 and C2 with respect to what is predicated of an object. Suppose I'm visually tracking a strange object as it passes in front of me. I cannot assign a sortal to it. It is continuously changing shape and colour. And I am continuously updating my beliefs about it as the visual channel delivers new information. So C1 and C2 have no predicational content in common. Despite the complete lack of overlap of predicational content, there is diachronic coordination between the representations in C1 and the representations in C2, in virtue of the fact that C2 was generated from C1 by an object-directed update process. This is just the limiting case of Perry's idea that there is *internal identity* between attitudes at different times when the right kind of causal process connects them.

So the continued filing approach achieves the characteristic result of the file approach (that diachronic coordination does not presuppose stability in predication). But framing the issue directly in terms of coordinated object representations, rather than in terms of files, brings to light a question that needs to be answered in order to give a satisfying account of diachronic coordination.

For diachronic coordination to exist between the representations in C1 and C2, does *every member* of C2 need to be the result of an object-directed update process acting on C1? Perhaps diachronic coordination does not require that much.

Take a slight modification of the visual-tracking case from above. Suppose that *one* of the predications in C2 was not generated by the visual tracking process. Perhaps, due to a cognitive glitch, a predication was added to C2 on the basis of no incoming information or inference. Other than this, however, the visual tracking process proceeds as normal, and the other predications in C2 are there because of its outputs. And, imagine, further, that both C1 and C2 contain predications to the effect that *o* has a yellow bill. This, we think, will count as a case of continued belief, from t_1 to t_2, *despite* the fact that C2 contains a predication that was not added by an object-directed update process.

The crucial thing to see is that this counts as a case of continued belief because many of the representations in C2 *have* been generated by an object-directed update process. And, this background condition also supplies the form of 'stability' against which we could see *differences* between C1 and C2 as not mere *difference,* but as *change*. The point we want to emphasise is that this required 'stability' is not itself sameness of predicational content (after all C2 could contain *none* of the same predicational content as C1 and still be diachronically coordinated with it). Nor need we understand it as sameness of a persisting mental particular. What is crucial is, rather, just that enough of the representations in C2 are generated via an object-directed update process operating on C1. The required stability for diachronic coordination—that is, the stability found in cases of both continued belief and change of mind—is stability of *process*.[21]

With this in mind, the official statement of the mental filing metasemantics for diachronic coordination should be something like:

> **GROUNDING DIACHRONIC COORDINATION**: an object representation, o_1, in a collection of synchronically coordinated representations, C1, at t_1, and an object representation, o_2, in a collection of synchronically coordinated representations, C2, at t_2, are *diachronically coordinated* only if a sufficient proportion of the members of C2 were generated via an *object-directed update process* acting on C1.[22]

This captures the kind of causal pathway between earlier attitudes and later ones that make the later attitudes rationally evaluable in relation to the earlier ones (in the way that is characteristic of coordination). We assume that if no part of some later body of attitudes is the result of an object-directed update process acting on any earlier body, then the later body is not diachronically coordinated with any earlier attitudes.

It will have been noticed that GROUNDING DIACHRONIC COORDINATION invokes the vague idea of a 'sufficient proportion' of members of a coordinated body of attitudes. One

[21] Of course, continued belief *also* requires stability of predication but that is a separate requirement that distinguishes it from change of mind.

[22] Note that we are not, here, signing on for a 'dominant causal source' theory of reference for mental states (see Evans (1973)). Firstly, we are not giving an account of reference determination but one of *coordination* between attitudes. Secondly, our requirement is not that a sufficient number of predications in C2 must have the same causal source as the predications in C1, but rather that a sufficient number of predications in C2 must be generated by the right kind of process acting on C1. And, the right kind of process is one that *makes it no accident* if the information used to form the representations in C1 and the information used to form the representations in C2 is causally sourced from the same object. This is not a *merely* causal requirement.

implication of this, which we are happy to embrace, is that there may be vagueness about whether diachronic coordination holds in a particular case (and thereby whether a subject counts as continuing to believe, desire, suppose, etc. the same thing). This strikes us as the right result. All the same, it would be nice to have a more principled way to characterize the mixture of object-directed and non-object-directed update that is consistent with diachronic coordination. We leave that project to future work.[23]

What we want to emphasise, however, that, as we see things, it is a *benefit* of the mental filing picture that it brings this question to light, rather than obscuring it by appealing to diachronic mental particulars with file-structure. For this is the sort of question one has to answer in order to give an explanatory account of diachronic coordination. Given that we do not see any explanatory payoff to positing files, we think it's better to face a question like this one head-on than to frame it as a question about the diachronic individuation of files.

Acknowledgements: The authors contributed equally to this work. They would like to thank an audience at the 'Reference and Remembering' conference at the Université Grenoble Alps, two anonymous referees, and James Openshaw.

References

Almotahari, Mahrad, & Gray, Aidan. 2020. Frege Cases and Bad Psychological Laws. Mind. https://doi.org/10.1093/mind/fzaa078.

Campbell, J. 1987. Is Sense Transparent? *Proceedings of the Aristotelian Society*, 88, 273-292.

Clarke, H. 2022. Mental Filing Systems: A User's Guide. *Ergo*, *8*.

Devitt, Michael. 1989. Against direct reference. Midwest studies in philosophy, 14(1), 206–240.

Dickie, I. 2015. *Fixing Reference*. Oxford University Press.

Evans, G. 1973. 'The Causal Theory of Names', *Aristotelian Society Supplementary Volume,* 47, 187-208.

Evans, G. 1982. *The Varieties of Reference*, *edited by J. McDowell*. Oxford University Press.

Fine, K. 2007. *Semantic Relationism*. Wiley-Blackwell.

Fodor, J. A. 1995. *The Elm and the Expert: Mentalese and its Semantics*. MIT Press.

Goodman, R. 2024. Mental Files. *Philosophy Compass*.

Goodman, R. & Gray, A. 2022. Mental Filing. Noûs. 54(1), 204-226.

Gray, A. 2017. Relational Approaches to Frege's Puzzle. *Philosophy Compass*, 12(10):e12429.

Heck, Richard Kimberly. 2002. Do demonstratives have senses? Philosophers' Imprint, 2(2), 1–33. (originally published under the name "Richard G. Heck, Jr").

Heck, R. K.. 2012. Solving Frege's Puzzle. *Journal of Philosophy*, 109, 132–174. (Originally

[23] A feature of this question that we have not addressed is the role of meta-cognition in establishing diachronic coordination. In sophisticated cases, the subject at the later time will have beliefs about whether some of their current attitudes are about an object which they previously had attitudes about. And it might be that the presence or absence of these beliefs is, itself, part of the grounding story for coordination. Though we don't want to deny this line of thought, it strikes us as theoretically useful to consider the 'simple case' first to understand what the basic phenomena looks like, before adding in the sophistication associated with meta-cognition. And we would certainly not want to hold that such second-order beliefs are infallible.

published under the name “Richard G. Heck, Jr.”).
Heck, Richard Kimberly. 2014. Intuition and the substitution argument. Analytic Philosophy, 55(1). (originally published under the name "Richard G. Heck, Jr").
Heim, I.R. 1982. The Semantics of Definite and Indefinite Noun Phrases. U. Mass Amherst dissertation.
Kahneman D. and Treisman, A. 1984. ‘Changing Views of Attention and Automaticity”, in R. Parasuraman and D.R. Davis (Eds.), Varieties of Attention, Academic Press, 29-61.
Kahneman, D., Treisman, A., and Gibbs, B.J. 1992. ‘The Reviewing of Object Files: Object-specific Integration of Information’, Cognitive Psychology, 24:2, 175-219.
Kamp, H. 1981. ‘A Theory of Truth and Semantic Representation’, in J. Groenendijk (Ed.), Formal Methods in the Study of Language, University of Amsterdam Press, 277-322.
Lawlor, K. 2001. *New Thoughts about Old Things: Cognitive Policies as the Ground of Singular Concepts*. Taylor & Francis.
Loar, Brian. 1988. Social Content and Psychological Content. In: Grimm, Robert H., & Merrill, D. D. (eds), Contents of Thought. University of Arizona Press.
Maier, E. 2016. ‘Attitudes and Mental Files in Discourse Representation Theory’, Review of Philosophy and Psychology, 7, 473-90.
Martin, C.B. and Deutscher, M., 1966. Remembering. *The Philosophical Review*, *75*(2), pp.161-196.
Murez, M., Smortchkova, J. & Strickland, B. 2020. The mental files theory of singular thought: a psychological perspective, in *Mental Files and Singular Thought,* edited by R. Goodman, J.Genone and N. Kroll. Oxford University Press, pp. 107-142.
Perry, J. 1980/1993. A problem about continued belief*,* in *The problem of the essential indexical: And other essays.* Oxford University Press, USA. pp. 69-90
Perry, J., 1997. Rip van Winkle and other characters. *European Review of Philosophy*, *2*, pp.13-40.
Pinillos, N. Á. 2011. Coreference and Meaning. *Philosophical Studies*, 154(2), 301–324.
Prosser, S. 2019. The Metaphysics of Mental Files. *Philosophy and Phenomenological Research*. doi:10.1111/phpr.12577
Pryor, J. 2016. Mental Graphs. *Review of Philosophy and Psychology*, 7(2), 309–341.
Recanati, F. 2012. *Mental Files*. Oxford University Press.
Recanati, F. 2015 ‘Replies’, Inquiry, 58:4, 408-437.
Recanati, F. 2016. *Files in Flux*. Oxford University Press.
Recanati, F. 2021 ‘Do Mental Files Obey Strawson’s Constraint?’, in C. Borgoni, D. Kindermann, and A. Onofri (Eds.) The Fragmented Mind, Oxford University Press.
Robins, S., 2016. Representing the past: Memory traces and the causal theory of memory. *Philosophical Studies*, *173*, pp.2993-3013.
Scholl, B.J. 2001. ‘Objects and Attention: The State of the Art, Cognition, 80:1-2, 1-46.
Tałasiewicz, M., 2022. Files are mental particulars: a rejoinder to the relationist challenge. *Synthese*, *200*(6), p.450.
Taschek, W. 1995. Belief, Substitution, and Logical Structure. *Noûs*, 29(1), 71–95.
Werning, M. 2020. Predicting the past from minimal traces: Episodic memory and its distinction from imagination and preservation. *Review of philosophy and psychology*, *11*(2), 301-333.