



HAL
open science

Do Mental Files Obey Strawson's Constraint?

François Recanati

► **To cite this version:**

François Recanati. Do Mental Files Obey Strawson's Constraint?. Oxford University Press. The Fragmented Mind., 2021. hal-02923236

HAL Id: hal-02923236

<https://hal.science/hal-02923236>

Submitted on 26 Aug 2020

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

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

Do Mental Files Obey Strawson's Constraint?

François Recanati
Collège de France

1. Frege cases and fragmentation

Rational subjects are supposed to avoid inconsistencies, but to avoid inconsistencies you need to be able to detect them. An inconsistency may be undetectable—hence it may escape the attention of a rational subject—in a variety of circumstances, of which I will mention two.

The first type of circumstance is well known. It corresponds to the class of examples which Jerry Fodor dubbed 'Frege cases' (Fodor 1995). What is objectively the same propositional content¹ may be presented through two distinct linguistic formulations. Unaware that this is the same propositional content, because of the difference in mode of presentation, the subject may accept the proposition (believe it) under one guise and reject it (disbelieve it) under the other guise. Thus the subject may accept the proposition that 'Phosphorus shines in the morning' while dissenting from the proposition that 'Hesperus shines in the morning'. Since Hesperus = Phosphorus, this is the same proposition (in the sense that the same property is ascribed to the same object)—a proposition which the subject both accepts and rejects. This is objectively a contradiction, but the subject is unable to detect it, so her rationality is not impugned. Indeed, as soon as she realizes that Hesperus = Phosphorus, she will stop holding these contradictory attitudes toward the proposition. That is because, when the subject realizes that Hesperus = Phosphorus, the inconsistency becomes detectable.

A caveat: the relevant notion of mode of presentation is not restricted to matters of linguistic formulation. There may be two distinct modes of presentation even if the linguistic formulation is the same. Thus, in a famous example due to John Perry, the subject assents to 'That ship is an aircraft carrier' (when looking at the ship through the right window) and dissents from 'That ship is an aircraft carrier' (when looking at what is in fact the same ship seen through the left window).² Or, in another famous example, due to Saul Kripke (1979):

¹ By saying that the propositional content is objectively the same I mean, e.g., that the same property is ascribed to the same object.

² Versions of this example are discussed in Perry (1993: 12–13) and Evans (1982: 84).

the subject assents to 'Paderewski is musical' (when thinking of Paderewski the musician) and dissents from 'Paderewski is musical' (when thinking of Paderewski the politician, whom he wrongly takes to be a different person). Modes of presentation are *ways of thinking of an object* (Evans 1982).³ The standard criterion of distinctness for modes of presentation is Frege's Constraint (Schiffer 1978: 180): if it is possible for a rational subject to believe of one and the same object both that it is *F* and that it is not *F*, that means that the subject is thinking of that object under two distinct modes of presentation (whether or not the modes of presentation in question are associated with specific linguistic formulations).

The other type of circumstance in which a contradiction may be undetectable corresponds to the fragmentation idea, that is, the idea that the mind is 'divided' or 'compartmentalised' (Lewis 1982; Davidson 1986; Egan 2008). Davies and Egan summarize the idea as follows:

Actual belief systems are fragmented or compartmentalised. Individual fragments are consistent and coherent but fragments are not consistent or coherent with each other and different fragments guide action in different contexts. We hold inconsistent beliefs and act in some contexts on the basis of the belief that *P* and in other contexts on the basis of the belief that not-*P*. Frequently we fail to put things together or to 'join up the dots'. It can happen that some actions are guided by a belief that *P* and other actions are guided by a belief that if *P* then *Q*, but no actions are guided by a belief that *Q* because the belief that *P* and the belief that if *P* then *Q* are in separate fragments. For example, some people are described as having their philosophical beliefs and their religious beliefs in separate compartments. David Lewis describes himself as having once had fragmented beliefs about the geography of Princeton. According to one fragment, Nassau Street ran north–south and was parallel to the railway track; according to another fragment, the railway track ran east–west and was parallel to Nassau Street (Lewis 1982, p. 436). (Davies and Egan 2013: 705–6)

³ There are also modes of presentation for properties and relations, but I am only concerned with modes of presentation for objects here.

In other words, the subject holds contradictory beliefs, but they do not interact (hence the contradiction remains undetected), because they are stored in *different compartments or 'territories' of the mind*. As Davidson writes,

People can and do sometimes keep ... opposed beliefs apart. To this extent we must accept the idea that there can be boundaries between parts of the mind; I postulate such a boundary somewhere between any (obviously) conflicting beliefs. ... The contradictory beliefs do not belong to the same territory; to erase the line between them would be to destroy one of the beliefs. (Davidson 1986: 91–2)

Elga and Rayo (Chapter * in this volume) illustrate the fragmentation phenomenon with the well-known division between two processing pathways for visual information. Owing to this division between the ventral stream and the dorsal stream, information about the orientation of a slot or the width of a block is available to a certain brain-damaged patient for the purposes of *acting* (e.g. inserting a card into the slot, or picking up the block), but not for the purposes of *reporting*—verbally or otherwise—what the block or the slot is like. As Milner and Goodale write, the brain-damaged patient ‘could guide her movements using visual cues of which she seemed completely unaware’ (Goodale and Milner 2006: 660). In a forced choice task, such a subject might be tempted to dissent from, or suspend assent to, a proposition (about the size of the block or the orientation of the slot) whose acceptance is nevertheless entailed by her behaviour. In many similar cases not involving brain damage, the subject’s behaviour justifies ascribing to her certain beliefs or propositional attitudes, but these beliefs are not beliefs the person consciously holds and is prepared to declare. Stalnaker mentions the case of ‘a shrewd but inarticulate chess player who may be able to access information for the purposes of choosing a move even if she is unable to access that same information for the purposes of answering a question, or explaining why she moved as she did’ (Stalnaker 1999: 254). In some extreme cases the subject explicitly (and forcefully) rejects the beliefs that her behaviour nevertheless betrays. Schwitzgebel, whom Elga and Rayo also cite, discusses the case of Ralph, who

is prepared to argue coherently, authentically, and vehemently for equality of intelligence [between men and women] and has argued the point repeatedly in the

past. And yet Ralph is systematically sexist in his spontaneous reactions, judgments, and unguarded behavior. ... I think it is not quite right to call Ralph an out-and-out sexist who simply believes that women are intellectually inferior. What Ralph says and how he reasons in his most abstract and most thoughtful moments is an important part of how he thinks and acts, even if it is only a part. ... His dispositions, his patterns of response, his habits of thought, are mixed up and inconsistent. It is neither quite right to say that he believes in the intellectual equality of the sexes nor quite right to say that he fails to believe that. (Schwitzgebel 2002: 191–2)

In all these cases, according to these authors, the mind is divided: one part of the mind hosts an attitude, another part hosts a conflicting attitude, and the two parts don't communicate with each other. That is why the inconsistency remains undetected.

One of my aims in this chapter is to show that an approach to modes of presentation which has gained momentum recently, namely the mental file approach, makes it possible to treat Frege cases themselves as an instance of fragmentation, that is, as cases in which conflicting pieces of information are stored in the subject's mind but remain insulated from each other, in such a way that the inconsistency cannot be detected. I will first spell out the relevant notion of mental file (§2) and the assumption that is needed to make Frege cases, understood in terms of mental files, an instance of fragmentation (§3). That assumption—the principle of informational segregation—is what motivates a constraint on files which derives from Strawson's work and which, for that reason, I call Strawson's constraint. I will then present the reasons why subsequent mental file theorists, from John Perry to my earlier self, have given up Strawson's constraint and, with it, the segregation principle (§4). In the last two sections of the chapter (§5–6), I will offer a defence of both the constraint and the principle.

2. *Mental files*

Mental files have a dual function: a file is a repository for information about a particular object (this is the *storage* function), and it (or its 'label') is used to stand for that object in thought (this is the *referential* function). It is not assumed that the reference of a file—the object it stands for—is determined *by* the information it contains, however. On the contrary, most mental file theorists think of the reference of files as being determined relationally

rather than satisfactorily. In my account, for example, mental files are based on ‘epistemically rewarding relations’ to entities in the environment (Recanati 2012). These relations are said to be epistemically rewarding because they make it possible for the subject to gain information from the object, information which can then be stored in the file. Thus if I see an object in front of me, my relation to that object enables me to gain information about it in perception, and that information is fed into the temporary file based upon the current perceptual relation (a demonstrative file deployable in thought: ‘that object’). But it is the relation, not the information stored in the file, which determines the reference of the file. The reference is the object I am actually seeing, and that object may fail to satisfy the predicates stored in the file, as in this example from Donnellan:

One is at a party and, seeing an interesting-looking person holding a martini glass, one asks, ‘Who is the man drinking a martini?’ If it should turn out that there is only water in the glass, one has nevertheless asked a question about a particular person, a question that it is possible for someone to answer. (Donnellan 1966/1977: 48)

The speaker deploys a mental file based on the perceptual relation to the interesting-looking person she is watching.⁴ That file contains the piece of (mis-)information ‘drinks a martini’, which the speaker uses to construct her description of the individual she has in mind, ‘the man drinking a martini’. But what the file refers to is the man in question—the man the speaker is watching—even though that man is not actually drinking a martini.

I draw a basic distinction between two types of mental file: a distinction I must introduce because it will play an important role in what follows. *Indexical files* are based on specific epistemically rewarding relations (henceforth: ER relations), in terms of which the files are typed. Thus there are demonstrative files for objects one is currently perceiving (as in the above examples), memory files for objects one has perceived in the past, deferential files for objects one has heard or is hearing about, recognitional files for objects multiple encounters with which have left us with a disposition to recognize them, and so on and so forth. In addition to these indexical files, based on specific ER relations, there are also *encyclopaedic files*, which are not based on specific ER relations but track a particular object

⁴ To understand what the speaker is saying, the hearer will have to look in the same direction, in order to open (if she hasn’t done so already) a file based on the same ER relation and coordinated with the speaker’s own file.

and exploit whichever ER relation to that object happens to be available in context to serve as an information channel. Thus my file about Mont Blanc, in contrast to my demonstrative file ‘that mountain’ (deployed while looking at Mont Blanc), is not tied to any particular way of gaining information, nor to any particular relation to the mountain. It may exploit a number of distinct ER relations to the reference in an opportunistic manner. Any relation will do as an information channel, provided it preserves the link to the object.

That they are not based on *specific* ER relations does not mean that encyclopaedic files are not based on ER relations *simpliciter*. In Recanati (2012) I describe them as being based on a *higher-order* ER relation, namely

$$\lambda x \lambda y [(\exists R) (Rx, y)],$$

where ‘*R*’ ranges over (specific) ER-relations. The higher-order relation holds between the subject *x* and some object *y* just in case there is some specific ER relation or other through which *x* can gain information from *y*.

Peter Strawson was the first mental file theorist, and he was only concerned with encyclopaedic files. Strawson construed the subject’s mental encyclopaedia (his representation of the world) as a system of files, where the files correspond to objects in a one–one manner. Millikan summarizes Strawson’s picture as follows:

On Strawson's picture, the identity of a particular is represented in the mind by the identity of another particular [viz., the mental file]. So long as you haven't made any mistakes, everything you know about your mother is attached to the same particular mental representation of your mother, the same token. Your understanding that all these facts are facts about the same woman consists in the representations of the logical predicates of each of these facts being attached to numerically the same ‘dot’ in your mind or brain. (Millikan 1997: 500)

On that picture, if there are two files in the subject’s mind, that means that (according to the subject) there are two objects out there. The architecture of the file system in the subject’s mind reflects her ontological commitments. Accordingly, two files should be merged when the subject realizes that they are about the same object, and a file should be split when the

subject realizes that it corresponds to two distinct objects. The role of identity statements is to trigger the first kind of cognitive restructuring: two ‘segregated bundles or clusters of identifying knowledge’ are ‘brought together and tied up into one for a given audience of an identity statement’ (Strawson 1974: 52).⁵

The idea that one should merge two files if one realizes that they have the same referent is what I call *Strawson’s constraint*. One might think that the main motivation behind it is a desire to avoid redundancy, that is, some kind of parsimony principle. But there is a deeper motivation, which takes us back to the fragmentation issue.

3. Informational segregation

A certain pattern of inference, which Campbell dubbed ‘trading on identity’, is illustrated by the pieces of reasoning on the left-hand side of Table 1:

(1)	(1*)
Cicero was Roman	Cicero was Roman
Cicero was an orator	Tully was an orator
Therefore, someone was a Roman orator	Therefore, someone was a Roman orator

(2)	(2*)
Chomsky is a linguist	Chomsky is a linguist
He _{anaphoric} is an activist	He _{deictic} is an activist
Therefore, someone is both a linguist and an activist	Therefore, someone is both a linguist and an activist

Table 1

⁵ See also Lockwood (1971: 209): ‘In the making of an (intendedly informative) identity statement, the speaker assumes that the hearer has more than one file on the subject of his assertion. Consequently, the respective roles of the two substantival terms used in the making of such a statement will not be sufficiently understood when all that is known is their intended reference. Further, it must be appreciated that these terms are intended to indicate, for the hearer, *distinct* mental files. The purpose of an identity statement, which will be fulfilled if it is accepted as true, is precisely to get the hearer to merge these files or bodies of information into one.’

Trading on identity (TI) has the following general form (where ‘*a*’ and ‘*b*’ stand for *token* singular terms, which may or may not be of the same type):

a is *F*

b is *G*

Therefore, there is an *x* which is *F* and *G*

The inference is valid only if the subject terms in the two premises of the argument are suitably ‘coordinated’, to use a phrase due to Kit Fine (2007). In the arguments on the right-hand side, the singular terms in the premises—‘Cicero’ and ‘Tully’ in (1*); ‘Chomsky’ and ‘he_{deictic}’ in (2*)—are not coordinated, even if they are coreferential, so trading on identity is illegitimate. To get to the conclusion (that some *x* is both *F* and *G*), we need an additional premise to the effect that the individuals talked about in the first two premises are identical. Thus to go from ‘Cicero was Roman’ and ‘Tully was an orator’ to ‘Someone was a Roman orator’, we need the additional premise ‘Cicero = Tully’. Likewise, to go from ‘Chomsky is a linguist’ and ‘he_{deictic} is an activist’ (or: ‘That man is an activist’) to ‘Someone is both a linguist and an activist’, we need an additional premise: he_{deictic} /that man = Chomsky. But in the arguments on the left-hand side, *no additional premise is necessary*. (Nor is it possible to provide one without launching a regress.)⁶ As Campbell says, ‘we trade directly upon co-reference, moving directly to the conclusion’ (Campbell 1987: 275–6). Identity here is given or presupposed; it does not have to be stated.

What does it mean to say that, in the arguments on the left-hand side, the singular terms are coordinated in such a way that coreference between them is given or presupposed? Descriptively, it means that ‘anyone who raises the question of whether the reference was the same would betray his lack of understanding of what you meant’ (Fine 2007: 40). Theoretically, there is no consensus as to what coordination boils down to. Some, like Fine, take it to be a primitive notion, while others analyse it in terms of identity: identity of expression at an abstract syntactic level (Fiengo and May 1994), or identity of sense in a Fregean setting (Campbell 1987). In the mental file framework, two expressions are said to

⁶ See Recanati (2012: 47–50) and the references therein.

be coordinated just in case they are associated with *the same mental file* (Recanati 2012, 2016).

According to the mental file story, trading on identity is only possible if the pieces of information corresponding to the predicates F and G in the two premises of the argument are stored in the same mental file. Via the file it is contained in, predicative information is bound to a particular object—the same object for all the predicative information within the file. That the two predicates F and G are predicated of the same object is therefore guaranteed as a matter of cognitive architecture when they belong to the same file, thus making TI possible. If the same information is distributed across files, it is not architecturally guaranteed that it will concern one and the same object, and TI is not possible.

The difference between the right-hand side and the left-hand side of the table is taken to be this: on the left-hand side, the pair of singular terms in the premises, whether they are the same singular term or not, are associated with the same mental file, while on the right-hand side they are associated with distinct files. When the pieces of information corresponding to the predicates in the premises are stored in distinct files (e.g. a Cicero file and a Tully file), trading on identity is *not* possible.

In that framework, an explanation in terms of fragmentation can be given for why the inconsistency is undetectable in a Frege case. The piece of information ‘shines in the morning’ is stored in the Phosphorus file, while the contradictory piece of information (‘does not shine in the morning’) is stored in the Hesperus file. If the two pieces of information were stored in a single file, trading on identity would be validated, and the contradiction (that some x is both F and not- F) would become apparent. But the information in one file is insulated from information in the other file, even though the two files corefer, so the contradiction cannot be derived.

As I said, this is an explanation of Frege cases in terms of fragmentation. Mental files are, to use Strawson’s terms, ‘segregated bundles’ of information. Segregation consists in the fact that information in one file is insulated from information in other files, even if the files corefer. In other words, trading on identity is only possible *within* files. This principle of informational segregation is what motivates Strawson’s constraint. Knowing of an object both that it is F and that it is G does not by itself put you in a position to exploit that information productively so as to derive the consequence that that object is F and G . The consequence can only be derived if the two pieces of information (that the object is F , and

that it is *G*) are stored in the same file. That is why proper acknowledgment of an identity requires the initially distinct files to be merged.

4. *The link/merge debate*

Later mental file theorists have given up Strawson's constraint and have proposed that two files be linked (rather than merged) when it is realized that they corefer. Linking preserves the distinctness of files but makes *free information flow* possible between them. The subject still maintains two distinct files for the same object, say a Hesperus file and a Phosphorus file, but the link established between the files makes it possible to transfer any information in one file into the other. In this way, inconsistencies become detectable: thanks to linking, the predicate 'does not shine in the morning' in the Hesperus file becomes visible from the Phosphorus file, which hosts the contradictory predicate 'shines in the morning'. In other words, linking achieves exactly the same effect as merging while preserving the distinctness of the files. In particular, linking is like merging in that it makes trading on identity possible, but in contrast to merging, it makes trading on identity possible *across* files. That means that the principle of informational segregation is rejected: it is not true that trading on identity is only possible *within* files.

The following type of example is supposed to illustrate the possibility of trading upon identity across files:

(3)

That vase was in good condition (based on memory)

That vase is broken (based on perception)

Therefore, something that was in good condition is broken

Here, arguably, two distinct files for the vase are deployed (one, in the first premise, based on memory; the other, in the second premise, based on perception), but despite the fact that they are distinct files, based on distinct ER relations, trading on identity is possible because the subject recognizes the vase he currently sees as the vase he has seen before and remembers. Owing to this act of (putative) recognition, the files are linked, and this makes TI possible across files. Or so the argument goes.

I will return to the analysis of (3) in Section 5. For the time being, let me spell out the three main arguments that have been offered in support of the linking proposal and against Strawson's merge model.

First, in contrast to merging, linking is a reversible operation. If, having been told that Hesperus is Phosphorus, I merge my files into a single, inclusive Venus file, what will happen if it turns out that I have been misinformed and that Hesperus is not identical to Phosphorus after all (Forbes 1990: 542–3; Millikan 1997: 508)? How will I *undo* the operation once the information in the initial files has been mixed and thrown into a 'single undifferentiated bundle' (Millikan 1990: 192)? Substituting linking for merging provides a solution to that problem:

An agent can retain two notions [= files] of an individual, while linking them, in the way one does when one recognizes that 'two' of one's acquaintances are actually a single individual. Why might two notions be retained when such a recognition takes place? One reason for this would be to allow the possibility of easy revision in case the 'recognition' was in error. (Crimmins and Perry 1989/1993: 256)

If the identification $[A = B]$ is tentative, the notions [= files] may retain their identity; if not, they may merge and become one. (Perry 2002: 196)

The second reason for making room for a linking operation less radical than merge is this. Even when we know for certain that $A = B$, it may be important to retain two distinct files, because *other people* may not know the identity. In order to describe these people's beliefs, or in order to understand them by mentally putting ourselves in their shoes, or in order to be able to rationally engage with them (and convince them of their mistake), we may need to simulatively deploy two distinct files for the same object. Suppose we want to inform them of the identity by telling them: 'Hesperus is Phosphorus'. To do so we need to deploy, vicariously as it were, the two files which we know refer to the same planet. So these files *still need to be available* after we discover the identity. They also have to be available for us to remember how we used to think of the matter before learning of the identity. As Bach writes:

Forbes mentions ... that when one comes to believe an identity, two files ... come to be merged, but he also points out that after the merging one can, at least for a while, distinguish what one believed about the individual in question under each of the two modes of presentation. That is, the merging does not immediately obliterate the two original files. (Bach 1994: 305)

Indeed, if we are to describe our earlier thoughts, we will have to deploy the initial files which, for that reason, must be retained. The following example, from Pinillos (2011: 314–16), provides a striking illustration:

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

In the first clause, the speaker espouses the point of view of the unenlightened (herself and her peers before they came to realize that Hesperus is Phosphorus), and she refers to Venus twice, via two distinct mental files (respectively associated with the names ‘Hesperus’ and ‘Phosphorus’) rather than via a single inclusive Venus file corresponding to her current point of view. The inclusive Venus file is associated with the demonstrative adverb ‘there’ at the end of the second clause. That second clause expresses the subject’s current point of view, while the first clause is phrased from the point of view of the speaker and her peers before they learnt the identity (when they thought there were two distinct stars).

The third reason for going for linking rather than merging is the most important one. Files are based on epistemically rewarding relations (ER relations), and we typically stand in distinct ER relations to the same object; so it is important to be able to think of the same object under distinct modes of presentation (through distinct mental files), even if one realizes that they are modes of presentation of the same object. Strawson’s constraint makes sense if we are exclusively concerned with *encyclopaedic* files, as Strawson was. Such files abstract from specific ER relations—they only depend on the *existence* of ER relations to the reference (whatever they are)—so there is no point in maintaining distinct encyclopaedic files about one and the same object: all such files would be based on the same (higher-order) relation to the object, so their multiplicity could only reflect the mistake of thinking that there are two objects where there is only one. But the constraint evidently

does not work for indexical files, which are tied to particular ER relations. Again, because our relations to the reference are contextually diverse and changing, it may be perfectly legitimate to take different perspectives on the same object simultaneously (by thinking of it through distinct indexical files, based on distinct ER relations). If I recognize the man I see in the mirror as myself, I entertain two representations of myself: as the man I see in the mirror, and as myself. When I think 'That man is me' or 'I am that man', I simultaneously think of myself under these two distinct modes of presentation, which correspond to different relations to myself, while being aware that they are modes of presentation of the same person. No norm is thereby violated. Accordingly, when the subject realizes that the two files concern the same object, she does not have to merge them. What she has to do, rather, is to *link* them so as to let information flow freely between them (as if they were a single file).

I said earlier that Strawson was only concerned with encyclopaedic files. Interest in indexical files arose a bit later, in reaction to the 'problem of the essential indexical', which gave prominence to the notion of an indexical mode of presentation (Perry 1993). Thinking of an individual as oneself and thinking of the same individual as 'that man' count as thinking of that individual under distinct modes of presentation, and when the two modes of presentation are exercised simultaneously, this can give rise to Frege cases. But, as John Perry and the other contributors to that literature have shown, the difference between these indexical modes of presentation cannot be spelled out in descriptive terms. The difference, rather, has to do with the relations in which the subject stands to the individual he is thinking about. It was thus hypothesized that indexical modes of presentation are (not descriptions but) mental files based on specific relations to the reference. By deploying the file (or its 'address' or 'label') in thought, the subject can think about the object *in virtue of standing in the relevant relation to it*.

Given the multiplicity of relations in which one can stand to one and the same object, there is no longer any reason to demand a one–one relation between the files in the subject's mind and the objects out there that the subject is thinking about. For any object of thought, there are many ways of thinking of that object that are available, and it is this plurality which gives rise to Frege cases. There is nothing intrinsically abnormal in the fact that different modes of presentation are available for the same object. The problem with Frege cases only arises because the subject *does not recognize* that these are modes of

presentation of the same object. It is that unawareness which leads to inconsistencies. But if the subject recognizes that two modes of presentation (two mental files) refer to the same object, she will link them, and that will be sufficient to get rid of the problem. Let me close this section with a citation from my own book, *Mental Files*, where this line of thought was vigorously pursued:

The 'merge' model is incompatible with the mode of presentation idea we are trying to cash out (Millikan 2000: 147–9). It is of the essence of modes of presentation that there can be a multiplicity of modes of presentation for the same object. On the picture I have presented, mental files qua nondescriptive modes of presentation correspond to various relations in which the subject stands to objects, and there is no doubt that a subject can and typically does stand in several relations simultaneously to the objects in his or her environment. Nor is this situation contrary to some normative ideal, as if the coexistence of several files for a single object was a defect to be avoided whenever possible. ... Taking seriously the idea that mental files are modes of presentation based upon contextual relations to objects demands that we accept the existence of a multiplicity of files for the same object even when the files are linked and the subject is aware that they stand for a single object. (Recanati 2012: 45–6)

5. In defence of the segregation principle

In arguing against the merge model in *Mental Files*, I invoked the phenomenon of recognition, construed as involving the simultaneous deployment of two distinct modes of presentation for the same object:

Imagine that I see a certain man cutting his grass and recognize him as Noam Chomsky. (Or imagine I learn he is Noam Chomsky, through an identity statement which I accept.) My perceptual file and my Chomsky file get linked, but there is no reason why either should disappear. Perry describes the perceptual buffer as being 'absorbed' into the more permanent file in such cases, but I think the buffer should only disappear when the ER relation on which it is based no longer holds. (Recanati 2012: 4–6)

But it is far from obvious that recognition involves the simultaneous deployment of two distinct modes of presentation for the same object, and I now believe this line of reasoning (in favour of the linking model) should be resisted.

If I recognize the man I see as Chomsky, then the file I deploy, after recognition, is my Chomsky file. (If I *immediately* recognize the man I see as Chomsky, then the file I deploy is my Chomsky file *from the start*.) I do *not* deploy the demonstrative file ‘that man’ after recognition. If you are not convinced, consider the structurally similar example in which I see myself in the mirror. Once I recognize that the man I see in the mirror is me, I no longer think of ‘him’ as ‘that man’; I think of him as me. This suggests that the demonstrative way of thinking is *no longer available* post-recognition. If that is right, then *it is not true that recognition involves the simultaneous deployment of two distinct modes of presentation*. Rather, one mode of presentation *supersedes* the other. This casts doubt on the linking model.

The example I mentioned in Section 4, in support of the idea that linking makes trading on identity possible *across* distinct files, also involves a form of recognition, and I think it too should be reanalysed. Let me repeat the example:

(3)

That vase was in good condition

That vase is broken

Therefore, something that was in good condition is broken

it was claimed that the first premise in (3) involves a file based on memory, while the second premise involves a file based on perception. Despite the fact that they are distinct files, it was argued, trading on identity is possible because the subject recognizes the vase and links the files. But one can deny that two distinct mental files are deployed in (3). Rather than involving two distinct files for the vase, one based on memory and the other one on perception, the reasoning in (3) arguably involves *two deployments of the same file based on a composite ER relation involving both memory and perception* (Recanati 2013, 2016).

In the passage from *Mental Files* cited at the beginning of this section, I too rapidly dismissed Perry’s suggestion to the effect that, after recognition, ‘the perceptual buffer [i.e.

the demonstrative file, based on perception] is ‘absorbed’ into the more permanent file’. In a footnote, however, I entered a caveat:

For qualifications, see chapter 7 on the ‘conversion’ of files. What I call ‘incremental conversion’ in that chapter can be described as a form of absorption (of one file by another). (Recanati 2012: 46 n.)

Incremental conversion is the process whereby, as more ways of gaining information about an object become available in context, a file based on a certain ER relation r_1 to that object dynamically evolves into a file based on a *set* of ER relations including r_1 among its members. For example, I hear of a certain person, named ‘Marcel’, and open a deferential file for him (a file purely based on testimony). Later, as I am introduced to Marcel in the flesh, new ways of gaining information about him become available, and my initially deferential file is transformed into a file based on more ER relations than just testimony. The same process arguably happens in the situation corresponding to the piece of reasoning in (3). Initially, the subject had a memory of the vase derived from a previous perceptual encounter, but then she encounters the vase again and recognizes it as ‘that vase’. She gets ‘reacquainted’ with the vase. The acquaintance relation in which she now stands to the vase is a composite ER relation involving both memory *and* current perception. Here, as in the deferential case discussed above, incremental conversion transforms an initial file (based on memory) into a different file based on a richer set of ER relations. *It is this output file, based on a composite ER relation (the acquaintance relation), which is deployed twice in (3), in association with each of the premises. Since that file simultaneously rests on two distinct ER relations, namely memory and perception, it can host the two predicates together* (‘was in good condition’, based on memory, and ‘is broken’, based on perception). That is why TI is possible. If this is correct, then trading on identity does not take place across different files, as suggested in my earlier presentation, but entirely *within* the file that is the *output* of incremental conversion.⁷

⁷ The same type of analysis can be provided for structurally similar examples involving trading on identity across modalities:

(4)

That glass is dirty (based on sight)
It/that glass is cold (based on touch)

In example (3) the post-conversion file, based on memory and perception, is a more inclusive file than the initial file based on memory alone, since it contains the information in the memory file plus new information gained through perception. Incremental conversion can indeed be described as a form of absorption of one file by another, more inclusive file. This is like merging, except that merging takes two initial files and yields as a result an inclusive file, while incremental conversion takes one initial file and yields as a result a more inclusive file. In both cases, the operation makes trading on identity possible, and in both cases, *trading on identity takes place within the inclusive file* (rather than ‘across’ files).

Earlier I said that if I *immediately* recognize the man I see as Chomsky, then the file I deploy is my Chomsky file *from the start*. It is important to realize that in a case of immediate recognition (e.g. when, at home, I see my wife reading a book in the living room), a single file is involved. The rich array of ER relations that file is based on includes the current perceptual relation, but no separate ‘demonstrative’ file for ‘that person’ is deployed and linked to the more permanent file for Chomsky or for my wife. Nor do we have to merge two files. Again, a single file is involved from the start, and trading on identity is possible within that file, as in the following train of reasoning:

Chomsky is cutting the grass (based on perception)

Chomsky has bought this house for 1 million dollars (based on testimony)

A man who has bought this house for 1 million dollars is cutting the grass

In non-immediate recognition, e.g. the original Chomsky example, the situation is different. A demonstrative file (‘that man’) is first deployed, based on the current perceptual relation. That file is distinct from the subject’s permanent Chomsky file, and it coexists with it in the pre-recognition phase. The permanent Chomsky file is an encyclopaedic file based on a rich and open-ended array of ER relations. When I recognize the man I see as Chomsky, two things happen: first, the current perceptual relation (on which the demonstrative file was based) is added to the pool of ER relations of the encyclopaedic file; second, the demonstrative file itself is suppressed, superseded by the encyclopaedic file. The permanent

Therefore, something is dirty and cold

Even though distinct modalities are involved in gaining information from the glass, a single file is deployed twice, fed by these different ways of gaining information. Again, there is no need to depart from the principle of informational segregation, according to which trading on identity is only possible within files.

Chomsky file is thus enriched and ‘absorbs’ the demonstrative file, as Perry insightfully suggested.

In the case I have just described, the two initial files of the pre-recognition phase (the demonstrative file for ‘that man’ and the permanent Chomsky file) are sort of ‘merged’, but this is not *exactly* merging, for the following reason. The output file resulting from the ‘merge’ is nothing other than the Chomsky file itself since the addition of a new ER relation does not change the identity of the file (qua encyclopaedic file based on the higher-order ER relation or, equivalently, on an open-ended set of first-order ER relations). So this type of case is intermediate between standard merging (which normally involves three files: the two initial files and the inclusive file) and the case of reacquaintance (which involves two files overall: the initial memory file, and the more inclusive file which is the output of incremental conversion). I call this ‘quasi-merging’. As in cases of standard merging, two distinct initial files are involved; this is in contrast to incremental conversion, which involves a single initial file. Unlike cases of standard merging, however, only two files are involved overall in the Chomsky example, as in the case of incremental conversion. Standard merge cases involve three files, as when the Hesperus file and the Phosphorus file are merged into a single Venus file, or when the demonstrative file for the ship seen from the left window and the demonstrative file for the ship seen through the right window are merged into a single demonstrative file when it is realized that there is only one ship. Figure 1 below summarizes the differences between the three types of case (merging, quasi-merging—the Chomsky case—and incremental conversion, illustrated by the case of reacquaintance).




Figure 1: three types of case

In all these cases, whatever their differences, a transition (corresponding to the phenomenon of ‘recognition’) occurs, leading from an initial phase featuring one or two

initial file(s) to an output phase featuring a more inclusive file.⁸ The inclusive file does not coexist with the initial file(s), but supersedes them. This argues against the linking model. And the reason why recognition makes trading on identity possible, in the three cases, is that the output file is inclusive enough to simultaneously host all the predicative information involved in TI. This is consistent with the principle of informational segregation, according to which trading on identity is only possible within files.

6. *In defence of Strawson's constraint*

In Section 4, I listed the three arguments that have persuaded mental file theorists to give up Strawson's constraint and to adopt the linking model instead of the merge model. In Section 5, I started to rebut the third argument in the list (the main argument, as I called it in Section 3).

The argument in question starts from two premises:

- (i) Mental files are based on epistemically rewarding relations to the reference
- (ii) We typically stand in several ER relations to the same object simultaneously

From (i) and (ii) it follows that several modes of presentation (mental files) are available for the same object. In identity judgments, and whenever we recognize an object *A* as the same as an object *B*, we deploy two such files simultaneously. So the argument goes.

In Section 5, I argued that recognition is not a state, involving the simultaneous deployment of two distinct files for the same object, but a transition from a (typically) two-file state⁹ to a one-file state. The two-file state is the pre-recognition state—when the subject believes or presupposes that there are two objects. This state comes to an end as soon as recognition occurs, or as soon as the subject becomes aware of the identity. The post-recognition state is the one-file state, in accordance with the merge model.

⁸ In quasi-merging, the inclusive file is not a distinct file but a more inclusive *stage* of the initial file.

⁹ I say 'typically' because I take *immediate* recognition to be a transition from a one-file state to a one-file state. Two cases were discussed: the case in which I immediately recognize the man cutting the grass as Chomsky, and the case in which I immediately recognize an object I see (e.g. the broken vase) as an object I remember from a previous perceptual encounter. In the former case, the permanent Chomsky file goes from a state in which it is 'detached' to a state in which it is 'attached', while remaining the same encyclopaedic file (Perry 2012: 86). In the latter case a memory file is transformed into a more inclusive file based on the acquaintance relation.

What about premises (i) and (ii)? Do they not entail that several modes of presentation (mental files) should be available for the same object? No, they don't, appearances notwithstanding. When we knowingly stand in several ER relations to the same object simultaneously, the mental file we deploy for the object is *itself based on several ER relations simultaneously*. Among ER relations, some are, as I said, 'composite' and involve a set of basic ER relations.¹⁰ I discussed the case of acquaintance: the acquaintance relation holds between the subject x and an object y just in case x remembers y from a previous perceptual encounter and is perceptually confronted with y once again. A mental file based on the acquaintance relation can host information from memory as well as information from perception (in a way that makes trading on identity possible). This is a case in which x stands in several ER relations to y simultaneously, but that does not entail that x , in that situation, deploys two distinct mental files (one based on memory, the other on perception). Such a two-file state would correspond to the situation in which the subject does *not* recognize the object she sees as the object she encountered before and remembers. If the subject recognizes the object, however, what she deploys is a *single* file based on *both* memory and perception (the two components of the acquaintance relation). All of this is perfectly consistent with Strawson's constraint.

What about the other two arguments against Strawson's constraint and the merge model? One of them had to do with the reversibility of linking, as opposed to merging. The other had to do with the fact that we often have to consider the point of view of 'unenlightened' people who are not aware of the identity: taking their point of view into consideration requires vicariously deploying the two coreferential files, something we could not do if, upon learning the identity, we had merged these files instead of keeping them distinct. I will call these arguments the reversibility argument and the metarepresentational argument, respectively. I start with the latter.

I fully agree that, when we attempt to describe another person's intentional states or actions, we often take that person's perspective in specifying the object toward which the state or action is directed. If we know that $A = B$ but they don't, then in reporting their thoughts or actions, we may vicariously deploy two distinct files, an A -file and a B -file,

¹⁰ In Hans Kamp's framework, a mental file based on a composite ER relation is an 'anchored entity representation' that has 'multiple anchors' (Kamp 2015: 286).

despite knowing that there is a single object. As an example, consider the following passage from Ruth Millikan's *Varieties of Meaning*:

When a kitten sees itself in the mirror, it tries to approach the other kitten, to smell it and touch it. Failing in this, it then tries to look behind the mirror. (Millikan 2004: 123)

What the kitten sees in the mirror is what it tries (but fails) to approach, to smell and to touch. In referring to it, however, Millikan adopts two different perspectives: first, from her own, external perspective on the scene, she describes what the kitten sees as 'itself'; then, in trying to make sense of the kitten's actions, she adopts the kitten's own perspective in describing what it is trying to approach, smell and touch: 'the other kitten'. This perspective shift is similar to that which occurs in the example from Pinillos (2011), which I quoted in Section 4:

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

Here also the speaker, in describing their attitudes, espouses the point of view of the unenlightened (herself and her peers before they learnt the identity): she refers to Venus via two distinct mental files (respectively associated with the names 'Hesperus' and 'Phosphorus'), despite knowing the identity. That means that the distinct Hesperus and Phosphorus files are still available and can be deployed for metarepresentational purposes. In the second clause, which expresses the subject's current point of view, the speaker deploys an inclusive Venus file (associated with the demonstrative adverb 'there' at the end of the second clause).

I am not denying any of this. On the contrary, I think we should draw a principled distinction between two types of mental file: regular files, corresponding to the subject's own point of view, and vicarious files, indexed to other people, which we deploy for metarepresentational purposes (Recanati 2012: 181–205, 2016: 35–9). The mental file associated with the description 'the other kitten' in Millikan's passage is *a vicarious file indexed to the kitten and representing the kitten's own way of thinking of the animal it tries*

to approach, to smell, and to touch. Likewise, the mental files associated with the names 'Hesperus' and 'Phosphorus' in Pinillos's example are vicarious files indexed to the unenlightened.

In this context we can formulate a suitably restricted version of the merge model which takes the metarepresentational argument into account. Strawson's constraint works for regular files. If we learn that $A = B$, then we merge our files into a single, inclusive file (e.g. the Venus file in the Hesperus/Phosphorus case). But that does not prevent us from keeping distinct *indexed* files, for metarepresentational purposes. We may even assume that when we merge two files, the two initial files do not disappear but automatically acquire the status of vicarious files indexed to our earlier self. So they disappear as regular files, as per the merge model, but remain available as indexed files, as per the metarepresentational argument.

This response to the metarepresentational argument protects the merge model by circumscribing its area of validity, but it additionally provides a response to the first argument: it kills two birds with one stone. Reversibility is now secured since the initial files are retained with the new status of indexed files, despite the merging operation, which leaves only one regular file in place: the inclusive file. If the identity turns out to have been mistaken, then to revert to the earlier state one only has to (i) suppress the inclusive file and (ii) restore the indexed files to their initial status as regular files.

On this version of the merge model, there is no linking of regular files. When an identity is discovered, the two files at stake are merged (while copies of the initial files are retained with the status of indexed files). So Strawson's constraint is respected. One form of 'linking' remains, however—what in *Mental Files* I call 'vertical' linking, that is, linking between a regular file and an indexed file (Recanati 2012: 193 ff.). For example, Millikan's vicarious file for 'the other kitten' is linked to her regular file about the kitten: she knows that what the kitten takes to be another kitten is *that kitten itself* seen in the mirror. She links her vicarious file 'the other kitten', indexed to the kitten, to her regular file about the kitten. Likewise, in Pinillos's example, the Hesperus and Phosphorus files are indexed to the unenlightened and linked to the speaker's regular file about Venus: the speaker knows that what the unenlightened take to be two distinct stars are actually a single planet.

Vertical linking is important because it secures a reference for indexed files. Thus we don't have to treat the description 'the other kitten' in Millikan's passage as referentially

vacuous (even though there is no other kitten). We can maintain that the description refers to what the kitten sees, namely itself, but through a vicarious mode of presentation, an indexed file, corresponding to the kitten's own perspective. That indexed file in Millikan's mind refers to what the regular file it is linked to itself refers to, namely the confused kitten. In other words, *an indexed file (which simulates in the subject's mind the way another subject thinks of an object) inherits the reference of the regular file it is linked to in the subject's mind*. In Pinillos's example, the Hesperus and Phosphorus files both refer to Venus because they are vertically linked to the speaker's regular Venus file. Of course, it is also possible for an indexed file not to be linked to any regular file in the subject's mind, but in that case the indexed file has a purely metarepresentational function and does not refer; it is a 'free-wheeling' indexed file (like the Vulcan file in: 'Leverrier thought that the discovery of Vulcan would make him famous').¹¹

Even though the restricted version of the merge model makes room for a form of linking between files, namely vertical linking, it is entirely compatible with the principle of informational segregation. While the other form of linking ('horizontal' linking between regular files) is supposed to allow information to flow freely between the linked files and to make trading on identity possible across files, vertical linking does not have this property: it *preserves* informational segregation. As I wrote in *Mental Files*,

¹¹ When an indexed file refers in virtue of its vertical link to a regular file in the subject's mind, I say there is *oblique reference* (Recanati 2016: 34–40; Recanati 2019: 69–75). There is oblique reference to some individual *x* whenever the mode of presentation through which the speaker refers to *x* is not the way the speaker herself thinks of *x* but the way some other person thinks of *x*. In such cases, to understand the utterance you need to understand that the speaker is, as it were, *echoing* some other person and referring to *x* as that other person would. That the referring expression is used echoically can be indicated by putting it within quotation marks, as in my old example (Recanati 1987: 63):

- (5) Hey, 'your sister' is coming over!

In this example, the speaker refers to Ann, who is not the addressee's sister (and who both the speaker and his addressee know not to be the addressee's sister) but who is believed to be the addressee's sister by a third party, James. The speaker is ironically mocking James by referring to Ann as the addressee's sister, thus echoing James's mistake. That is possible because James's mistake is common knowledge between the speaker and his addressee. In that situation, the speaker says *of Ann*, vicariously thought of as the addressee's sister, that she is coming over. The speaker succeeds in referring, even though he does so by deploying a vicarious file, indexed to James. This is possible because the indexed file in question is linked to a regular file (about Ann) in the speaker's mind. The regular file refers in virtue of the ER relation(s) it is based on, and the indexed file inherits the reference of the regular file to which it is linked.

Linking between regular files makes it possible for information to flow freely between the linked files. But indexed files are used to stand for some other subject's body of information about some object, and that function could not be served if, through linking with the subject's regular files, the indexed file was contaminated by the subject's own information about that object. ... So vertical linking between regular files and indexed files (or between indexed files with different degrees of embedding) preserves the informational encapsulation of files, which standard (horizontal) linking has the effect of suppressing. (Recanati 2012: 193)

7. Conclusion

I have provided a defence of Strawson's constraint, to the effect that two coreferential files should be merged. I criticized the idea that the merge requirement only applies to encyclopaedic files, not to indexical files, because indexical files are based on particular ER relations, and we can stand in several ER relations to the same object simultaneously. As I argued, the merge requirement also applies to indexical files. When we discover that we stand in two ER relations to the same object (e.g. when we realize that the woman we see in the street is the woman we are talking to on the phone, as in Mark Richard's famous phone booth example),¹² we automatically stop thinking of that object through two distinct mental files and instead think of it through a single mental file based on a composite ER relation. That is exactly what the merge model predicts. The linking model, widely accepted as a substitute for Strawson's merge model, rests on the mistaken construal of recognition as a state, when in fact it is a transition between states.

The only cases that are not subject to Strawson's constraint are those that involve indexed files. We deploy indexed files to track alien perspectives on some object. Each perspective is governed by Strawson's constraint, but the constraint does not apply across perspectives. That means that, even though there is only one object, we can simultaneously deploy several files referring to that object, provided there is at most one regular file among these multiple files. In addition to that unique regular file, we can simultaneously deploy as many indexed files, linked to our regular file, as there are alien perspectives we care to mentally simulate. In particular, I concede that an identity statement such as 'Hesperus is

¹² See Richard (1983).

Phosphorus' involves the simultaneous deployment of two distinct mental files referring to Venus, but I do not think this supports the linking model; for, as I have mentioned, these two files are vicarious files indexed to the unenlightened people whom the identity statement is meant to inform.¹³ The simultaneous deployment of these two files is therefore compatible with Strawson's constraint, properly understood.

This defence of Strawson's constraint, and of the segregation principle it is based on, suggests that the two types of case in which a contradiction may be undetectable for a rational subject have more in common than one might have thought. In Frege cases ('Hesperus is *F* but Phosphorus is not *F*'), a contradiction is undetectable because of a difference in mode of presentation. In fragmentation cases, a contradiction is undetectable because the conflicting pieces of information are stored in different compartments of the mind. The mental file framework, construed in the Strawsonian way, makes it possible to view Frege cases themselves as cases in which a contradiction escapes the notice of a rational subject because the conflicting pieces of information are stored in different compartments of the mind.

This result evidently calls for a more systematic comparison of the frameworks (the mental file framework and the standard fragmentation framework) and of the phenomena they were introduced to deal with. An apparent difference between the phenomena is that, as Andrea Onofri put it in his comments on a previous version of this paper:

In classic Frege cases, it seems entirely correct to say that the inconsistency is undetectable—undetectable without further empirical investigation. Things are different in Lewis's case [the Nassau Street example], where the inconsistency is undetected but detectable. Lewis *could* notice the inconsistency if he thought about the matter more carefully—indeed, he says that he *did* notice it after a while and that this led him to revise his beliefs.

That a contradiction is undetected yet detectable in a fragmentation case is what the passage from Davidson quoted earlier suggests. Davidson says that a 'boundary' has to be postulated between distinct territories of the mind whenever a rational subject holds 'obviously conflicting beliefs' (Davidson 1986: 91). It is because of the boundary and the

¹³ See the citation from Lockwood in footnote 4.

fragmentation it induces that the subject is unable to 'join up the dots', despite the obviousness of the contradiction:

The contradictory beliefs do not belong to the same territory; to erase the line between them would be to destroy one of the beliefs. (Davidson 1986: 92)

Were the contradiction not obvious, it might escape the notice of a rational subject even if the beliefs were stored in the same compartment of the mind. That is the case, according to Onofri, for the belief that Hesperus is *F* and the belief that Phosphorus is not *F*. That these beliefs conflict and yield a contradiction is not obvious to the subject because of the difference in mode of presentation, so this sort of case does not generate a need for fragmentation.

But I have argued that Frege cases are also a matter of boundaries. This is so even if boundaries are understood along Davidsonian lines. The contradiction in a Frege case *is* obvious at a certain level: the predicates *F* and not-*F*, which the rational subject ascribes to the same object, *are* obviously contradictory, and it is only because they are stored in different files (the Hesperus file and the Phosphorus file) that the subject ends up predicating both of Venus. Both in Frege cases and in standard fragmentation cases, what makes the contradiction undetectable is the fact that the conflicting bits of information are stored in different places and remain insulated from each other.

Still, there is a difference between the two types of case, as Onofri correctly points out. That the predicates *F* and not-*F* are stored in different files affects the *content* of the subject's respective beliefs. The subject believes that Hesperus is *F* or she believes that Phosphorus is not-*F*—that is not the same thing. Because of this difference at the level of content, there is no *obvious* contradiction between believing that Hesperus is *F* and believing that Phosphorus is not-*F*. As a result, we do not need a boundary between these objectively contradictory beliefs, in contrast to David Lewis's Nassau Street case, where the contradiction is obvious and would be detected if the beliefs belonged to the same fragment. But this difference is compatible with the idea that a kind of fragmentation is also at work in the Hesperus/Phosphorus case, provided this kind of fragmentation takes place at a prior level and, in contrast to standard fragmentation, is partly constitutive of belief content.

In the standard fragmentation framework, issues of content are one thing, issues of accessibility (where the information is stored) are another. The mental file framework blurs that division. In the mental file framework, the place (the file) where a predicate is stored affects the content of the belief involving that predicate by determining both what the predicate is predicated of (the reference of the file) and how the object in question is thought of. In the Hesperus/Phosphorus case, the reference is the same because the two files corefer, but the modes of presentation (the files themselves) are different, and this is what makes the contradiction non-obvious at the level of belief content. Since, in the fragmentation framework, the information that is stored in different compartments of the mind directly corresponds to the content of the attitudes, where the information is stored has no similar impact on the content of the attitudes in question (though it may have an impact on the attitudinal 'mode').

References

- Bach, K. (1994), *Thought and Reference*, (2nd edn, Clarendon Press).
- Campbell, J. (1987), 'Is Sense Transparent?', in *Proceedings of the Aristotelian Society* 88: 273–92.
- Crimmins, M., and Perry, J. (1989/1993), 'The Prince and the Phone Booth: Reporting Puzzling Beliefs', in J. Perry, *The Problem of the Essential Indexical and Other Essays* (Oxford University Press), 249–78.
- Davies, M., and Egan, A. (2013), 'Delusion: Cognitive Approaches — Bayesian Inference and Compartmentalization', in K. Fulford et al. (eds), *The Oxford Handbook of Philosophy and Psychiatry* (Oxford University Press), 689–727.
- Davidson, D. (1986), 'Deception and Division', in J. Elster (ed.), *The Multiple Self* (Cambridge University Press), 79–92.
- Donnellan, K. (1966/1977), 'Reference and Definite Descriptions', in S. Schwartz (ed.), *Naming, Necessity, and Natural Kinds* (Cornell University Press), 42–65.
- Egan, A. (2008), 'Seeing and Believing: Perception, Belief Formation and the Divided Mind', in *Philosophical Studies* 140: 47–63.
- Evans, G. (1982), *The Varieties of Reference*, ed. J. McDowell (Clarendon Press).
- Fiengo, R., and May, R. (1994), *Indices and Identity* (MIT Press/Bradford Books).

- Fine, K. (2007), *Semantic Relationism* (Blackwell).
- Fodor, J. (1995), *The Elm and the Expert* (MIT Press/Bradford Books).
- Forbes, G. (1990), 'The Indispensability of Sinn', in *Philosophical Review* 99: 535–63.
- Goodale, M., and Milner, D. (2006), 'One Brain — Two Visual Systems', in *The Psychologist* 19: 660–63.
- Kamp, H. (2015), 'Using Proper Names as Intermediaries between Labelled Entity Representations', in *Erkenntnis* 80: 263–312.
- Kripke, S. (1979), 'A Puzzle About Belief', in A. Margalit (ed.), *Meaning and Use* (Reidel), 239–83.
- Lewis, D. (1982), 'Logic for Equivocators', in *Noûs* 16: 431–41.
- Lockwood, M. (1971), 'Identity and Reference', in M. Munitz (ed.), *Identity and Individuation* (New York University Press), 199–211.
- Millican, P. (1990), 'Contents, Thoughts and Definite Descriptions', in *Proceedings of the Aristotelian Society* supp. vol. 64: 167–203.
- Millikan, R. (1997), 'Images of Identity', in *Mind* 106: 499–519.
- (2000), *On Clear and Confused Ideas* (Cambridge University Press).
- (2004), *Varieties of Meaning* (MIT Press/Bradford Books).
- Perry, J. (1993), *The Problem of the Essential Indexical and Other Essays* (Oxford University Press).
- (2002), *Identity, Personal Identity, and the Self* (Hackett).
- (2012), 'Thinking about the Self', in J. Liu and J. Perry (eds.), *Consciousness and the Self: New Essays* (Cambridge University Press), 76–100.
- Pinillos, A. (2011), 'Coreference and Meaning', in *Philosophical Studies* 154: 301–24.
- Recanati, F. (1987), 'Contextual Dependence and Definite Descriptions', in *Proceedings of the Aristotelian Society* 87: 57–73.
- (2012), *Mental Files* (Oxford University Press).
- (2013), 'Perceptual Concepts: In Defence of the Indexical Model', in *Synthese* 190: 1841–55.
- (2016), *Mental Files in Flux* (Oxford University Press).
- (2019), 'Modes of Presentation in Attitude Reports', in A. Sullivan (ed.), *Sensations, Thoughts, Language: Essays in Honor of Brian Loar* (Routledge), 54–77.

- Richard, M. (1983), 'Direct Reference and Ascriptions of Belief', in *Journal of Philosophical Logic* 12: 425–52.
- Schiffer, S. (1978), 'The Basis of Reference', in *Erkenntnis* 13: 171–206.
- Schwitzgebel, E. (2002), 'Self-Ignorance', in J. Liu and J. Perry (eds.), *Consciousness and the Self: New Essays* (Cambridge University Press), 184–97.
- Stalnaker, R. (1999), *Context and Content* (Oxford University Press).
- Strawson, P. (1974), *Subject and Predicate in Logic and Grammar* (Methuen).