



# Files are mental particulars: a rejoinder to the relationist challenge

Mieszko Tałasiewicz<sup>1</sup> 

Received: 21 March 2022 / Accepted: 12 October 2022 / Published online: 30 October 2022

© The Author(s) 2022

## Abstract

The paper is a rejoinder to a challenge against the particularist version of the mental files framework (MFF) posed by the relationist approach based on the notion of content coordination [such as recent attempt by Rachel Goodman and Aidan Gray in (Noûs <https://doi.org/10.1111/nous.12354> (2020))]. Relationists argue that important explanatory goals of MFF: (1) could be achieved without positing files as mental particulars, as there is a relationist notion of *content coordination* at hand that can be aptly used for “filing without files”; and (2) should be so achieved, as there are difficulties that afflict the particularist approach to MFF and the relationist account is simply better. However, both claims should be rejected. The particularist approach to MFF, properly interpreted, would not get into the troubles it is accused of generating. Indeed, it is the relationist approach that gets in trouble. Specifically, it lacks resources for explanation of nuances, which can be easily accounted for in terms of particularist interpretation, and, furthermore, it lowers the interdisciplinary standing of the whole framework. The particularist version is therefore better.

**Keywords** Mental files · Relationism · Coordination · François Recanati · Mental representation · Mental particulars

## 1 Introduction

The theoretical framework of mental files (MFF) was conceived for the task of conceptualizing a broad array of ideas within the philosophy of language, including syntax, semantics and pragmatics. Stemming from the works of Paul H. Grice, Michael Lockwood, Peter F. Strawson and John Perry, MFF was developed and extensively

---

✉ Mieszko Tałasiewicz  
m.talasiewicz@uw.edu.pl

<sup>1</sup> Faculty of Philosophy, University of Warsaw, Warsaw, Poland

elaborated by François Recanati in a series of publications, most notably in an influential monograph, *Mental Files*, and subsequent books and papers.<sup>1</sup> According to the classic formulations of this framework, mental files are posited as mental *particulars*, real cognitive entities. Alternatively, some theorists attempt to frame the idea of mental files in terms of content coordination: *relational*, non-objectual interdependency between some representational chunks of information. A recent example of the latter approach has been given by Rachel Goodman and Aidan Gray (Goodman & Gray, 2020).

The aim of my paper is to argue that the particularist interpretation of MFF withstands criticisms and has greater explanatory power than its relationist alternative. I intend to proceed along four lines (outlined in Sects. 2–5). First, I will argue that currently debated “worries” about the notion of containment of information in files-as-particulars are not really serious, and I will propose an interpretation of MFF that avoids them. Second, I will argue that the charge of the representational insignificance of files as particulars is not accurate. Third, I will show that there are serious limitations to the relational approach to MFF, which substantially reduce its explanatory power, compared to the particularist version. Fourth, I will advocate a methodological view: that even if relational accounts could be considered *legitimate* alternatives to particularist ones in a certain domain—which I consider a counterfactual assumption—they would not be *preferable*, *ceteris paribus*.

## 2 There is nothing worrying about containment

Among the objections raised against the particularist version of MFF is the so-called “containment puzzle.” Specifically, it is believed that the particularist version of MFF involves a potentially troubling form of reciprocal containment of beliefs in files and files in beliefs. As Goodman and Gray would put it, “there is a sense in which a file about *a* contains beliefs about *a* [and] there is also a sense in which my beliefs about *a* contain my *a*-file” (Goodman & Gray, 2020, p. 18). Such a worry, if sound, would be serious, for while containment is just a metaphor within MFF, it plays important clarifying role in this framework. A collapse of this metaphor might create a threat of circularity in explanations within MFF or lead to ultimately unsatisfying complications of MFF.

This worry is not sound, though. There is a straightforward answer available to particularists about what is contained in what. The sense according to which beliefs are contained in files, and the sense according to which files are contained in beliefs, while both metaphorical, are substantially different senses, and thus cannot contribute to any sort of circularity.

To see that this answer is plausible, we need to acknowledge some rudimentary insights regarding syntax that exist within MFF. Struck by MFF’s explanatory power in semantics, we sometimes lose sight of the fact that MFF originated as a part of *syntactic* theory. In particular, it featured in Strawson’s attempts to elucidate the nature of the subject-predicate distinction. Name-hood, for Strawson, is not a semantic but

---

<sup>1</sup> Recanati (2012, 2013), Murez and Recanati (2016), and Recanati (2016, 2021a, b).

a syntactic notion, roughly equal to being the subject of a sentence. The syntactic role of a name, the role of *referring*, can be performed by many kinds of expressions: proper names, descriptive names, general names, regular descriptions, pronouns, etc. No particular semantic scheme distinguishes a name (cf. Strawson, 1974, chp. 2). Thus, a problem arises: how this syntactic notion of referring can be disentangled from various semantic mechanisms of reference, characteristic of particular types of names. A germinal mental file picture was devised to (among other purposes) provide an answer to this problem: names, of whatever kind, are connected to certain mental representations of objects—mental files—and this is what is meant by their playing the “referring” role. How exactly particular representations of different types capture their real-world designates is a different question (a semantic one). Syntax, in this account, mirrors the cognitive architecture of thought. It is supposed to “reflect some fundamental features of our thought about the world” (Strawson, 1974, p. 11).

I am not going to discuss Strawsonian grammar in full detail here, for one need not rely on any particularly Strawsonian conception: any syntactic conception will do, as long as it keeps the relations of constituency between sentences and their parts clear. The point of mentioning Strawson is to show that syntactic insights are not “epicycles” to the mental files framework. On the contrary, they belong to the very core of it, and the MFF theorists are fully entitled to make use of them in response to challenges. Let us make such use, then.

Propositions—and propositional attitudes, of which propositions are contents—syntactically correspond to sentences. In the simple case of a subject-predicate sentence, the subject (name) corresponds to a *whole* file. The *content* of the file, in turn, corresponds to the (set of available) predicate(s). Thus, the content of a file is literally neither a nominal object representation (for it is the file as a whole what counts as an object representation) nor a propositional belief. The content of a file is *unsaturated, predicative* “information about  $\xi$ ,” where  $\xi$  is not some indefinite reference but a definite blank slot, ready to be filled, in the course of further cognitive processes, by the referent of the file.<sup>2</sup> Processes of belief-forming are precisely such further higher-level processes, in which some information taken from the content of a file is predicated about the referent of the file. As Récanati would put it: “a file is deployed in thought only if it has a sufficient degree of activation (whatever that amounts to in neural terms)” (Récanati, 2016, p. vii). Mere containment of a predicate in a file would not count as *activation* of the file and thus it could not be squared with entertaining *occurrent* beliefs.

However, it is *standing* beliefs of a person *s* about object *x* that might be plausibly taken to characterize the set of predicates about *x* available to *s*. But standing beliefs cannot be identified with the contents of files, either. At best, they might be identified with *files having their contents*. It is plausible to stipulate the following equivalence:

<sup>2</sup> “Files only contain *predicative* elements [...] the reference of the file is, to a very large extent, independent of the information it contains” (Récanati, 2016, p. viii). For a discussion of the unsaturatedness of predicates, see e.g. (Dummett, 1995, 1981; Frege, 1892; Geach, 1975, 1980; Hale, 1979; Husserl, 2001; Tałasiewicz, 2017b). For the present purpose, I keep Frege’s notation, in which the slots in the predicates are represented by Greek letters “ $\xi$ ” and “ $\zeta$ ” (in contrast to nominal variables “*x*”, “*y*”, “*z*”, which represent objects that are supposed to fill the slots in the process of sentence-forming: on Strawson’s account nominal variables are syntactically names). Alternatively, lambda-notation is always possible. In this notation, a one-place predicate is represented by a formula  $\lambda x.Gx$ .

(Standing Belief) A person  $s$  has a standing belief  $F(x)$  if and only if  $s$  has a file for  $x$  and the file contains the predicate  $\lambda x.Fx$ .

It does not follow, however, that the whole belief  $F(x)$  is contained in the file for  $x$ . Consider the predicate “ $\xi$  is a philosopher.” Suppose I have two beliefs: that Plato is a philosopher and that Aristotle is a philosopher. I have two distinct files, for Plato and for Aristotle, respectively. These files, while distinct, contain the same predicate, “ $\xi$  is a philosopher.” However, they do not contain the same belief.

Despite Goodman’s and Gray’s worry, then, there is no threat of circular containment: while files are contained (deployed) in *occurrent* beliefs, they contain predicates that reflect the subject’s *standing* beliefs about the referent of the file. However, even in this latter case we might say that beliefs are contained in files only figuratively, if we identify predicates with standing beliefs in a *pars pro toto* manner, abstracting from the syntactic difference between a sentence and an unsaturated predicate contained in this sentence.

Another concern that might be raised with regard to the particularist notion of containment—apart from its alleged circularity, discussed above—is connected to the idea that co-containment of predicates in the same file is supposed to account for the validity of inferences called “trading on identity.” Typically, such inferences are understood as inferences from  $Fa \ \& \ Ga$  to  $\exists x (Fx \ \& \ Gx)$ , and the problem is that their validity depends on the co-referentiality of both occurrences of the term “ $a$ ” being somehow “encoded” in the premises (while it cannot be expressed explicitly as an additional premise under the threat of an infinite regress). Within MFF, it is precisely the co-containment of predicates that encodes their coreference in such inferences.<sup>3</sup> However, it has been suggested that “trading on identity” should be understood more broadly, so that this notion could be applied to transitions between different attitudes toward the same object; for instance, to a transition from a *belief* that Hesperus is visible and a *desire* to see Hesperus to an *intention* to look at Hesperus. Arguably, trading on identity in this broader sense would not be able to be accounted for in terms of co-containment of predicates in files (Goodman & Gray, 2020, pp. 3–5).

Quite naturally, though, this extended notion of trading on identity cannot be framed in terms of *co-containment* of predicates in the same file, because the reasoning here has nothing to do with any predicates. The particularists have different resources to account for this extension of trading on identity, namely the notion of *co-deployment* of the same file in different attitudes. Let us attempt the following characterization:

(Trading Across Attitudes) If a person  $S$  has an attitude  $A$  with the referential content  $\varphi(a)$  and an attitude  $B$  with the referential content  $\psi(a)$ ,  $S$  can trade on the identity of  $a$  in those attitudes if and only if there is a mental file  $F$ , such that the attitude  $A$  has the referential content  $\varphi(a)$  in virtue of a deployment of the file  $F$  and the attitude  $B$  has the referential content  $\psi(a)$  in virtue of a deployment of the file  $F$ .

<sup>3</sup> For a fuller discussion of this problem on the grounds of MFF, see (Recanati, 2012, pp. 47–50).

This characterization arguably only covers the cases in which occurrent attitudes are considered (for the notion of file-deployment applies to active—occurrent—states).<sup>4</sup> However, this seems sufficient to elucidate the extended notion of trading on identity. Trading is a sort of *inference*, not just an entailment, and inferences—in contrast to entailments—are mental activities in which occurrent attitudes are involved.

There are problems in the vicinity, though, which prompt closer examination of standing attitudes.<sup>5</sup> One of such problems is a problem raised against MFF in (Ninan, 2015): how to account for non-doxastic attitudes (such as desires or imaginings) in which some predicate is counterfactually produced about the referent of a given attitude. Suppose we believe that Napoleon was dictatorial and wish that he had not been. According to the formula (Standing Belief) above, our standing belief that Napoleon was dictatorial can be squared with our having a Napoleon-file and the file containing the predicate “ $\xi$  is dictatorial.” It is not that easy in the case of the standing desire that Napoleon hadn’t been dictatorial.

However, there are multiple solutions already available. One has been proposed by Ninan himself. According to this solution, separate sets of files in our minds come along with different attitudes: we have one Napoleon file for our beliefs about Napoleon and another Napoleon file for our desires about Napoleon (and the latter contains the predicate “ $\xi$  is not dictatorial”). Another solution, inspired by (Forbes, 1990), has been proposed by Recanati in his response to Ninan (Recanati, 2015). According to this solution, separate sets of predicates in our files come along with different attitudes: files store not just predicates, but rather classified predicates, or pairs: attitude-predicate. We have just one Napoleon-file, but this file contains both  $\langle \text{Bel}, \xi \text{ is dictatorial} \rangle$  and  $\langle \text{Des}, \xi \text{ is not dictatorial} \rangle$ .

Yet another solution, formulated in terms of discourse representation theory (although with the explicit aim of contributing to the MFF), has been suggested in (Maier, 2016). According to this solution, we might say that standing attitudes are higher level *structures* in which files are contained, just as occurrent attitudes are higher level *processes* in which files are deployed. Files-constituents of a standing attitude serve the purpose of identifying the referent of the attitude. What is predicated about this referent within this attitude is specified independently, in subsequent constituents of the attitude, and has nothing to do with the content of the file. Such structure of an attitude reflects the idea of singular thinking: the file is a mode of presentation of the referent of the attitude, but the content of the file does not affect the conditions of satisfaction of singular thoughts about this referent. Specifically, such attitude as *a wish that Napoleon had not been dictatorial* can be represented as the following structure:  $\langle \text{Napoleon-file } (x), \text{Des } (\text{not-dictatorial } (x)) \rangle$ .<sup>6</sup> The fact that the Napoleon-file contains the predicate “ $\xi$  is dictatorial” does not affect the conditions of satisfaction of the desire. The file provides only referent for the desire, not the content.

<sup>4</sup> For recent arguments that occurrent states are active states, see (Bartlett, 2018).

<sup>5</sup> Thanks to the anonymous referee for bringing these problems to my attention.

<sup>6</sup> Compare Maier (2016, p. 489).

Summing up, MFF seems to have resources to get through the problem of trading on identity across attitudes as well as the problem of different attitudes involving incoherent predicates.<sup>7</sup>

Yet another worry about containment is that it is unclear how a file theorist should understand a belief that  $a$  stands in relation  $R$  to  $b$ . “Is it a presence of the predicate ‘ $x$  stands in  $R$  to  $b$ ’ in the  $a$ -file? Or the predicate ‘ $a$  stands in  $R$  to  $x$ ’ in the  $b$ -file? Or both? If both, are these two independent states?” (Goodman & Gray, 2020, p. 4). However, there is nothing worrying in having the predicate “ $\xi$  is in  $R$  to  $b$ ” stored in the  $a$ -file and “ $a$  is in  $R$  to  $\xi$ ” in the  $b$ -file simultaneously. For, as highlighted above, predicates are not beliefs. The belief that  $a$  stands in relation  $R$  to  $b$  is a higher-level state that deploys both files but is not contained in either of them. Besides, as MFF is supposed to model the informational resources of a subject, it also allows for having just one of the predicates expressing relative properties stored in a file depending on informational needs, which might be very asymmetrical.<sup>8</sup>

<sup>7</sup> In case someone insists that trading on identity across attitudes somehow involves standing attitudes, despite appearances, the solution to the former problem will certainly depend on the solution of the latter—but in a quite straightforward manner. For example, if one takes Recanati’s solution to the problem of standing attitudes with incoherent predication (according to which files store classified predicates), she might choose the following modification of the (Trading Across Attitudes) formula:

If a person  $S$  has an attitude  $A$  with the referential content  $\varphi(a)$  and an attitude  $B$  with the referential content  $\psi(a)$ ,  $S$  can trade on the identity of  $a$  in those attitudes if and only if there is a mental file  $F$  such that the attitude  $A$  has the referential content  $\varphi(a)$  in virtue of the file  $F$  containing the classified predicate  $\langle A, \varphi \rangle$  and the attitude  $B$  has the referential content  $\psi(a)$  in virtue of the file  $F$  containing the classified predicate  $\langle B, \psi \rangle$ .

<sup>8</sup> Suppose it is important information about me that I am indebted intellectually to Aristotle. Certainly, it is not an important thing to know about Aristotle that I owe him my intellectual legacy. The predicate “ $\xi$  is indebted to Aristotle” might thus be an important part of my SELF file, while the predicate “Mieszko is indebted to  $\xi$ ” is not an important part of my ARISTOTLE file (and plausibly it is not a part of it at all). Again, there is nothing awkward or puzzling in this. Along with the problem of storing relational information in files, it is worthwhile to consider the question of whether, in certain circumstances, some instantaneous beliefs in the form “ $a$  stands in relation  $R$  to  $b$ ” can occur without any deployment of files (or other representations of objects) at all. Some authors contend that some beliefs can be formed with the use of information that comes directly from some non-representational cognitive resources. John Searle argues, for instance, that we need a background of “nonrepresentational mental capacities that enable all representing to take place” (Searle, 1983, p. 143). Among such capacities Timothy Williamson counts, for example, “various propensities to form expectations about what happens next: for example, to project the trajectories of nearby moving bodies into the immediate future (otherwise we could not catch balls)” (Williamson, 2007, pp. 148–49). Presumably, a capacity to instantly estimate spatial relations between simultaneously perceived objects could be counted here as well. Such a capacity would, in turn, open a possibility that a belief like “ $a$  is now closer to me than  $b$ ” can be formed in one’s mind without any previous information about the relative location of the objects  $a$  and  $b$  in the respective files. In such a case, the predicates “ $\xi$  is in  $R$  to  $b$ ” and “ $a$  is in  $R$  to  $\xi$ ” can be written in the respective files for future use only as a *result* of this particular instance of belief-forming, prompted by immediate visual cognition (or, if not important enough, they can be dropped altogether, leaving the  $a$ -file and  $b$ -file relieved from the burden of keeping unwanted information). In such cases, instantaneous beliefs would serve as triggering mechanisms for opening files, without prior deploying any file—except, perhaps, the SELF file. Such cases would go beyond a standard range of MFF, yet not necessarily against the gist of this framework. Be it as it may, this remark has no bearing on the question raised by Goodman and Gray about *standing* relational beliefs, which certainly must involve files (on the grounds of MFF).

To sum up the arguments in Sect. 2, containment is still a metaphor, of course, for files are not literal containers and predicates are not the stuff within them. But it is a metaphor that keeps clear what is contained where and gives a correct general picture of what is going on in predication. Containment worries should not affect any self-confident version of the particularist MFF.

### 3 *Haeceitas* is too much to ask for

Another argument against the particularist version of MFF says that containment of information in a particular file is not representationally significant (and thus that files as entities are theoretically dispensable in favor of relational notions such as coordination of units of information). As Goodman and Gray put it: “if file-containment is a representational feature, permutations that preserve the structure of co-filing but change which predicates are stored in which files could alter the representational features of an attitude state. But the file-framework doesn’t appear to make use of meaningful permutations of this kind” (Goodman & Gray, 2020, p. 12).

I will discuss two interpretations of this argument, and I argue that neither hits its target. One of them imposes on MFF a requirement that is easily met, and the other a requirement that is impossible to meet for many theoretical frameworks in science (and thus irrational to impose on any).

According to the first interpretation, the requirement is straightforward: MFF should allow for a permutation that preserves the structure of co-filing, but changes which predicates are stored in which files, and thus alters the representational features of an attitude state. MFF meets such a requirement easily. It is enough that a certain cluster of predicates would get into a file that had a different type (say: encyclopedic, or visual, or SELF), or structure (say: relational file, or satisfactorial file),<sup>9</sup> or belonged to a different linking network. Since the types, structure and linking network of the files are representationally significant, a permutation that changed the type, structure or linking network of a file, while preserving the coordination of predicates within, would *ipso facto* change the representational features of the overall cognitive state. For instance, transferring the content from a singular-relational file to a descriptive-satisfactorial file would obviously change the representational features of the cognitive state, as it would change the truth conditions of the thoughts deploying the respective files. A permutation that switched the content between files of different linking networks would change of representational features of the cognitive state, too. Such linking might license a rational presumption of coreference (given the trustworthiness of the link), so that such a permutation, which cut the links, would directly affect inferential resources of the cognitive state (see the next section for details). According to this interpretation of the requirement, files as particulars simply are representationally significant.

According to the second interpretation, the requirement might be read as suggesting that any change in the representational features of cognitive states connected to any

---

<sup>9</sup> Differentiation of files into singular (relational) and descriptive (satisfactorial) is advocated in (Goodman, 2016) and (Tałasiewicz, 2017a), and acknowledged in (Recanati, 2016, 2021b).

difference in files depending on any sort of specifiable general characteristics whatsoever would not count as a significant permutation. In this view, such differences are not rooted in the files' true identities, but merely in their satisfying certain conditions. Thus, they are accountable, theoretically, in purely relational terms (say, by specifying certain conditions on the relation of coordination), without resorting to files as particulars. According to this interpretation, the only case that would count in favor of mental files as particulars would be one in which representational differences in an attitude state were induced by a permutation of content between *qualitatively identical* files, sharing not only the structure of co-filing, but literally all conceivable characteristics, differing *solo numero*.<sup>10</sup>

But this is too much to ask for. Presumably, there is no way in which the *haecceitas* of a file could play an explanatory role in semantics (arguably, for instance, differences *solo numero* would not suffice for generating Frege cases). However, this is true of many theoretical objects of many sciences, and it would not normally prevent us from admitting the existence and entitativity of these objects. For example, the *haecceitas* of electrons does not play any role in any sort of physical explanation; yet we are not inclined, because of this fact, to consider electrons as mere spatio-temporal coordinations between some mass and electric charge.

It may well be, then, that files can be individuated merely as “pegs” for certain characteristics—let us call these characteristics “metadata.”<sup>11</sup> Type, structure and opening circumstances (such as time of opening as well as time of subsequent modifications, linking/merging history, type of information channel governing the file, etc.) are some of the most important and salient ones.<sup>12</sup> Mental-file theorists, when speaking of mental files as mental particulars, do not claim that files have their *haecceitas* and that we can intelligibly talk about a permutation that switches two files which are identical under any aspect whatsoever, but still “numerically” distinct, and furthermore that such a permutation would make any representational difference (at least, I am not aware of anyone espousing such a radical form of particularism). What they say is that two identical clusters of predicates may be stored in two distinct files, which are thus identical *in the aspect of their informational content*, but which differ in their metadata, such as their type, structure, linking network, opening circumstances, activation history, etc. These non-informational differences might make a representational difference or a semantic difference (including say, a difference between singular and descriptive truth-conditions).

An excellent illustration of this point is Peter Pagin's “moth example”:

<sup>10</sup> There are reasons to think that Goodman and Gray have this second interpretation in mind in their challenge against mental files as particulars. They suggest that the mere type change of a file would be irrelevant for establishing the representational significance of file containment, and explicitly talk about a “purely haecceitic switch” as the intended requirement (Goodman & Gray, 2020, p. 12).

<sup>11</sup> For a fuller account of files' metadata see (Talasiewicz, 2017a).

<sup>12</sup> Files, like all representational objects that have some content, require observing the difference between properties that are “contained” in the content of a file, and properties of the file as an individual (these are what I call “metadata”). For instance, my ARISTOTLE file contains predicates, say, “ξ is a philosopher” and “ξ is an author of *Nicomachean Ethics*,” and possesses properties, say, of being opened in 1986 in the course of syntactic processing of a certain sentence in a certain copy of a certain textbook, and of being activated recently in the course of writing this paper.



It seems to me quite possible that a subject can have two files that do *not* differ in their acquaintance relations [ER relations]. Suppose X takes herself to see two moths flying around in her kitchen. She opens a file for each, alpha and beta, thinking of them as “A” and “B”, respectively. She takes herself to see now A, now B. The acquaintance relations are indeed different in case there *are* two moths, one causing the opening of alpha, the other the opening of beta. But in case the subject in fact is mistaken, and there is only *one* moth causing the opening of both files, there does not seem to be any difference between the acquaintance relations of alpha and beta. X opens first demonstrative files, which are converted into memory files, when a moth is taken to go out of sight, and then converted to recognitional files, when a moth is taken to be seen again. There is, we may assume, no particular feature in the external aspects of the causal relations that explains the difference. X takes herself sometimes to see A, sometimes to see B, and most often not to know whether she is seeing A or B (Pagin, 2013, p. 140).

This example shows something *prima facie* problematic: that we can have two representations governed by the same ER relation, and yet we cannot trade on the identity of the object thereby represented (for we take ourselves to believe that there are two different objects).

When mental files are posited as particulars, the problem is easily solved. The channels of information connecting objects to their respective files are constituted not merely by “the external aspects of the causal relations” but also by the internal properties of the files themselves, such as the circumstances of their opening. Let us stress this again: it is the properties *of* the files (their metadata), not the predicates *contained* in the files, that matter here. In the moth example, Mrs. X clearly opens two distinct files in two distinct contexts; whether the objectual source of these files is one moth or two moths is another question.<sup>13</sup>

A relationist’s response would require an elaboration, without positing files as mental particulars, of a difference between two distinct instances of some ER relation in a situation in which there is no feature of the external aspects of this situation that explains the difference.

However, relationists limit their accounts to mere *content* coordination. For instance, Goodman and Gray characterize coordination by a connection between the functional role of trading on identity and the functional role contributed by epistemically rewarding (ER) relations. As they say, “[t]he idea is simple: representations *a* and *b* are governed by the same ER relation if and only if any object representations that are generated by that ER relation will stand in the trading on identity functional role to both *a* and *b*” (Goodman & Gray, 2020, p. 15). But, as “moth example” shows, such equivalence doesn’t hold in general.

The particularists, then, need not endorse a radical claim that differences in files *solo numero* are representationally significant; they might take a moderate claim that differences in files’ metadata are representationally significant. What relationists would have to do in order to show the dispensability of files as particulars is to characterize a total coordination of predicates *within* a file *and* of properties *of* the file (the

<sup>13</sup> For further discussion, see Recanati (2021a) and Talasiewicz (2017a).

file's *metadata*). Only such theory of coordination could successfully replace particularist interpretations of MFF, for the properties of the files are as representationally significant as properties contained in the files.

#### 4 Content coordination is not enough

As we have seen in the previous section, content coordination is not enough for explaining such cases as Pagin's moth example. Moreover, there are other concerns about the notion of content coordination in connection with "trading on identity."

One thing is the synchronicity of coordination in the relationist account, notably in Goodman and Gray's proposal. They justify this limitation by saying that "co-filedness of information is meant to explain the rational permissibility of inferences that trade on identity—and it is synchronic co-filing that explains this" (Goodman & Gray, 2020, p. 8). But there are more things the files are meant to explain than merely the rational permissibility of inferences that trade on identity.

In the core of MFF there is a distinction of regular files and indexed files, the latter standing "in the subject's mind, for another subject's file about an object" (Recanati, 2012, p. 183). In particular, a subject might possess a temporally indexed file for an object, reflecting the subject's own previous mode of presentation of this object. In recent developments of MFF (Recanati, 2016, 2021a), an orthogonal distinction is introduced: of dynamic files and static files, which allows for an even more fine-grained picture. We might say that regular files are dynamic representations, which, in the process of accommodation of some new information, undergo certain changes, such as simple conversion, incremental conversion, fusion, fission, etc. Temporally indexed files, on the other hand, are static files: *stages* of the former ones. Different static stages of a given dynamic file are not related to each other through strict identity but rather through dynamic continuity (which is an intransitive relation).<sup>14</sup>

Accordingly, two notions of coreference *de jure* (CDJ) are introduced: strong CDJ and weak CDJ. A subject in the strong CDJ state with respect to M and N knows that  $\forall x(\text{Ref}(N, x) \equiv \text{Ref}(M, x))$ .<sup>15</sup> A subject in the weak CDJ state with respect to M and N knows that  $\forall x\forall y((\text{Ref}(N, x) \ \& \ \text{Ref}(M, y)) \rightarrow x = y)$ . Strong CDJ rules out all cases of referential divergence between M and N, while weak CDJ allows for cases in which one of the utterances/deployments M or N refers and the other one fails to refer (Recanati, 2016, pp. 26–27).

Simultaneous deployments of the same file are coreferential *de jure* in the strong sense. Only such coreference licenses trading on identity. However, strong CDJ is only part of the story, for in the case of two non-simultaneous deployments of the

<sup>14</sup> "I take modes of presentation to be static files or file-stages, and I take such files to undergo dynamic operations such as conversion, incremental conversion, fusion (file merging), fission (file splitting), and so on. Sequences of file related by such operations are dynamic files. These operations do not preserve file identity in the strict, Leibnizian sense: in contrast to identity, dynamic continuity between files is an intransitive relation, just like dynamic continuity between person stages" (Recanati, 2016, p. 84).

<sup>15</sup> Typically, "M" and "N" are different utterances of referential terms. However, within MFF, such linguistic entourage is not required: "M" and "N" might be considered more generally as different deployments of mental files.

same regular file (or simultaneous deployments of different temporally indexed files considered as static stages of the same regular file), only weak CDJ is involved.<sup>16</sup>

Recanati gives an example of such deployment after Pinillos. He considers the following utterance:

(HPT) We were debating whether to investigate both *Hesperus* and *Phosphorus*; but when we got evidence of their true identity, we immediately sent probes *there*.

with the following commentary:

[The] example involves three coreferential files: two indexed files (the HESPERUS file and the PHOSPHORUS file, indexed to the ascriber) and a regular file (the inclusive VENUS file, corresponding to the speaker's current point of view) [...]. The problem is that, even though the terms 'Hesperus' and 'Phosphorus' are only coreferential *de facto* [...], each of them is coreferential *de jure* with the inclusive term 'there' in the second clause (the term associated with the inclusive file): the speech protagonists know that either 'there' fails to refer to a unique location (if the identity Hesperus = Phosphorus is not true), or (if the identity is true) it refers to the location of the single planet which Hesperus and Phosphorus turn out to be. That piece of knowledge corresponds to weak CDJ: for each of the two terms 'Hesperus' and 'Phosphorus', the subject knows that *that term corefers with the inclusive term 'there' if they both refer*. In other words: 'Hesperus' is in the (weak) CDJ relation to 'there', 'Phosphorus' is in the (weak) CDJ relation to 'there', yet 'Hesperus' and 'Phosphorus' do not stand in the weak CDJ relation to each other: they are not coreferential *de jure*, but *de facto* (Recanati, 2016, pp. 37–38).

This example shows that diachronic individuation of files and corresponding weak CDJ, while insufficient for licensing trading on identity, is still relevant to the explanation of certain aspects of coordination between files. In (Recanati, 2012), such coordination between indexed files was dubbed "vertical linking". While recognition cases indeed can be construed as synchronic, thanks to the notion of incremental conversion of persisting files,<sup>17</sup> there are cases—such as HPT—in which a rational subject needs to deploy simultaneously her persisting regular file and her time-indexed stage-file representing her previous state of knowledge.<sup>18</sup>

<sup>16</sup> "Two deployments of the same file at different times are not coreferential *de jure* in the strong sense; they are only bound to refer to the same thing if they both refer (weak CDJ)" (Recanati, 2016, p. 88).

<sup>17</sup> Such is the case of a yellow-billed cuckoo, discussed in Goodman and Gray (2020, pp. 6–7).

<sup>18</sup> A similar account has recently been proposed by Simon Prosser, who argued for an analogy in metaphysical status of mental files and persons (Prosser, 2020). There is one important difference though between Recanati's view and Prosser's: Recanati insists on keeping static stage-files alongside dynamic persisting files, while Prosser is skeptical about their utility. In (Prosser, 2020, pp. 662–64) several cases are shown to be accountable without resorting to stage-files, indeed, and it is convincingly argued that the dynamic files do not call static ones for explanation. Thus perhaps Recanati's statement that "we need the static files, in particular, to make sense of dynamic phenomena like fusion and fission of files" (Recanati, 2016, p. 83) is too hasty. However, to my mind, examples such as HPT—Prosser did not discuss any of this kind—do reveal the need for the static files. The key to such cases is simultaneous deployment of a dynamic and a static file, indexed by time. Besides, this sets some limitations to the analogy of mental files and persons, advocated by Prosser. Suppose, for instance, that when I am no longer capable of climbing tall fences, I call

Weak CDJ is a case of coordination between files that is not based upon and does not license trading on identity. We may think of even looser kinds of coordination between files, which could be dubbed “putative linking.” This kind of linking—we are going here beyond the framework proposed in (Recanati, 2016) or (Recanati, 2021b), but not against the major tenets expressed there, I presume—emerges between obviously different files when they get some salient characteristics in common (like circumstances of opening). Suppose that I hear a whiz and simultaneously get hit in my head. I don’t know whether the same thing that produced this whiz hit me; in fact, I am curious whether or not it is so. I open two distinct mental files for the object that made the whiz and for the object that hit my head, and I think of their putative coreference. Such coreference is neither strong CDJ nor weak CDJ; it is no coreference *de jure* at all. But it is not *merely de facto*, either.

By *merely de facto* coreference, I mean a coreference of two files that have nothing in common and just happen to co-refer. Suppose I have a neighbor to whom I say good-morning every day on the train, but I have no other relationship with him. I do not know his name, and I am not even sure I could recognize him in different circumstances. Certainly, I have a mental file for him, albeit a rather meager one. One day, I learn from newspapers that the new President of the Academy of Sciences has just been elected. Since I am interested in matters concerning higher education and the organization of science, I open a file for the new president. I have not the faintest idea that my neighbor and the president are in fact the same person. Yet such is the truth. They are the same person. My files are *de facto* co-referential, although I don’t know about it and I have no reason to suspect that they are. These files are not coordinated anyhow.

The whiz-and-hit case is certainly different. Although such a case would not license trading on identity, it would license a somewhat tentative identification: next time I hear a whiz, I move my head in fear of a hit—and I have good reasons to do so. A certain form of *coincidence* between two distinct files creates a sort of a link between them, a link that by no means *establishes* any form of coreference, but *suggests* one, and thus enables me to make some tentative inferences I could not make in the absence of such a link. Of course, putative linking cannot guarantee any co-reference at all. Ultimately, it may turn out that those were two different things: one that made the whiz and one that hit my head. Yet, initially, the files are in a sense *coordinated* through such a link—and in this sense, again, coordination falls beyond the reach of relationist accounts such as Goodman and Gray’s.

As the preceding discussion has made clear, the notion of content coordination as limited to synchronic cases and based upon the notion of trading on identity seems to be lacking explanatory powers needed to account for an array of phenomena that particularist versions of MFF handle relatively easily, notably weak co-reference *de jure* (vertical linking) and putative co-reference (putative linking).

---

Footnote 18 continued

for my earlier self, from the time I had such abilities, to give me a hand and help me to the other side of a fence. That is certainly impossible, for there is never any chance to deploy *now* a physical object in one of its past stages with all its physical powers transferred to the present moment. Yet we can deploy a mental representation in one of its past stages with all its inferential powers transferred to the present moment. We cannot always physically do now what we could do before, but we can always infer now what we could infer at a given stage of our past.

## 5 Realism about files is a virtue

Mental files as particulars are theoretical objects. Theoretical objects are objects that we do not have extra-theoretical access to, but which we assume to have existence “out there,” in the world, and certain properties. Such an assumption is supposed to explain, on the grounds of a given theory, a certain range of phenomena, or conceptualize a certain range of ideas.<sup>19</sup> Different theoretical perspectives might assume different properties of such objects, and thus draw different pictures of these objects. In particular, assuming that mental files exist and have certain properties is supposed to conceptualize certain ideas in truth-conditional semantics, while assuming they exist and have some other set of properties is supposed to explain certain phenomena in some areas of psychology. Eventually, we encounter a kind of ambiguity in thinking of theoretical objects in general, and of mental files in particular. First, we can think of them as of real entities that we attempt to model from different perspectives. Second, we can think of them as of proposed models of some entities, encompassing a certain perspective and highlighting, perhaps in a somewhat metaphorical way, certain features of these entities. The metaphor of files belongs to the layer of modelling real mental representations from a certain perspective. There are alternative models on the market, e.g., the model of mental *graphs*, anticipated in (Strawson, 1974, pp 45–46) and elaborated in (Pryor, 2016). However, neither the figurative character of the theoretical models of mental particulars nor the mere possibility of styling such models according to different templates would justify the denial of the existence of these particulars as such. On the contrary, as long as properties assumed within different perspectives are not outright inconsistent, a kind of realism about theoretical objects licenses the question about the identity of such objects posited by different theories and prompts interdisciplinary research. As long as it is fruitful and yields novel results, that is a virtue: “The [philosophical] theory of mental files [...] has connections to these various [linguistic, psychological] uses of the notion of file. These connections are well worth exploring, since they are what ultimately gives the theory its empirical bite” (Récanati, 2012, pp. vii–viii).

Indeed, the interdisciplinary dimension of MFF has proved to be fruitful. We can find a recent example in (Murez et al., 2020). The authors convincingly argue against a certain hypothesis within MFF (namely, the hypothesis that “mental files are a wide-ranging psychological natural kind underlying all and only singular thinking” (Murez et al., 2020, p. 135)), on the grounds of the empirical data taken from visual psychology. And while this particular outcome is of a negative character, it positively proves that the interdisciplinary exchange is valuable (corrections in one discipline based upon findings from another are particularly instructive). As the authors conclude, “[m]ental files thus constitute a particularly promising field of interdisciplinary investigation,

<sup>19</sup> “Explaining phenomena” seems fit rather to the empirical domain, while “conceptualizing ideas” is intended for a broader meaning, encompassing also theory-forming processes in philosophy. For instance, mental files are posited in order to—among other things—reconcile, within a unified framework, singular thinking and the distinction between coreference de jure and coreference de facto. Loosely, per analogiam to empirical sciences, we can say that singular thinking and the distinction CDJ-CDF are some *phenomena* and MFF is a theory that attempts to explain these phenomena. However, it would be more accurate to think of them as some *ideas* we consider important for our understanding of semantics, and try to arrange together within some coherent conceptual frame.

at the intersection of psychology, linguistics, and philosophy” (Murez et al., 2020, p. 136).<sup>20</sup>

In this light, the relationist approaches, which are likely to ignore the connections between philosophical, linguistic and psychological notions of mental files,<sup>21</sup> are inferior, *ceteris paribus*, to the particularist ones, which are not. Thus, even if we counterfactually assume, for the sake of argument, that thanks to some more or less “epicyclic” reformulations of some relationist approach, it would acquire roughly equal explanatory power to particularist MFF in the philosophical domain, we should not prefer it against the latter. On methodological grounds, such a theoretical transition would be a loss rather than a gain, a fragmentation of our knowledge instead of unification. It is a great advantage of particularist MFF that it creates a platform of unification for such diverse fields as semantics, discourse representation or visual psychology. There is no point in ruining this wide perspective without really good reasons.

## 6 Conclusion

In this paper, I have attempted to defend a particularist version of the mental files framework against a relationist challenge. In particular, I have argued for the following theses: (1) MFF, properly understood, has no problems with the notion of containment of information in files; on the contrary, this framework helps us keep clear important syntactic and cognitive insights. (2) It is not the case that file containment is representationally insignificant; on the contrary, it is significant as long as we allow a sense of significant permutation broader than a purely haecceitic switch. This significance, in turn, (3) undermines any relationist approach based on the notion of content coordination, as it is incapable of reflecting important distinctions of representational states. In particular, equivalence between enabling trading on identity and being governed by the same ER relation fails to reflect important aspects of the domain, such as those described in Pagin’s moth example, in which two representations are governed by the same ER relation, yet trading on their objects’ identity is blocked. In addition, the notion of content coordination cannot encompass the forms of coordination that do not license trading on identity, such as vertical linking or putative linking. Eventually, relationist accounts based on content coordination prove to be insufficient for replacing particularist MFF, as they lack explanatory power. The final thesis, (4), is that there are methodological considerations, connected with the idea of theoretical unification of different disciplines, which suggest that even if relationist accounts were explanatorily sufficient, we shouldn’t, *ceteris paribus*, abandon particularist MFF just for this reason.

<sup>20</sup> Other examples of the use of MFF in interdisciplinary research are (Newen & Wolf, 2020; Perner & Leahy, 2016; Perner et al., 2015).

<sup>21</sup> Goodman and Gray declare, for instance, that they are “concerned [...] with what we think of as *the philosopher’s notion* of a mental file [...]. Insofar as files play a role in visual psychology and linguistics these notions are not our target” (Goodman & Gray, 2020, p. 2).

**Acknowledgements** I am very grateful to Rachel Goodman, Aidan Gray, Michael Murez and François Recanati for valuable comments and discussion of earlier versions of this paper. Many thanks also to the anonymous referee, for prompting me to clarify my position in some important points and for helping me to avoid some embarrassing mistakes. However, all remaining faults are wholly mine.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Bartlett, G. (2018). Occurrent states. *Canadian Journal of Philosophy*, 48(1), 1–17. <https://doi.org/10.1080/00455091.2017.1323531>
- Dummett, M. (1981). *The interpretation of Frege's philosophy*. Duckworth.
- Dummett, M. (1995). *Frege: Philosophy of language* (2 ed. [3. Print]). Harvard University Press.
- Forbes, G. (1990). The indispensability of sinn. *The Philosophical Review*, 99(4), 535. <https://doi.org/10.2307/2185616>
- Frege, G. (1892). On concept and object. In P. T. Geach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege* (2nd ed., pp. 42–55). Blackwell.
- Geach, P. T. (1975). Names and identity. In S. Guttenplan (Ed.), *Mind and language* (pp. 139–158). Clarendon Press.
- Geach, P. T. (1980). *Reference and generality*. Cornell University Press.
- Goodman, R. (2016). Against the mental files conception of singular thought. *Review of Philosophy and Psychology*, 7(2), 437–461. <https://doi.org/10.1007/s13164-014-0209-0>
- Goodman, R., & Gray, A. (2020). Mental filing. *Noûs*. <https://doi.org/10.1111/nous.12354>
- Hale, B. (1979). Strawson, geach and dummett on singular terms and predicates. *Synthese*, 42, 275–295.
- Husserl, E. (2001). *Logical investigations*. Edited by Dermot Moran. *International Library of Philosophy*. Routledge.
- Maier, E. (2016). Attitudes and mental files in discourse representation theory. *Review of Philosophy and Psychology*, 7(2), 473–490. <https://doi.org/10.1007/s13164-015-0296-6>
- Murez, M., & Recanati, F. (2016). Mental files: An introduction. *Review of Philosophy and Psychology*, 7, 265–281.
- Murez, M., Smortchkova, J., & Strickland, B. (2020). The mental files theory of singular thought: A psychological perspective. In R. Goodman, J. Genone, & N. Kroll (Eds.), *Singular thought and mental files* (pp. 107–142). Oxford University Press. <https://doi.org/10.1093/os0/9780198746881.003.0006>
- Newen, A., & Wolf, J. (2020). The situational mental file account of the false belief tasks: A new solution of the paradox of false belief understanding. *Review of Philosophy and Psychology*, 11(4), 717–744. <https://doi.org/10.1007/s13164-020-00466-w>
- Ninan, D. (2015). On Recanati's mental files. *Inquiry*, 58(4), 368–377. <https://doi.org/10.1080/0020174X.2014.883751>
- Pagin, P. (2013). The cognitive significance of mental files. *Disputatio*, 5(36), 133–145.
- Perner, J., Huemer, M., & Leahy, B. (2015). Mental files and belief: A cognitive theory of how children represent belief and its intensionality. *Cognition*, 145, 77–88. <https://doi.org/10.1016/j.cognition.2015.08.006>

- Perner, J., & Leahy, B. (2016). Mental files in development: Dual naming, false belief, identity and intentionality. *Review of Philosophy and Psychology*, 7, 491–508.
- Prosser, S. (2020). The metaphysics of mental files. *Philosophy and Phenomenological Research*, 100(3), 657–676. <https://doi.org/10.1111/phpr.12577>
- Pryor, J. (2016). Mental graphs. *Review of Philosophy and Psychology*, 7, 309–341.
- Recanati, F. (2021b). 'Mental files'. In P. Stalmaszczyk (Ed.), *The Cambridge handbook of the philosophy of language*. Cambridge Handbooks in Language and Linguistics. Cambridge University Press. <https://www.cambridge.org/core/books/cambridge-handbook-of-the-philosophy-of-language/6283AABA41E79084E9EA2E09DD459130>.
- Recanati, F. (2012). *Mental files* (1st ed.). Oxford University Press.
- Recanati, F. (2013). Mental files: Replies to my critics. *Disputatio V*, 36, 207–242.
- Recanati, F. (2015). Replies. *Inquiry*, 58(4), 408–437. <https://doi.org/10.1080/0020174X.2015.1033868>
- Recanati, F. (2016). *Mental files in flux, 1st edition. Lines of thought*. Oxford University Press.
- Recanati, F. (2021). Do mental files obey Strawson's constraint? In C. Borgoni, D. Kindermann, & A. Onofri (Eds.), *The fragmented mind*. Oxford University Press.
- Searle, J. R. (1983). *Intentionality, an essay in the philosophy of mind*. Cambridge University Press.
- Strawson, P. F. (1974). *Subject and predicate in logic and grammar*. University Paperbacks; 538. Methuen; distributed by Harper & Row, Barnes & Noble Import Division.
- Tałasiewicz, M. (2017a). Mental files. Triggering mechanisms, metadata and "Discernibility of Identicals." *Studia Semiotyczne*, 31(2), 13–34.
- Tałasiewicz, M. (2017). The sense of a predicate. *Philosophy and Logic of Predication*, 7, 129–43.
- Williamson, T. (2007). *The philosophy of philosophy. The Blackwell/Brown Lectures in Philosophy 2*. Blackwell Pub.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.