

Against the Mental-files Conception of Singular Thought*

(Published in *The Review of Philosophy and Psychology*. Please cite the published version.)

In this paper, I argue against a conception of singular thought that has become dominant in the current literature, which I call the mental-files conception of singular thought (MFC, for short). According to proponents of the MFC, we can theorise the notion of singular thought by appeal to the psychological notion of a mental-file because, to paraphrase terms employed by Recanati (2012), to have and use a mental-file is to have and use a singular object concept.¹ In contrast, my view is that the notion of a mental-file and the notion of a singular thought ought not to be identified. Each is an important notion for the philosophy of mind, but the two notions play different theoretical roles, and they are not co-extensive.

Proponents of the MFC hold that one thinks a singular thought about an object *o* iff one's thought employs a mental-file on *o* [Jeshion (2010), Recanati (2012)]. I will show this is false by showing that file-thinking is not always singular. This is because there are what I will call *descriptive mental-files*: mental-files for which a description plays certain governing roles to be spelled out in the paper. Certain central features of mental-files have led theorists of files and singular thought to assume that the idea of a descriptive file is incoherent. But I argue that the possibility of a descriptive file is consistent with the central features of file-hood, and thus show that file-thinking ought not to be conceptually equated with singular thought. I then apply the notion of a descriptive file to some cases, and thus show that descriptive file-based thoughts are not only possible, but actual: some file-based thoughts are not singular thoughts.

After establishing that file-based descriptive thoughts are both possible and actual, however, I note that the notion of a descriptive file leaves us with several questions to address. First, descriptive files are subject to certain limits, which do not apply to non-descriptive files. Second, complications arise when we apply the notion of a descriptive file to certain particular cases (which appear nonetheless not to fit the model of a singular file). To address these issues, I introduce a distinction between a *fully descriptive file* and a *holistically descriptive file*, which helps us to better understand both the nature of files in general, and the possibilities for descriptivism about file-based thought.

I want to be clear that the upshot of my argument is not to discount the importance of the notion of a mental-file for the philosophy of mind. Instead, I claim that, by rejecting the MFC—and thereby distinguishing the notion of singular thought from that of file-

* Thanks to Imogen Dickie, Aidan Gray, Simon Prosser and Robbie Williams for discussion of this and related material. Thanks also to Jeff King for comments on an earlier version of the paper, to an audience at the Leeds Philosophy Department Senior Seminar for their questions, and to two anonymous referees at *Review of Philosophy and Psychology* for helpful comments.

¹ Recanati (2012), 34, writes, 'A non-descriptive mode of presentation, I claim, is nothing but a mental-file'. 1

thinking—we can better understand the nature of mental-files *and* the nature of singular thought.

I Singular Thought

The claim I wish to argue against is that thinking a *singular thought* is to be identified with *thinking with a mental-file*. I'll do this by arguing that the bi-conditional endorsed by MFC proponents—that one thinks a singular thought about an object *o* iff one thinks with a mental-file on *o*—is false. But, to assess this bi-conditional, we need to start with a working definition of singular thought.

Until recently, it has been most common to define singular thoughts in terms of their content: as mental states with singular (as opposed to general) content.² Two distinct contents can both be true or false at the actual world depending on whether the very same object does or does not possess the very the same property if one of those contents is singular, and the other is descriptive. (1) and (2) are both true at the actual world iff Barack Obama was born in America, but their semantic content is different:

- 1) $\exists x$ [44th President (x) & $(\forall y)$ (44th President (y) \rightarrow ($x = y$)) & born in America (x)]
- 2) born in America (*Barack Obama*)³

(1) picks out Obama for predication by laying down a descriptive condition that he satisfies—it picks out Obama by generalizing over objects. (2) picks out Obama because it includes a terms which is an individual constant that contributes Obama to the truth conditions of the whole content—in this sense, the content itself constitutively involves Obama, the individual. One way to define singular thoughts is as mental states that have content like (2), rather than like (1).^{4 5}

There are difficulties with using this traditional definition to assess the credentials of the MFC, however. Firstly, it entails that singular thoughts are object-dependent—a claim that some MFC proponents deny [Jeshion (2002), (2010)].⁶ In arguing against the MFC, it will be

² This definition entails that singular thoughts are object-dependent. The definition has lost popularity in the recent literature in part for this reason, since it has become common to dispute this claim. See, for example, [Sainsbury (2005)].

³ Let's assume that the italicized name here plays the role of an individual constant in the formalisation of this content. Its semantic role is to introduce an object for predication.

⁴ Or indeed like: $(\forall y)$ [(American President (y) \rightarrow born in America (y))], which is also general. (1) can be thought of as a limiting case of general content.

⁵ One can give this definition of singular thought but remain neutral as to whether the metaphysics of singular content is such that it contains an object itself as a constituent, or an object-dependent concept. In other words, both Russellians and Fregeans about singular content can adopt this definition.

⁶ The extent to which Jeshion rejects the definition of a singular thought as a mental states with singular content is unclear, even though this definition sits uneasily with her account. She is committed to denying that singular thoughts are object-dependent insofar as she allows that empty file-based thoughts are singular (see for example her 'Vulcan' example in [Jeshion (2002) & (2010)]). On the other hand, she does *tentatively* endorse the claim that singular thoughts are mental states with singular content, along with an object-dependent conception

preferable not to assess the position against a definition of singular thought that is controversial, or that is rejected by some MFC proponents. Secondly, even if all parties accepted the above definition, it would not serve as a definition that enables us to assess the claim that singular thought is file-based thought. This is because it is too abstract to give us traction when asking if particular cases of file-based thought are singular or descriptive.⁷ When we assign mental states with content, we are using abstract objects, like the contents expressed by (1) and (2), to map the representational or intentional properties of those states. If the only grip we have on the notion of singular thought is that singular thoughts are mental states to which we assign contents like (2) rather than (1), for any given case where *I* attempt to claim that a given file-based thought is descriptive, rather than singular, the *MFC proponent* can simply disagree, and it will be unclear how we ought to adjudicate the dispute.⁸

To address both these difficulties, what we need is a working definition of singular thought that is relatively neutral—at least, one to which MFC proponents would agree—and which has built into it some sense of what features of a mental state a proponent of the traditional definition would be trying to *capture* by assigning it with a content like (2), rather than (1), or vice versa.

A working definition meeting these desiderata can be built on the following idea. The notion that a thought has descriptive content is used to capture the idea that it involves employing a set of general concepts, arranged in a certain way such as to net an object or objects, which satisfy the condition laid out by those concepts.⁹ Therefore descriptive thoughts involve thinking about an object *via its properties*, or *merely as the possessor of certain properties*. Singular thought can then be defined negatively by contrast: thinking a singular thought involves thinking of an object, but *not* merely as the possessor of properties *x, y, z*, or *not* via its properties, but in some *other* way. We can call this way, *non-satisfactorily*.¹⁰

Our question about the MFC now becomes, ‘*Is it true that file-based thinking is the same thing as thinking of an object non-satisfactorily?*’. That is, when a thought employs a mental-file, does this mean it is a case of thinking of an object, not *via its properties* or *merely as the possessor of certain properties*, but in some other way?

of singular content in [Jeshion (2010), 108]: ‘The content of singular thoughts are singular propositions, containing as their constituents individuals and properties.’ Therefore, although she explicitly but tentatively accepts the definition of singular thought as thoughts with singular content, the definition is inconsistent with her view.

⁷ A similar point about the limitations of defining singular thoughts by their content is made by Jeshion (2010). Her point here is that agreeing on the idea that singular thoughts have singular content does not settle questions about what it takes, cognitively, epistemically, etc., to *think* such thoughts.

⁸ The suggestion that we can use our everyday practices of ascription to settle this debate is mistaken because ordinary purposes attitude ascriptions do not systematically track the content of the attitudes they are used to ascribe. In Goodman (unpublished. a), I argue that the truth conditions of ordinary-purposes attitude ascriptions underdetermines even *the truth-conditional content* of the mental states they are used to ascribe. For other discussions of the way that ordinary ascriptions fail to track the content of the mental states they are used to attribute see [Recanati (2012) 150-54], [Bach (1987), (1997)] & [Taylor (2002)].

⁹ Jeshion (2010), 108, agrees: ‘For descriptive thought there is widespread agreement [about what it takes to entertain such thoughts]: One must possess and grasp those constituent concepts in the general proposition, must do so in the way in which they are structured in the proposition.’ See also, [Jeshion (2010), 129].

¹⁰ This terminology is derived from Bach’s distinction between thoughts whose objects are determined *satisfactorily* vs. *relationally* (see [Bach (1987)]). Bach is also an MFC proponent.

II The MFC

The proponent of the MFC answers this question affirmatively. She holds the following two commitments. Firstly, when a mental state is singular or non-satisfactional, what this means *in psychologically real terms* is that the mental state is a case of thinking with a mental-file. Secondly, the notion of a mental-file allows us to understand the *difference* between singular and descriptive thought. Singular thoughts involve the use of mental-files, whereas descriptive thoughts don't. Thus, she is committed to the central claim of the MFC: one thinks a singular thought iff one's thought employs a mental-file.

A mental-file is essentially a cluster of predicates believed to be co-instantiated, which are stored together, and used together, in a mental economy.¹¹ Part of their importance in the *philosophical* literature comes out of their usefulness in helping us to understand both the phenomenon of continued belief, and also the streamlining of inferences and automatic updating of belief that occurs in certain cases.¹² ¹³ They can play these explanatory roles because they are taken by file theorists to individuate cognitive perspectives on objects, or to play the role of *object-concepts*.

For an example of the way that files help us to understand the automatic updating of beliefs, take a case in which the guy who lives next door to you also works in an office across the hall from your office, and you know this. If you store all your information about this person in a single mental-file, you will be able to presuppose sameness of the identity of this individual in a way that allows you to move efficiently and without the need for further justification from a belief formed at the office, which might be expressed with,

3) The person who works in the office across the hall is a psychologist,

to a claim you might make at home, to your family that,

4) The person who lives next door is a psychologist,

, *without* having to appeal to the identity premise—that is, the *extra step*:

5) The person who lives next door = the person who works in the office across the hall.¹⁴

¹¹ This fits with Recanati's claim that files 'are a matter of information clustering' ([Recanati (2012), 42]).

¹² Again, this is in line with Recanati's recent conception of files according to which the existence of a file licenses the integration and exploitation in automatic inference patterns of information stored in that file (See, for e.g., [Recanati (2012), 41, 96]).

¹³ They have also been used to understand the psychology of name-use. See, e.g. [Grice (1969)], [Lockwood (1971)].

¹⁴ It is important that, in spelling out this example, I intend no commitment to the idea that (3) & (4) state the *contents* of beliefs or thoughts. (3) and (4) are rather *statements* that might be used to *express* the beliefs or thoughts involved. According to proponents of the mental-files framework, the thought one has at the office about the individual whose office is across the hall and the thought one has at home about the individual who lives next door, both involve the use of the *same* mental file (as well as predication of the same property). On 4

The existence of a mental-file as part of your psychological economy explains the fact that, on receiving the information stated by (3), you automatically (and in a justified way) update your beliefs concerning the person who lives next door to you. This is often stated by saying that mental-files allow thinkers to ‘trade on identity’.¹⁵

Furthermore, the notion of a mental-file gives us a way of understanding how, when the psychologist across the hall tells you a funny joke on *Monday*, and you form a belief that you might express to your colleagues by saying ‘the psychologist across the hall is a funny guy’, your utterance at home to your family on *Wednesday* of ‘the guy next door is a funny guy’ could be an expression of a continued belief—an expression of the same belief you formed on Monday.¹⁶

The role of files in explaining both of the above phenomena goes hand in hand with the basic role they are claimed to play as dynamic object-concepts, or cognitive perspectives on objects.¹⁷ Sameness of file explains continued belief because, in a case of continued belief from t^1 to t^2 , the very same file is used to think that the object is, say, \emptyset . This is what makes it the case that changes in context or descriptive information (say, the change from thinking of the object as *today* to thinking of it as *yesterday*) do not interfere with the internal continuity of the putative belief. Use of the same file—the same object concept—is what persists through these changes and captures the sense in which a single belief is maintained.

Likewise, possession of a file explains justified *trading on identity* because the inference from, say, ‘*this* is F’ and ‘*this* is G’ to ‘something is both F and G’, employs the same object-concept throughout (which in turn is grounded in continued, successful tracking of the object, or a continued perspective on the object). This allows for the ‘assumption’ that *this=this*, without further justification or reasoning. Thinking of files as object concepts in this way also shows why they are useful for theorizing the presence or absence of Frege

the mental-files framework, this means that these two thoughts have the *same* content. As I explain further below, this is because, according to file-theorists, a mental file *is* an object-concept, and the two thoughts employ the same object-concept. Thanks to an anonymous reviewer for requesting that this point be clarified. ¹⁵ For more discussion of what this means, see Goodsell (2013). The sense of ‘trade on identity’ I use here is (roughly) like Goodsell’s ‘current-reasoning reading’, according to which one trades on the identity of a and b if one is disposed to reason as if $a=b$ without using an additional identity premise in one’s reasoning. I don’t mean to suggest, by appealing to this role for mental-files, that we have at our disposal a fully worked-out sense of what it means to reason without an extra premise, or an adequate understanding of *when* this occurs. This is work that must be done in order to fill out the theory of mental-files. What does seem clear, however, is that there is that reasoning of this kind must occur for, without it, it is unclear how justified inference would be possible (see [Campbell (1987)]).

¹⁶ See [Perry (2000), 69]. It is important to note the following: Although in some of Perry’s examples there is no continuing descriptive information available to the subject, which could be used to specify the putative (continued) belief, this does not entail that it is *only* in cases in which there is no continuing descriptive information that files can be used to theorise continuing belief. Perry’s claim is that continuation of either descriptive information (or in his terms ‘text’) is not a necessary or sufficient condition on continued belief (Ibid., 81). This is consistent with the claim that continuity of a file accounts for continuity of belief in cases where there *is* continuing descriptive information available, or a continuity of ‘text’. This point is often overlooked, and is in fact related to my later claim that it is a mistake to think that the file framework is *only* suited to theorising cases of non-descriptive thought.

¹⁷ This is sometimes put by saying that files play the role of modes of presentation (See, [Recanati (2012), Ch. 3])

cases. A case where it is *informative* to be told that *this=that* is theorized as a case where the subject is in possession of two (co-referential) mental-files. A case where this is *not* informative is one where *sameness of file* explains the transparency of the identity fact. To *learn* that *this=that*, and integrate this knowledge into one's cognitive perspective, is to *merge* one's two files into one.¹⁸

But, what do mental-files have to do with *singular thought*? What motivates the identification of file-based thought and singular thought?

Files are a way of keeping track of particular things, either within or across contexts. As I explain above, this invites the reasonable suggestion, made for example by Recanati (2012), that a mental-file itself is the mode of presentation by which one thinks of a particular object. To step away from the Fregean terminology of modes of presentation, it invites the reasonable suggestion that the file one keeps on an object is one's *object-concept*. But why think files are *non-descriptive* object concepts?

Broadly speaking, a range of the features of mental-files have led naturally to the assumption that file-based thoughts must be singular thoughts.

Firstly, files are not individuated simply by the information contained or stored in them. My file on an object persists through new information being introduced into the file, and information being discarded from it. In some extreme cases, a complete overhaul of the file's contents may even occur—'ship of Theseus'-style—while the same file persists. Any theory of file-thinking that conceived of the content of a file-based thought in terms of, say, a conjunction of the descriptive information stored in the file, would not accommodate this feature. It would also deny file-thinking the role of accounting for continued belief. For example, on a view on which the content of a file-based thought was determined by the descriptive information stored in the file, files could not accommodate any sense of a *single* continued belief in any case where the file's contents have changed over time.¹⁹

Secondly, and perhaps most centrally for the MFC proponent, the content of a file-based thought is claimed to be singular because files can contain false information, and this does not affect their referential success. It is usually emphasised that a file can be a file on an object, *a*, even if the file contains predications that are not true of *a*.²⁰ Relatedly, the sum of the information stored in a file may not add up to a descriptive condition that picks out the

¹⁸ This is somewhat complicated by the possibility (contra Strawson) of *knowingly* keeping two distinct but co-referential files. There are several reasons why this might occur (without irrationality): for example, relatively low credence in the identity fact (which nonetheless rises to the level required by knowledge), or it being useful to keep track of the perspectives of others (or one's former self) on the object [Recanati (2012), 44-45]. A natural way to theorise such cases is in terms of files that are linked, but not merged. There are several interesting issues attached to the possibility of such cases, but I set them aside here.

¹⁹ That my file on an object *o* loses or gains a predicate 'is F' from *t*₁ to *t*₂ does not preclude that the file accounts for my continuing to believe, from *t*₁ to *t*₂, that *o* is *G*. However, in cases of *complete* overhaul of descriptive information stored in a file, that file's accounting for continued belief will be foreclosed. It is clearly a constraint on a file's being used to theorise continuation of the belief that *o* is *G* that 'is *G*' remains in the file for the duration of the belief (otherwise there is no continued commitment to *o* being *G*). In this sense, 'Ship of Theseus' cases are marginal, in that they are cases in which a file cannot play *one* of its theoretical roles. In such cases, the file can account for trading on identity in inference, and continuation of object-concept, but not for continued belief.

²⁰ It is arguable, but not obviously true, that a file could be a file on *a* even if it contained *no predicates true of a*. 6

object of the file uniquely. It is *taken* to be a condition on these possibilities that the file does not pick out its object in virtue of the fact that the object *satisfies* the predicates contained in it, but rather *non-satisfactionally*. This suggests (in accordance with our criterion for the distinction between singular and general thought outlined in §I) that files, and therefore the content of the thoughts that employ them, are non-descriptive. In other words, there is a version of Kripke’s ‘semantic argument’ that suggests that the meaning of a file, like the meaning of a name, is not descriptive [Kripke (1980)].

Finally, the idea that files are *singular* object-concepts feeds off the association of mental files with cases in which one has a proper name for an object.²¹ If mental files give us an understanding of the psychology of thinking with a name, and if proper names are singular referential devices of natural language, then it can seem as if mental files play the role of spelling out the psychology of direct reference. However, the connection between name-use and mental-files should not, even at first glance, yield commitment to the claim that all file-based thoughts are singular. Even if the use of a mental-file is a feature of the psychology of name-use, this does not entail that, in all cases of file-based thoughts, the thinker has a proper name for the file’s object. Even allowing that all name-use goes along with file-use, this leaves open that some file-use doesn’t go along with name-use (indeed, few would deny this). Therefore, even granting that all thoughts associated with name-use are singular—a claim I would in fact dispute—some file-based thought might still be non-singular.

III Descriptive Files

I will argue that the MFC proponent is right about the central features of mental-files, but wrong to think these features entail that all mental-files are singular, or non-satisfactional. In fact, file-based thought can be descriptive thought—it can be thought in which the object of the file is thought of *merely as the possessor of certain properties, or via its properties*. This possibility is consistent with the features of mental-files that mark them *as* files. I will show this by introducing the notion of a *descriptive file*: a file that is *satisfactional*, despite bearing the marks of file-hood.

Before doing this, let’s review the central features of mental-files, so we can check the notion of a descriptive file against them. A mental-file is a cluster of predicates, believed to be co-instantiated, which is stored together and used together in thought. They have features (F1)-(F5):²²

²¹ See [Evans (1985)] for an early example, and [Jeshion (2009) (2010)] for more recent examples.

²² A note on two absences from this list: First, the connection between files and name-use does not register as an item on the list of the features of files. This is for a reason discussed in the final paragraph of §II. Even those who emphasize the connection between files and name-use do not argue that *all* files are associated with names. Second, the reader may wonder whether the idea that files involve *rigid* reference to objects ought to be included as a feature of mental-files (Thanks to Heather Logue for raising this possibility). This, one might suggest, is another reason to think that file-based thoughts are singular. My position is both that files do *not* always refer rigidly to their objects, and also that rigidity clearly does not entail singularity. Although my argument in this paper does not focus on the notion of rigidity, I argue that files don’t always refer rigidly in [Goodman (unpublished. b)] where I give a case of a non-rigid file.

F1) Files are used to store information about their objects and, thereby, to keep track of their objects over time.

F2) A file is a cognitive perspective on an object, or an object-concept. The facts about sameness of file, or distinctness of files, account for the transparency, or informativeness, of certain identity judgments.

F3) The presence of a mental-file is marked by certain patterns of automatic inference and updating that involve the thinker's ability and disposition to 'trade on identity' (in the sense described in §II).

F4) A file is not individuated simply by the information stored in it: the same file can persist through changes to information stored in it.

F5) The semantic content of a file-based thought cannot be identified simply with the descriptive information stored in the file, because a file can contain false information about an object *o* whilst still referring to *o*. It can also be used to theorise continued or persisting beliefs in cases where the information content of the file has changed.

That all files have these characteristics does not entail that all file-based thoughts are singular. Despite the apparent connection between files and singular thought—articulated especially by (F4) & (F5)—the notion of a *descriptive file* is consistent with the file-theoretic framework. A descriptive file is a file that is *satisfactional*. Employing the file involves thinking of its object *via its properties*, or *merely as the possessor of certain properties*. A descriptive file is a file that meets the criteria (F1)-(F5) for file-hood, but is *governed* by a description.²³ This means that there is a description that plays a privileged role with respect to the file. The description plays roles (G1)-G4):

G1) It determines which object the file is about, if any.

²³ The particular notion of *governing* I spell out here is partly inspired by points stressed by Dickie (2011), in giving a theory of how proper names refer, although my conclusions are not the same as hers. In Dickie (2011), a distinction is made between cases in which mistaken descriptive information (in particular, mistaken kind, or sortal, information) interfere with the referential properties of a name, and cases in which such mistakes do not. My notion of *governing* spells out a relation that a piece of information can bear to a file containing it, such that it does determine the referential features of the file along with other central features of the file (to be spelled out in what follows). Dickie's cases in which a piece of descriptive information is mistaken but this does *not* interfere with the referential properties of the containing file are cases in which that information does *not* govern the file. A note on terminology: Despite the fact that my notion of *governing* is inspired by Dickie's cases, it is very different from, and therefore not to be confused with, the notion of *governance* used by Dickie herself (see, again [Dickie (2011)]). My notion of *governing* spells out a (privileged) relation that a piece of descriptive information can bear to a file in which it is stored. Dickie's notion of *governance* is a relation that an object can bear to a file, or body of beliefs.

G2) It sets the limits on possible mistakes that fall within the scope of successful reference for the file.²⁴

G3) It acts as a ‘gatekeeper’ to the file.

G4) It determines the persistence conditions of the file.

Let me explain these roles in turn.

(G1) is the most straightforward. For a description, say ‘the F ’, to determine which object a file is about is for it to play the role of determining that the file’s object is the F , whichever thing *that* is. If there is no unique F , the file has no object.²⁵ Note that proponents of the MFC largely agree that a description can play role (G1) insofar as they allow for cases of files for which the ‘referent is fixed’ by description.²⁶ A way of stating the difference between my view and theirs is that they claim that the role a description can play with respect to a mental file *begins and ends* with the file’s *meta-semantics*. In a case of descriptive reference-fixing for a mental file, the MFC proponent holds that the description determines an object as the referent of the file but that, in all other respects, the file behaves non-satisfactionally (the description does not play any roles that make it part of the semantics for the file). This is where I disagree. My claim is that in *some* cases where the file’s object is determined by description, the description also plays *other* roles (i.e. (G2)-(G4)), which make it right to say 1) that the file is a case of thinking *satisfactionally* and therefore 2) that the description plays not merely a meta-semantic role but a semantic role.

(G2) is that the description sets the limits on possible mistakes that fall within the scope of successful reference for the file. This can be understood by thinking about what (F5) above requires of a file, and what it does *not* require. (F5) says that the semantic content of a file-based thought cannot simply be the descriptive information stored in the file, because part of what is distinctive about files is that they can contain false information about their objects whilst still referring to those objects. The fact that an object o is not G does not preclude the predicate ‘is G ’ from being stored in a file that succeeds in picking out o . A file thereby allows for mistakes that fall within the scope of successful reference. That some such errors are possible without defeating referential success is a general requirement on file-hood. However, when a description *governs* a file, it limits these possibilities without violating the general requirement. If ‘the F ’ *governs* a file, (G2) states that the file’s thinker cannot

²⁴ When I use the term ‘reference’ in relation to mental-files, I do not use it in the ‘strong’ sense (sometimes contrasted with *denotation*) that implies *singular* reference.

²⁵ I’m going to largely set aside questions about the status of thoughts that employ empty files for the purposes of this initial exploration of descriptive files. However, we can work with the assumption that these thoughts are like other descriptive thoughts: in the case where there is no unique F , descriptive thoughts that employ a file governed by the description ‘the F ’ are false.

²⁶ There is disagreement among MFC proponents about the constraints on descriptive reference fixing for files. Jeshion (2010) allows for descriptive reference fixing without acquaintance as long as her *significance condition* is fulfilled. Recanati (2012) has a more complex position. He thinks that acquaintance is a normative requirement on opening a file but, simplifying somewhat, allows for descriptive reference-fixing on the condition of either expected or ‘imagined’ acquaintance (see [Recanati (2012), 167 & 168], respectively).

mistakenly think that the file's object is F. That the file's object is F is a condition on the file being a file on that object.

(G3), the 'gatekeeper' role, is an epistemic-cum-justificatory role that a governing description plays with respect to a descriptive file. Since a file is a cluster of predicates, believed to be co-instantiated, which is maintained over time and used to store information about an object, it is a condition on the rationality of every file that it must be associated with an 'information marshaling strategy':²⁷ a strategy for determining which information will be stored in the file. Otherwise, the thinker maintaining the file would not be justified in presuming that information stored in the file concerned a single object, and the file would not purport to be in good standing as a cognitive perspective on an object or object-concept. For example, for a perceptual demonstrative file, which is based on an attentional, perceptual connection to an object, a perceptual link with the file's object is the file's information marshaling strategy. A piece of information enters the file, and is justified in entering the file, only when it is acquired through the perceptual link associated with the file.²⁸ A way of thinking about this is that the perceptual connection acts as the 'gatekeeper' to the file in the case of a perception-based file. In contrast, in the case of a *descriptive file*, the file is not associated with a perceptual, informational connection to its referent, so no such connection can play the role of determining *which* information is to be stored in the file. But something must play this role. In a file governed by a description, the governing description plays this role: if the description 'the F' governs a file, 'is G' can only (rationally) be stored in the file when the thinker has some reason to believe that 'is F' and 'is G' are co-instantiated.

(G4), that a governing description determines persistence conditions for the file, will be the most controversial role. However, it follows from the theoretical role of files as accounting for sameness of thought content across time, along with the claim that the semantic content of thoughts employing a descriptive file is specified by the file's governing description.²⁹ I will outline (G4) here, but defer discussion its implications for later in the paper.

We saw above (see (F4)) that it is a condition on file-hood that a mental-file can persist through changes to the information in it, and is thus not *individuated* by the information stored in it. Arguably, in cases where none of the descriptive information stored in a file governs the file—that is, in the case of a singular file (e.g., a perceptual demonstrative file)—the information in the file creates *no* limits on the persistence conditions of the file: a complete overhaul is possible while the same file continues to exist. Since no particular piece of information plays a governing role, *any* piece of information could be discarded from the file, and the thinker could go on thinking thoughts with the same semantic content by employing the same file. However, this is not the case with a descriptive file. A descriptive file cannot continue to exist as the same file if its governing description is discarded. This

²⁷ I take this helpful terminology from Dickie (2011) & (MS).

²⁸ The perceptual connection associated with a perception-based file is an instance of what Recanati (2012) calls an 'epistemically rewarding relation'.

²⁹ That is, abstracting away from the predicative part of their content: the semantic content of a thought employing a descriptive file governed by the description 'the F' is a descriptive content: 'The F is ...'

follows from 1) the commitment that thoughts employing the same file have the same semantic content and 2) the fact that a descriptive file's governing description specifies the semantic content of thoughts employing it.³⁰

A descriptive file governed by 'the *F*' could not discard and replace the predicate 'the *F*' (while still remaining the same file) because this particular information update will entail a difference of semantic content of the associated thoughts and beliefs. Say I employ a descriptive mental file governed by 'the *F*' in thinking a series of thoughts (at t_1 , t_2 and t_3) to the effect that the *F* is *G*. The content of these thoughts is:

6) The *F* is *G*

Now, imagine that, at t_4 , 'the *F*' is discarded from the file. From t_4 onward, to the extent that I am able to continue to use the same information cluster—so, the cluster now *missing*, 'the *F*'—the thoughts that result will *not* have the content specified by (6).³¹ This is because they are not thoughts that employ a file governed by the description 'the *F*' (after all, 'the *F*' has been *discarded*). To the extent that it is part of the file theoretic picture that sameness of file entails sameness of content (because it entails sameness of object-concept), we must theorise the thoughts after T_4 as ones that employ a *different* mental-file.

With a basic grip on roles (G1)-(G4) in place, here is a schematic description of the kind of case I have in mind, with an explanation of how a description plays the roles that add up to it *governing* a file.

A file is formed to keep track of an object, which is known at the time the file is opened only as 'the *F*'. The file's object, if it has one, is the *F* (whichever thing that is). The file's object is determined by description (G1). Over time, information about the *F* is stored in the file. Some of the information may be false of the file's referent: the fact that the *F* is in fact *not G* does not preclude the predicate, 'is *G*' from being stored in the file. If the file's semantic content were identical with the descriptive information stored in it, this mistake would not be possible. However, given that the description, 'the *F*' determines which object the file is about, it is *not* possible that the thinker is mistaken that the object is the *F*. Thus, the description sets limits on possible *misinformation* in the file (G2). Since the file is not associated with, say, a perceptual connection to its object, no informational connection can play the role of determining *which* information is to be stored in the file.³² Instead, the description forms the basis for the information marshaling strategy associated with the file: a

³⁰ Again, I abstract away from the predicative part of the content here.

³¹ I am allowing myself to talk here as if the information cluster can be distinguished from the file in order to coherently describe the envisioned scenario. To the extent that this causes discomfort or tension, this suggests that the idea of a descriptive file entails counterintuitive persistence conditions for the files it is applied to. This challenge is discussed later in the paper, in section §VI. However, it should also be noted, firstly, that the diachronic identity conditions for files are a difficult issue requiring separate treatment and, secondly, that there are arguments that combination of continued information cluster with difference of file should be allowed and can be accounted for within the file-theoretic framework in terms of a so-called 'conversion operation' by which information from one file is transferred into another (see Recanati (2012) ch.7 for discussion).

³² Contrast the case of a perceptual, demonstrative file, where the criterion for which information is stored in the file is that it must be information accessed through the perceptual connection associated with the file.

predicate, 'is G' should only be stored in the file when the thinker has some reason to believe that 'is F' and 'is G' are coinstantiated. Thus, the description has a 'gatekeeping' role (G3). The file can persist through changes in informational content—information can be introduced and discarded from the file without destruction of the file (Thus, (F4) is satisfied). But, given that the description, 'the F' plays roles (G1)-(G3), this description also sets a limit on this possibility. The file is, by its nature, a file on the F: to the extent that an arrangement of information in file-form could continue when 'the F' is discarded and replaced, this constitutes this *particular* file being closed and a new one being opened (G4).

The point is that, in a case with these features, the marks of file-hood are in place but the reference determination for the file, the information marshaling strategy of the file, the possibilities for error and misinformation in the file, and the identity conditions over time of the file are determined *satisfactorily*, rather than *non-satisfactorily*. Given our understanding of descriptive thought as satisfactorial, and singular thought as non-satisfactorial, we have reason to say that the content of thoughts employing such a file is descriptive. This, we can now see, is perfectly consistent with the file *being* a file.

This brings us to the first conclusion of the paper: There is nothing inconsistent—nothing at odds with the constraints of the file-theoretic framework—about the notion of a *descriptive file*. File-based descriptive thoughts are *possible*. This may seem like a modest claim, but it is not insignificant. To the extent that MFC proponents wish to claim that there is an essential or *conceptual* connection between file-based thought and singular thought, this contradicts their position.³³ Given the kinds of motivations for thinking that file-based thoughts are singular, a tempting conclusion is that it is part of the very *nature* of mental files that the thoughts one thinks with them are singular. However, we have now seen that this is not the case.

IV Real-Life Descriptive Files

Having illustrated that file-based descriptive thought is possible, let's now apply the notion of a descriptive file to an actual case.

To do this, we can look to cases where mental-files are initiated to keep track of the satisfier of some descriptive condition, for example (but not exclusively), cases of files associated with so-called 'descriptive names'.³⁴ These cases will be a site of dispute between the MFC proponents and me.³⁵ Robin Jeshion, for example, has emphasized these cases [Jeshion (2002), (2009), (2010)], using them to argue that there can be singular thought

³³ Recanati (2012) seems to imply this kind of view at times (e.g., see 34-35 & 165)

³⁴ Some descriptive files may be opened and maintained to keep track of the satisfier of a definite description, without the introduction of an accompanying name.

³⁵ Let me note in advance that, by claiming that some mental-files associated with descriptive names are descriptive files, I do not thereby claim that descriptive names (the linguistic devices belonging to a public language) are not singular referential linguistic terms. I also need *not* claim that *all* files associated with descriptive names are descriptive, rather than singular.

without acquaintance, but I think what they really show is that there are descriptive files in the sense outlined in §III.³⁶

We can use Evans's 'Julius' case as the basis for an example.³⁷ Evans (1982) gave the case of a name, introduced into the language through an act of descriptive reference-fixing: *'let us call whoever invented the zip, Julius'*. What I wish us to imagine is a corresponding mental file, opened with the description 'the inventor of the zip'.³⁸ Adding some details to the case will serve to establish that ours is indeed a case in which a thinker opens and maintains a mental file on the satisfier of this description.³⁹

In 1948, the Swiss engineer George de Mestral invented the hook and loop fabric fastener, Velcro, apparently conceiving of his invention as a replacement for, among other things, the zipper.⁴⁰ We can imagine de Mestral, in the years before his invention achieved commercial success, frustrated by the lack of appreciation of his invention, and fixated on the inventor of the rival (and already successful) product, the zipper. Imagine that de Mestral's disappointment and frustration are directed at the idea that the inventor of this other product, the zipper, whoever he may be, is somewhere enjoying the success that should be his.⁴¹ He probably lives in the largest mansion in town, he decides; he probably spends his winters in the south of France. We can also imagine that de Mestral has convinced himself that the inventor of the zip is at the helm of an elaborate plan to stop Velcro from achieving its rightful success. After each failed meeting with a potential factory producer or marketer for his invention, de Mestral reaches the conclusion that the inventor of the zip must have paid off the person in question. After all, de Mestral believes, this man no doubt has all the right connections in manufacturing. When Velcro eventually becomes popular as a fastener used in scuba gear and children's clothing, usurping revenue from the

³⁶ In this paper, I am not concerned with all the details of Jeshion's own account of singular thought, but rather with the broader question of whether file-based thoughts are always singular. In [Goodman (unpublished. b)] I give an argument against Jeshion's *particular* version of the MFC.

³⁷ Jeshion has her own stock of favorite examples of file-based thoughts involving descriptive reference fixing [Jeshion (2010), 116-117]. I choose not to use one of these particular cases because they all introduce complicating factors such as the central role of causal traces of the putative referent of the file (e.g. *Unabomber*, *Bearprint*), emptiness (e.g. *Vulcan*), or future directed reference (e.g. *Desert Sensations*). I choose to focus instead on a simpler case to elucidate the notion of a descriptive file. However, I discuss one Jeshion's own cases later in the paper.

³⁸ If one were committed to the claim that a mental-file on an object can only be formed when one has a causal, informational connection to that object, one might react by denying that this could be a case of file-thinking, but neither I, nor the principle MFC proponents find this denial plausible. To the extent that a mental-file is a matter of information clustering, it is unclear why a case of information clustering could not begin with an act of descriptive reference fixing, rather than acquaintance. Even MFC proponents who hold that mental-files require acquaintance (e.g., [Recanati (2010) & (2012)]) also attempt to accommodate the possibility of this kind of case.

³⁹ Jeshion (2009) (2010) has disputed both that (Evans's) 'Julius' is a genuine descriptive name, and that Evans's case would meet the conditions for opening a mental file on the inventor of the zip. She argues that there is a *significance* condition on both naming and singular thought, which is not satisfied in Evans's (1982) description of the case. Here, I amend the case such that Jeshion's *significance condition* is satisfied.

⁴⁰ Although de Mestral did invent Velcro, the rest of this story is fictional embellishment.

⁴¹ If one is concerned that the descriptive nature of the case is undermined by de Mestral's contact with causal traces of the zipper's inventor in the form of perception of instances of his invention, we could stipulate that de Mestral has never seen a zipper.

zipper, de Mestral rejoices in the belief that his rival is defeated, and hopes his French holiday house is falling into disrepair.

It is psychologically plausible that, in the case described, de Mestral opens and maintains a mental file, using the description ‘the inventor of the zipper’. It is clear that de Mestral is concerned to, and actively engaged in, the psychological project of keeping track of the satisfier of the description, ‘the inventor of the zip’ in thought. His pattern of thoughts, interests and affective states suggest his body of beliefs satisfies Jeshion’s *significance condition* on file-possession [Jeshion (2010), 125-126]. Furthermore, we can imagine that his body of thoughts and beliefs are marked by the kind of automatic updating and inference patterns that go along with the use of a mental file.⁴²

Assuming that de Mestral’s thoughts involve the use of a mental-file, what do we learn from cases of this kind? On the *assumption* that the MFC is correct, cases like this one would seem to show that there are instances of singular thought without so-called ‘acquaintance’. This is what Jeshion (2010) takes cases like this to show. But, I believe this is the wrong conclusion. Instead, by thinking about the possibilities and limits of the use of de Mestral’s mental-file, we see it has the features of a *descriptive file*. Plausibly, the description ‘the inventor of the zip’ does not *merely* fix the referent of the de Mestral’s mental-file (as the MFC proponent must claim). It plays roles (G1)-(G4) from §III. I give a description of how the description can be thought to play each of these roles in turn.

Firstly, the description determines which object the file is about (G1). De Mestral’s file concerns the inventor of the zip, whoever that person is. If there were no such person, the file would be an empty file.

Secondly, the kinds of mistakes or misinformation the file allows for, and those it precludes, are set by the description, ‘the inventor of the zip’ (G2). De Mestral could be mistaken in his belief that his rival lives in the biggest mansion in town. He could be wrong that he is well-connected in manufacturing circles and that he is at the helm of a plan to prevent de Mestral’s success. ‘Lives in the biggest mansion in town’, ‘has all the right connections in manufacturing’, and ‘is at the helm of a plan to thwart me’ can be *misinformation* about the inventor of the zip stored in the file. But, despite the fact that the file can contain misinformation—and so cannot simply be identified with the information stored in it—certain mistakes are *not* possible. De Mestral could *not* be wrong, that his rival is the inventor of the zip.⁴³

Thirdly, the description plays the ‘gatekeeper’ role with respect to the file (G3). As we saw in §III, every file must have an ‘information marshaling strategy’ that determines which information enters the file (it answers the question, ‘*does this information concern the referent of this file?*’). For a perception-based file, only information accessed through a particular perceptual channel enters the file. For a file where there is no informational connection to govern the file something must play this role. In the case of de Mestral’s file, I submit, a *description* plays

⁴² We can grant that, *without* the intention or possibility of forming beliefs about the object in question, and in the *absence* of any particular psychological dispositions with respect to reasoning about that object, this would not be a case of forming and maintaining a mental-file.

⁴³ It is *a priori* for him that, if his rival exists, he is the inventor of the zip.

this filtering role. De Mestral is concerned with the object of his file only insofar as that person is the inventor of the zip. The kinds of beliefs he forms about this person are driven by his fairly myopic conception of him as the person who invented a particular product, and what follows, according to him, from this fact. His criteria for forming beliefs about the person in question is roughly that the person who invented the zip—a successful, rival invention to his own—is likely to have some given property, or not.⁴⁴

Furthermore, the predicate, ‘is not the inventor of the zip’ could not enter de Mestral’s file (without irrationality), because it is a file on *whoever* is the inventor of the zip. There’s a *contrast* here between a governing description and other predicates stored in the file. Consistency of information is a norm on all files, and any file is defective to the extent that it contains the predicates ‘is *G*’ and ‘is *not G*’. The norm of *consistency* implies that, if a file contains ‘is *G*’, then it’s a normative condition on adding ‘is *not G*’ that ‘is *G*’ is discarded. However, it is always allowable to add some bit of information that is inconsistent with information already stored in the file by using a ‘discard and replace’ strategy. The predicate ‘lives in a penthouse apartment’, or ‘lives in a garrett’ could always in principle enter the file without irrationality, as long as ‘lives in the biggest mansion in town’ is discarded. But, this strategy is *not* available with respect to the description that governs the file, so it is not available for, ‘the inventor of the zip’.

This is related to our final governing role from §III: the description determines the file’s persistence conditions (G4). It is not possible for the description, ‘the inventor of the zip’ to be discarded and replaced *within the context of continuation of the same mental-file*. To the extent that one discards and replaces a governing description, one discards an old file and opens a new one. That the description ‘the inventor of the zip’ plays this role fits well with picture of the nature of de Mestral’s mental file and its purpose in his mental economy. The very point and purpose of his file is to manage his thoughts and beliefs about the inventor of the zip, whoever this may be. To the extent that this descriptive information is discarded and replaced, it makes sense to theorise this change as the end of one file and the beginning of another.

All this supports the claim that the de Mestral case as described is an instance of a *descriptive file*. It is a file on *whichever thing* happens to be the inventor of the zip. The file’s possibilities for error and persistence, as well as the facts about what information is stored in it, are constrained by this fact. In short, in a case like this one, a description plays not just a meta-semantic or reference-fixing role, but also a constitutive or governing role. Furthermore, none of this is inconsistent with the file having the central features of a mental-file.

Before moving on, it is important to note the following point. To argue that de Mestral’s file as described here is a *descriptive file* is not to argue that we couldn’t conceive of a *different* case, wherein de Mestral *initiated* a file using the description, ‘the inventor of the zip’,

⁴⁴ The fact that de Mestral’s conception of the inventor of the zip is somewhat fanciful and distorted adds to the intuition that a fairly thin description, and a peculiar sense of what follows from this, plays its gatekeeper role. As we’ve described the case, de Mestral has not made any real effort to consider or learn about his rival’s identity beyond his status as the inventor of the zip.

but for which this description did *not* govern the file. For example, perhaps a case in which de Mestral's concern for the inventor of the zip was, as it were, *broader* or less myopic, might yield a file where the description did not play such a central role with respect to the resulting file. My purpose in this paper is not to claim that all files initiated by description are *descriptive files* in my sense, but only to claim that some are, and to illustrate the features that mark them out as such.

V Preliminary Morals

What does the existence of descriptive files show about mental-files and singular thought?

It shows that there are cases for which it is plausible to say *both* of two things: First, the thoughts involved are file-based. Second, the thoughts involved are *satisfactional* or *descriptive*.

Let me pause to reinforce the motivation for the second of these claims. Why, you might wonder, am I entitled to say on the basis of showing that there are cases of *descriptive files*, in my sense, that the thoughts employing these files are descriptive thoughts, or have *descriptive content*?

If we think back to our introduction of the notion of singular thought, the idea was that descriptive thoughts involve thinking of an object as the object that satisfies a certain descriptive condition. Singular thought was then defined negatively by contrast: singular thoughts are cases where one thinks of an object, but not merely as the object that satisfies a certain descriptive condition, or not merely as the object with certain properties. It is hard to think of a decisive argument that we are *necessitated* to conceive of thoughts employing descriptive files in my sense as satisfactional thoughts but, given our criterion for the singular/descriptive thought distinction, it seems very reasonable to do so. It is reasonable to think that, insofar as a descriptive condition plays certain constitutive roles with respect to thoughts employing descriptive files—not just fixing their referents, but also playing other governing roles with respect to the files they employ—these thoughts fit the bill for being satisfactional thoughts. When we're doing theory of mind and assigning contents to mental states, we have good reason to assign these thoughts with *descriptive* rather than singular content. If what we are trying to capture by assigning a mental state with descriptive rather than singular content is that it is a case of thinking of an object *o* as the satisfier of a certain descriptive condition, then, if the object-concept (in these cases, the file) used to think of *o* is governed by a description, we seem to be capturing the right kind of features by ascribing the mental state with descriptive content.

I've already said that the MFC proponent will tend to describe cases of descriptive file-based thought as ones in which a description 'fixes the referent' of the file, but does *not* give its semantic content.⁴⁵ In claiming such file-based thoughts are *singular*, she will appeal to a sharp distinction between the *meta-semantics* for the file, and its *semantic content*. Now, it is certainly true that any argument that concluded that the semantic content of thoughts employing a given file is descriptive *simply* on the basis that the file had its reference *fixed* by description would beg the question against the MFC. But it is important that this is not the

⁴⁵ This is certainly how Jeshion (2009), (2010) conceives of these cases.

kind of argument I have given. The question we must ask is whether the MFC proponent's use of the distinction between 'fixing the referent' and 'giving the meaning' is *justified* in all cases of files whose referents are fixed by description. I take myself to have argued that she has not justified in her use of this distinction for these cases—she is not justified in her assumption that the description plays a merely meta-semantic role. I have done this by showing that a description plays *more* than a mere reference-fixing role with respect to a descriptive file: it plays a constitutive and governing role.

By recognising the possibility of descriptive files we come to a better understanding of mental-files, and of their relation to questions about singular thought. Mental-files are opened and employed when it is useful and important to keep track of a particular thing, regardless of whether one's of thinking about that thing is descriptive or non-descriptive. A mental-file is a *means of mental organization*. Files should be defined in terms of a certain functional role they play in a mental economy, and this functional role—or means of mental organization—ought not to be associated exclusively with either descriptive thought *or* singular thought. There is nothing in the very notion of a file that is essentially connected to singular thinking, and some actual files are used to think descriptive thoughts.

VI Limitations and Complications

The story does not end here, however. Firstly, descriptive files are limited in certain ways that singular files may not be. Secondly, in applying the notion of a descriptive file to certain cases, complications arise. These complications are instructive, and they motivate the introduction of a distinction between *fully descriptive files* and *holistically descriptive files*.

VI. I *Limitations of descriptive files:*

I will address the *limitations* of descriptive files first. These don't pose difficulties for my argument against the MFC. Rather, they illustrate *differences* between singular and descriptive files.

I have argued against the MFC by establishing that there are file-based descriptive thoughts, but I don't wish to deny that descriptive files are limited in certain ways that singular files may not be. In the case of a descriptive file, there is a particular piece of information that has a privileged, governing role with respect to the file, and this restricts certain possibilities for the file: first, certain errors are precluded and, second, so are certain possibilities for information change consistent with persistence of the same file.

First, for a file governed by the description 'the *F*', the thinker could not be wrong that the file's referent is the *F*. This means that a descriptive file cannot contain *entirely* false information about its object. This *perhaps* represents a difference from singular files. In the case of a singular, perception-based file, for example, since there is no particular piece of information that plays a privileged, governing role, it is in principle possible that all the information stored in the file is misinformation.⁴⁶

⁴⁶ This is not a totally straightforward claim, since there may be sortal or categorial restrictions on referential success, even for the case of many singular files. This *might* entail that even singular files can't contain *entirely* false information. This would depend on the precise nature of the sortal or categorial restrictions: it is only

Second, a descriptive file governed by ‘the *F*’ could not discard and replace the predicate ‘the *F*’ (while still remaining the same file). Why this is so was explained in my discussion of (G4) in §III, but is worth reviewing. Since the content of a thought employing a descriptive file is given by its governing description, changes in the file’s governing description will entail a difference of semantic content of the thoughts and beliefs based on the use of that file.⁴⁷ But, on the view whereby mental-files are conceived as object-concepts accounting for sameness and difference of content, a difference of content entails a difference of mental file. The theoretical role of files as object-concepts and content-determiners commits us to the implication that, when we have two instances of file-based thought employing files with different governing descriptions, we have two different files. Thus, the governing description of a descriptive file cannot be discarded while the file remains the same file.

This constraint on file identity and persistence does not apply to singular, or non-satisfactorial files. Since no particular piece of information plays a governing role in such cases (paradigmatically, the governing role is instead played by an informational relation to the file’s referent), there are in principle no information-change limits set on the file’s persistence over time. In principle, *any* piece of information could be discarded from the file, and the thinker could go on thinking thoughts by employing the *same file*.

Should these two differences make us think there is no such thing as a descriptive file? We can imagine someone suggesting what is *distinctive* about mental-files are the very features with respect to which descriptive files are limited. Files allow for successful reference in the context of *radically* false belief; they allow for sameness of content and continued belief *unconstrained* by any continuing descriptive information. To the extent that these features are limited in the putative case of descriptive files, why call these cases of files at all?⁴⁸

It seems to me that this objection is unconvincing because it conflates conditions on what it takes to be a file with certain possibilities that might be of interest to us about *particular* cases of file-hood. Note that we can grant the limitations of descriptive files outlined above, but still point to the fact that the main features of file-hood (spelled out in §III, by (F1)-(F5)) apply in these cases. This was in effect what I aimed to show in §III- §IV. The point, therefore, is that some cases warrant both the title of *mental-file* (insofar as they are marked by certain patterns of update and organization, and a distinction between their semantic content and the information they contain), and the title of *descriptive* (insofar as their limits are set by the governing role of some particular piece of descriptive information). The ‘objection from limitations’ to the idea of a descriptive file seems to needlessly assume just what the argument for descriptive files shows we shouldn’t: that there couldn’t be genuine files that are nonetheless limited and constrained in particular ways.

VI.II *Complications for application:*

entailed if the form of the sortal restriction was not just that certain sortal errors could *not* occur, but rather that certain sortal information was *required* to achieve reference. For argument against the view that this is a requirement on all singular thought, however, see [Goodman (2012)].

⁴⁷ Again, I abstract here and in the rest of this paragraph from the predicative part of the content: Thoughts employing a descriptive mental file governed by ‘the *F*’ will have the content, ‘the *F* is ...’

⁴⁸ Thanks go to Simon Prosser for helpful discussion of this question.

The fact that descriptive files have limitations does not generate problems for the claim that there *are* descriptive files, but these limitations are related to complications that arise when we apply the notion of a descriptive file to *some* particular cases.

To spell this out, we can use a favorite case of Robin Jeshion's. The case is one that seems at the outset to fit the framework of a descriptive file, but also causes some strain to that framework.

Jeshion (2010) tells us that, some time ago, her father had a plan to open a cake-delivery business. Well before the business was in place, he constructed a range of plans and projects in connection with his future business, and named the cake delivery business he would one day open, 'Dessert Sensations'. This is a real-life case of a name whose referent was fixed by an act of descriptive-reference fixing (as in, 'I hereby name the cake-delivery business I will one day open 'Dessert Sensations').

Jeshion claims—I think plausibly—that her father's thoughts about his future cake delivery business involve the use of a 'Dessert Sensations' mental-file, but let's fill out some of the psychological features of the case that justify the claim that there is a mental-file in use.⁴⁹

Let's imagine Jeshion's father, thinking about his future business.⁵⁰ We can imagine that he believes it will make and deliver French pastries. It will make the best éclairs in Pasadena. It will be the only business in Pasadena to deliver fresh éclairs right to your front door. He believes, and intends, that this cake delivery business will be the venture into which he pours his life's savings. In short, such details are meant to show that Jeshion's father is concerned to, and engaged in, the project of keeping track of 'Dessert Sensations'—the cake-delivery business he will one day open—in thought. Further, imagine that Jeshion's father's thoughts about his future business involve the kind of automatic inference and updating patterns that are the mark of the use of a mental-file.

MFC proponents conceive of this case as one of singular thought without 'acquaintance'. Granting the MFC, establishing that Jeshion's father has a 'Dessert Sensations' mental-file, entails he has singular thoughts about his future cake delivery business. But, having brought to light the possibility of *descriptive files* in my sense, we might wonder whether Jeshion's case is a case of *descriptive* file-based thought. And, in fact, this claim does not look implausible. Since, like our de Mestral case, the file is opened to keep track of the satisfier of a descriptive condition, and since it does not involve an information connection to its object, it is at least intuitive to think that the description 'the cake delivery business I will one day open' will do governing work with respect to the file. It is in any case intuitive to think that this description plays at least roles (G1)-(G3) from §III.

Firstly, it is plausible that the description plays the role of *determining which object the file is about* (G1). I take this intuition not to require explanation. Secondly, if this description

⁴⁹ Jeshion's claim is not that the mere fact of an act of descriptive reference-fixing entails that a file is opened and maintained. She claims that 'significance' is a condition on the possession of a (non-perceptual) mental-file (See [Jeshion (2010), 136]).

⁵⁰ Here, I add details to Jeshion's example that are not present in Jeshion (2010), but these details are in keeping with what she says about the example, and the reasons she thinks she is justified in attributing a 'dessert sensations' mental-file to her father (see [Jeshion (2010), 117 & 127]).

determines which object the file is about, it also seems to set constraints on the kind of mistakes or misinformation the file can contain. Jeshion's father could be wrong that his éclairs will be the best in Pasadena and wrong that his business will be the only one in Pasadena delivering éclairs right to your front door. But, if the object of the file is the cake delivery business that Jeshion's father will one day open, whichever thing that is, it is not possible that he is mistaken that, if Dessert Sensations does come to exist, it is the cake delivery business he will one day open.⁵¹ Thirdly, it seems reasonable to think the description plays the 'gatekeeper' role for the file (G3). Given the lack of an informational connection to the file's referent, and given the need for some information marshaling strategy for the file, it is at least plausible that the description on the basis of which the file was opened would play the role determining which information enters the file—that it would play the role of answering the question, '*does this information concern the referent of this file?*'.

But, in working through the suggestion that the description, 'the cake delivery business I will one day open' *governs* the Dessert Sensation mental file, we also experience some strain.

Firstly, although the intuition that the description, 'the cake delivery business I will open' determines which object (if any) the file is about hardly seems to require defense, this claim has the following implication. If this *particular* description plays the role of fixing the referent of the file, then, if Jeshion's father opens a different *kind* of business, that business will not be the referent of the file—it won't be Dessert Sensations (and this is the case, even if he calls that business 'Dessert Sensations'). This cuts against the intuition that, if Jeshion's father decides to open a gelateria instead of a cake delivery business *but in other respects maintains continuity with respect to his plans, projects and beliefs concerning his future business*, the gelateria he opened *could be* Dessert Sensations, and the thoughts and beliefs he previously entertained employing his 'Dessert Sensations' mental-file would be thoughts and beliefs about that gelateria.

Secondly, the application of our final governing role—that the description determines the file's persistence conditions (G4)—has similar counterintuitive results. If the description plays this role, then it is not possible for the description, 'the cake delivery business I will open' to be *replaced* with 'the gelateria I will open' *within the context of continuation of the same mental-file*. To the extent that one discards and replaces a governing description, one discards an old file and opens a new one. This means that, if Jeshion's father decides to open a gelateria instead of a cake delivery business, this change will constitute his coming to think with a new mental-file, despite the fact that his plans, projects and beliefs may otherwise be marked by continuity.

VI.III *The 'which description?'—problem*

The difficulties of applying the descriptive file framework to the 'Dessert Sensations' file can be summed up by what I shall call the '*which description?*'—*problem* .

⁵¹ It is *a priori* for him that, if Dessert Sensations exists, it is a cake delivery business.

If a descriptive file is governed by some description, then we should be able to ask in any given case, ‘*which description governs the file?*’. But, in some cases, this question will be hard to answer, and will entail counterintuitive implications.

For example, in my discussion of the Dessert Sensations case, I worked with the hypothesis that the governing description for the file was ‘the cake delivery business I will one day open’. But, if *this* is the governing description for the file, then it isn’t possible that Jeshion’s father is mistaken that Dessert Sensations is *a cake delivery business*. However, as I’ve said, we might wonder whether Dessert Sensations could turn out to be a business that *makes cakes* but doesn’t deliver them, or it could turn out to be a *gelateria*, depending on how Jeshion’s father’s plans evolve. Relatedly, the specification of ‘the cake delivery business I will one day open’ as the governing description for the ‘Dessert Sensations’ file results in counterintuitive persistence conditions for the file. A descriptive file can’t discard and replace its governing description while continuing to exist *as the same file*. But, this constraint entails that, in the case where Jeshion’s father decides to open a *gelateria*, and therefore discards ‘*the cake delivery business I will open*’ and replaces it with ‘*the gelateria I will open*’, he closes his ‘Dessert Sensations’ mental-file, and opens another. In short, the problem is that the *specification* of any particular governing description cuts against our intuitions about the possibilities of error, and of change, for this particular file. This case looks like it strains the notion of a descriptive file because it looks difficult in this case to specify a *particular* description that plays all the required governing roles for a descriptive file.

There *are* responses to the above complications to which we could appeal, but is unclear how convinced we should be by them. One solution to the ‘which description?’-problem is to look for a governing description that is *general* enough to accommodate our intuitions about possible kinds of error and change (for example, ‘the business I will open’). However, it may in the end look difficult to determinately specify a particular description that gives the right results. A second option is to appeal to vagueness: even if we are committed to a file being governed by a description, it is reasonable to think there may be some *vagueness* as to which description this is. Thirdly, we might wonder how damaging it really is if implications about file-persistence and possible error are counterintuitive.⁵² The notion of a mental-file is a *theoretical notion* after all. Perhaps this theoretical tool gets the required clarity to do its theoretical work at the cost of not tracking certain intuitive features.

VII Holistically Descriptive Files

These responses may not look sufficiently convincing to allow us to classify the ‘Dessert Sensations’ file as a descriptive file (at least, without serious reservations). On the other hand, I also don’t think the complications that arise in applying the notion of a descriptive file to this case serve to do away entirely with the intuition that the governing work for this

⁵² In particular, counterintuitive persistence conditions can be dealt with via a story similar to Recanati’s story about *conversion*—an operation on files by which information from a file that ends is placed in a newly created file. This kind of story allows us to accommodate intuitions about continuation of an information-clustering pattern without appealing to continuation of file. Although I lack space to address this strategy here, there may be good reasons to think we will need to appeal to this kind of story in accounting for the diachronic identity conditions of files (I explore this issue in work that is currently in progress).

file *must* be done descriptively. After all, if the distinction between singular and general thought is, in essence, the distinction between thinking of an object *satisfactorily* and *non-satisfactorily*—between thinking of it merely as the possessor of certain properties, and thinking of it but *not* merely as the possessor of certain properties—there is a remaining intuition that Jeshion’s father is only in a position think of it as the possessor of certain properties, or via its properties. This suggests that, if we are faced with a choice between 1) dismissing the ‘which-description?’-problem and classifying the ‘Dessert Sensations’ file as a descriptive file, and 2) accepting it as a case of non-satisfactorial thought, then perhaps we aren’t considering all our options.

For purposes of understanding the nature of mental-files, and also the relationship between mental-files and singular thought, we’re going to look at a wide range of cases of putative file-thinking. It seems to me that different cases are going to land in different places on a spectrum on which we have clear cases of singular files on one end, and clear cases of descriptive files on the other, and perhaps some difficulty classifying cases that lie in between. The case of a perception-based file looks like a good candidate for being clearly singular. The file is an information cluster that is based on a perceptual relation. The perceptual connection, or perceptual channel, fixes the referent of the file if there is one, and does the other governing work for the file (it dictates what information goes in to the file, and so forth). On the other hand, take our de Mestral case. In this case, a description fixes the referent of the file and plays the other governing roles for the file, and it’s pretty clear which particular description that is. This case fits the model of a descriptive file as I’ve outlined it pretty well. Our intuitions about the case do not strain the model of a descriptive file very much and, to the extent that they do, I’m inclined to think that appealing to vagueness about the governing description (or one of our other responses to complications from the end of §VII) is a reasonable response to this strain. So far, this means that some files are singular and some are descriptive. Thus, the MFC is false, and we understand both singular thought and file-thinking a bit better than we did before. Now take the Dessert Sensations case. This case does have some natural fit with the notion of a descriptive file—we’re inclined to suspect that reference fixing and governing is done by description—but certain intuitions we have about the case strain the model, in that we have trouble specifying *which particular description* governs the file. It is with respect to cases like this that we might feel some third category would be theoretically illuminating.

But, what would a third category look like, which could address the strain placed by certain cases on the notion of a descriptive file without conceding that they are simply cases of singular or non-satisfactorial thought? Although I lack space in this paper to provide a full account, I finish with some preliminary suggestions.

The cases we are trying to theorise are ones where, although there seems like some reason to say the file is fixed and governed *descriptively*, the claim that some *particular description* plays the governing role looks shaky. Given this, let us introduce a distinction between two kinds of descriptive files. A *fully descriptive file* is a descriptive file as I have so far described it: it is a file that is governed by a *particular* description. A *holistically descriptive file*, on the other hand, is a file that is governed by descriptive information, but it is not governed by a

particular description. It is a file for which there is no information link that fixes the referent and does the other governing work for the file, but also no *one* description that does all of this. Instead, a core cluster of descriptions—a bit like a core *theory of the object*—governs the file.

How does this help with our difficult cases? And why isn't a holistically descriptive file just a singular file? In answering these two questions, we can fill out a preliminary exploration of the notion of a holistically descriptive file in a little more detail.

Firstly, the notion of a holistically descriptive file helps with difficult cases because holistically descriptive files are such that, in principle, they allow for complete information overhaul—that is, they allow for 'ship of Theseus'-style changes. In the case of a holistically descriptive file, there is no one governing description for the file, so the process of information overhaul is such that no particular piece of information is held sacred. This means that, in cases where we want to accommodate intuitions about the persistence of the same file despite discarding and replacement of governing descriptive information, we are able to do so. However, having said this, at any *given* time, a core cluster of descriptions (or the core theory of the object) performs the roles of fixing reference and doing other governing work for the file. Reference and possibilities for error are determined by the logical sum, or disjunction, of the descriptions, and the core cluster of descriptions performs the gatekeeping function. In this way, I am appealing to a distinction between synchronic and diachronic possibilities for a holistically descriptive file. The synchronic possibilities for the file are constrained by the privileged cluster of descriptions for the file at that time. The diachronic possibilities for the file are such that the cluster could change.

Secondly, why isn't a holistically descriptive file just a non-satisfactional file? Why am I justified in introducing a third category? If we've admitted that there is no particular description that constrains mistakes and persistence for a holistically descriptive file, then isn't a holistically descriptive file just one that is unconstrained in the same way that singular files are unconstrained? Another way of putting the question is this: it seems like the mark of a descriptive file—and of satisfactional thought—is that certain errors within the scope of success *aren't* possible, and that complete overhaul of information *isn't* possible. So, why isn't the kind of case I'm describing just a singular file?

First, let me address the question about overhaul. For a holistically descriptive file, complete overhaul is possible, but we can see the point of *distinguishing* this kind of file from a singular file by asking about the *process* by which the overhaul takes place. Assuming complete overhaul is possible in the case of a perceptual file, the process of overhaul is governed by an informational link to the object.⁵³ The process of overhaul depends on the information the perceptual link provides. Now, contrast the Desert Sensations case. In this case, there *is* no informational link, so the overhaul can only take place by a process that appeals to inference concerning descriptive information that is stored in the file. It seems

⁵³ As I've noted, it isn't straightforwardly the case that this is possible for perception-based files, but let us assume for purposes of argument that it is. See n.45 of this paper, and also the considerations raised in [Dickie (2011)]

like this is a contrast worth marking, and the notion of a *holistic* but still *descriptive* file allows us to mark it.

Second, let me address the question about possible errors. There is, again, a contrast to be noticed here between paradigm singular files and holistically descriptive files. And, again, appealing to the distinction between synchronic and diachronic possibilities, which is central to the notion of a holistically descriptive file, helps us mark it. It is true that, *over time*, there is no error that is out of bounds for a holistically descriptive file. This is because no descriptive information is held sacred with respect to the overhaul possibilities for the file over time. However, *at any given time*, it is *not* true that any error at all is possible. Since the core descriptive theory of the object at a particular time does the governing work for the file *at that time*, the disjunction of the descriptions in that core theory *has* to be true of the referent of the file at that time (if it has one). Furthermore, you cannot *at that time* store information in the file that is inconsistent with this disjunction, on pain of irrationality. Therefore, it is not the case that, at any given time, that a holistically descriptive file could contain *entirely* false information. And, again, insofar as this distinction is worth marking in our theory of files, the notion of a file that is holistic, but still descriptive, allows us to mark it.

VIII Conclusion

I conclude by saying something about what I take myself to have shown in this paper, and where this leaves us with respect to our understanding of mental-files, and their relationship with singular thought.

Most centrally, I hope to have shown that proponents of the MFC have not entitled themselves either to the claim that descriptive file-based thoughts are impossible, or that all file-thinking is singular thought. By introducing the notion of a *descriptive file*, I hope to have illustrated that the central features of file-thinking are consistent with the possibility of descriptivism about *some* files. Thus, it is a mistake to assume that all file-based thought is singular thought. This argument against the MFC stands, regardless of our conclusions about whether the model of a fully descriptive file fits a wide range of cases, or whether fully descriptive files are relatively rare.

Cases in which a file is opened and maintained in order to keep track of the object that satisfies some descriptive condition provide examples of file-based descriptive thought. This is not to say, however, that the notion of a *fully descriptive file*—that is, of a file governed by a *particular* description—fits all instances of descriptively governed file-based thought equally well. Some cases may cause strain to this notion, for reasons connected to the ‘*which description?—problem*. But, these complications should not necessarily incline us to give up on the descriptive files framework when theorising these cases. Rather, some of these cases motivate introducing the notion of a *holistically descriptive file*, which I have briefly outlined here, but have not fully explored. Once we introduce this notion, we are able to adopt a view on which, although there are some singular files and some fully descriptive files, some file-based thoughts may be holistically descriptive.

More work remains to be done both in fully spelling out the notion of a holistically descriptive file, and providing a detailed classification of cases of file-thinking into singular,

fully descriptive and holistically descriptive. However, in this paper, we have made some progress with respect to our understanding of the relationship between mental-files and singular thought by recognising, firstly, that the MFC is false because fully descriptive files are consistent with the file-theoretic framework and there are some fully descriptive files and, secondly, that we can appeal to a distinction between *fully* and *holistically* descriptive files in theorizing the full range of cases of files governed descriptively.

Our central lesson is the following. Singular thought is *non-satisfactional thinking*—it is thought that involves thinking of an object, but not merely as the possessor of certain properties. Mental files are a *means of mental organization*—they show up any time information is clustered in thought so as to allow certain patterns of automatic inference and updating. This means of mental organization can be exploited both when an object is thought of non-satisfactionally, *and* when an object is thought of merely as the possessor of certain properties. By answering the question of whether a particular thought employs the use of a mental file, we do not answer the question of whether the thought is satisfactional or non-satisfactional.

Bibliography

Bach, K. (1987) *Thought and Reference*. Oxford: Oxford University Press.

- Bach, K. (1997) 'Do Belief Reports Report Beliefs?', *Pacific Philosophical Quarterly*, 78p. 215-241
- Evans, G. (1982) *The Varieties of Reference*. J. McDowell (ed.), Oxford: Oxford University Press.
- Evans, G. (1985) 'The Causal Theory of Names', reprinted in *Gareth Evans: Collected Papers* Oxford: Oxford University Press.
- Campbell, J. (1987) 'Is Sense Transparent?' *Proceedings of the Aristotelian Society* 88: 273-92
- Dickie, I. (2011) 'How Proper Names Refer' *Proceedings of the Aristotelian Society* 111: 43-78
- Goodman, R. (2012) 'Why, and How, Not to be a Sortalist about Thought' *Philosophical Perspectives* Volume 26
- Goodman, R. (unpublished. a) 'Do Acquaintance Theorists Have an Attitude Problem?'
- Goodman, R. (unpublished. b) 'Cognitivism, Significance and Singular Thought.'
- Goodsell (2013) 'Mental Files and Their Identity Conditions', *Disputatio*, Book Symposium on Recanati's *Mental Files* V, 36, 177-191.
- Grice, H. (1969) 'Vacuous names' in D. D. J. Hintikka (ed.) *Words and Objections*, 118–145. Reidel.
- Harman, G. (1977) 'How to use Propositions' *American Philosophical Quarterly* 14: 173 – 6.
- Hawthorne, J. & Manley, D. (2012) *The Reference Book*. Oxford: Oxford University Press
- Heim, I. (1982) *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- Jeshion, R. (2002) 'Acquaintanceless de re Belief', In J. Campbell, M. O'Rourke and D. Shier (eds), *Meaning and Truth: Investigations in Philosophical Semantics*. New York: Seven Bridges Press.
- Jeshion, R. (2009) 'The Significance of Names', *Mind & Language*, 24, 372 – 405.
- Jeshion, R. (2010) 'Singular Thought: Acquaintance, Semantic Instrumentalism, and Cognitivism' In R. Jeshion (ed.), *New Essays on Singular Thought*. Oxford: Oxford University Press.
- Kripke, S. (1980) *Naming and Necessity*. Harvard University Press.
- Lockwood, M. (1971) 'Identity and reference' in M Munitz (ed.) *Identity and Individuation*, 199-211. New York: New York University Press

Perry, J. (2000) 'A Problem About Continued Belief', in *The Problem of the Essential Indexical and Other Essays*. Stanford: CSLI.

Recanati, F. (2012) *Mental Files*. Oxford: Oxford University Press

Sainsbury, R. M. (2005) *Reference without Referents*. Oxford: Oxford University Press.

Taylor, K. (2002) 'De Re and De Dicto: Against the Conventional Wisdom', *Nous* Vol. 36, 225-26.