

Eliciting and Conveying Information

Routledge Handbook of Linguistic Reference

(Penultimate draft)

As a young boy I often spent time in the shop of our local printer. The printer sometimes tried to supplement his income by printing comics. He would receive reams of unbound comics that contained the pictures without words together with an original master copy that contained a dialogue. He then translated the original into Icelandic, set the text in lead blocks using a spider-like machine that always seemed to be falling apart and, finally, ran the reams through the press, thus adding the newly translated text to the pictures.

During one of my visits I picked up on some pages of a new comic, pictures only. I browsed through several pages, trying to make sense of the story. When the printer saw what I was up to he stopped his work, walked over to me, placed a couple of pictures from the comic on a table, looked at me knowingly, smiled and winked as only he could do, and then he moved the two pictures toward each other until they overlapped. Not a word was said. But I got it. These two were the same! Many days later, when the printer had finished his work and the comics were ready, I learned that the names of the characters involved were “Clark Kent” and “Superman.”

Almost without exception philosophers think of Frege’s puzzle in linguistic terms. Solving the puzzle means finding a semantic solution. However, my childhood experience shows that the puzzle can be raised and that one can discover the relevant identity without relying on semantics, which indicates that perhaps the focus when dealing with the puzzle should be on, not semantics, but rather some underlying mental states. After browsing through the pictures at the print shop I had gathered various information about, on the one hand, the person wearing a cape, and on the other hand, the office worker. What the printer helped me realize was that, because the two characters were the same, the two sets of information were about the same individual. What is important here is that a semantic solution to the puzzle will not be able to explain my discovery.¹ There were no terms involved in my case and so there is no meaning of the relevant terms involved that captures and explains my discovery.

The lesson that I want to draw from the story is that since Frege puzzle cases can extend beyond semantics, the attempted solutions to the cases should not focus solely on semantics. Solutions that do so can at best provide a partial solution to the puzzle. They will not provide a solution that explains the broader phenomena; the one that includes my childhood case. Below I will provide a solution that accounts for the typical Frege case as well as my childhood case. The solution will, accordingly, not be a semantic solution. Instead it will focus on information we have on objects and how we organize and access the information. The solution I will provide is psychological and not semantic in nature. At the same time I will show that the solution I provide aligns itself well with Millianism.

The Importance of Simple Sentences

For well over two decades the discussion of cognitive significance and the substitutivity of coreferential names focused on names as they appear in embedded sentences, and then particularly in belief reports. Since the Superman stories provide a clear and vivid example that most are familiar with, assume with me that the story details real events, i.e., assume that Lois and Clark Kent are as real as the rest of us and that Superman has, once again, saved Metropolis from a meteor impact.

How, then, can we explain Lois assenting to (2) and not to (1), given that that “Clark Kent” and “Superman” refer to the same person?

1. Lois believes that Clark Kent saved the city.
2. Lois believes that Superman saved the city.

There is no denying that we do have anti-substitution intuitions when it comes to sentences like the ones above. How do we explain those intuitions? Those who accept a Frege-like semantic view can explain the anti-substitution intuition by arguing that the embedded sentences express different propositions. But those who accept direct reference generally cannot resort to that kind of a solution, since the direct reference theorists generally accept the view that the semantic meaning of a name is just the object named. It is particularly difficult to account for the anti-substitution intuition if one is a Millian, as I am, as the Millian further assumes that names only contribute their referents to propositions expressed by simple sentences in which they occur, that the resulting proposition is a singular proposition that can be represented as an ordered pair of an object and a property, and that embedded simple sentences express the same proposition as when not embedded.² Given this,

3. Clark Kent saved the city

and

4. Superman saved the city

express the same proposition that can be represented as

<Clark Kent, saving the city>.

Further, the embedded sentences in (1) and (2) express the same proposition and so (1) and (2) have the same truth value.

There are good reasons to believe that the attention given to embedded sentences was to a large degree misplaced. In the introduction to *Naming and Necessity* Kripke pointed out that he never intended to argue for a doctrine of universal substitutivity of proper names. He pointed out that the sentence “Hesperus is Phosphorus” can sometimes be used to raise an empirical issue while “Hesperus is Hesperus” cannot be used in the same way. Sentences expressing identity that are not embedded, simple sentences, resist substitution as well as do sentences embedded in attitude contexts. That is, our intuitions about substitutions are not limited to embedded sentences. But our anti-substitution intuitions are not limited to simple sentences expressing identity. Chances are that people will resist substitution in (3) and (4), as well as in (5) and (6)

5. Clark Kent is Clark Kent.
6. Clark Kent is Superman.

I believe that an approach that treats simple sentences and belief reports as raising fundamentally different problems, is mistaken. Instead, I believe that the preferred way to deal with substitution problems is to view substitution in simple sentences and embedded sentences as raising the same problem, in which case we should be able to find a unified solution. We have intuitions that tell us that we cannot substitute freely in simple sentences, and we have intuitions that tell us that we cannot substitute freely in belief contexts. I believe that the same intuitions are at work in both instances and the same kind of an explanation of these intuitions should apply to both cases. The solution that I will present focuses on singular propositions and ways of believing them and it embraces Millianism. I will try to honor our pre-theoretic intuitions against free substitutions and at the same time agree with the Millian view that truth value is preserved when substituting coreferential names, both in simple sentences and in belief reports. The key move in accomplishing this is to distinguish what our intuitions are working with: semantics or information. Once we recognize that our intuitions are working on two levels, namely with information and with semantics, we can provide an account that recognizes both the Millian view that substitution preserves truth value and our uneasiness with substitutions. The idea that we can believe singular propositions in different ways is, I believe, fundamentally sound. But it is possible to advance this

basic idea in various ways or tell different stories about how one can believe singular propositions in different ways. The story I will tell will not focus on a semantic solution, and it will not rely on one systematically mistaking one proposition for another, i.e., it will not be a story about pragmatic implicature. Instead, it will be a psychological story about how we represent objects and elicit information about objects when so prompted. But before I tell my story I will look at an account developed by David Braun and Jennifer Saul.

Braun and Saul's Mistaken Evaluation.

David Braun and Jennifer M. Saul have, independently as well as together, presented an attractive account of beliefs of simple sentences.³ The crux of their view is that substitution failures are ultimately to be explained as *mistaken evaluations*. While I do like some aspects of their account, I will argue that the mistaken evaluation view does not work because it does not adequately account for our anti-substitution intuitions.

Braun and Saul develop a view that does not rely on pragmatic implicature.⁴ Consider the following sentences:

7. Superman leaps more tall buildings than Clark Kent.
8. Superman leaps more tall buildings than Superman.

It appears that (7) is true while (8) is false. Braun and Saul suggest that one may maintain two cognitively separated sets of beliefs, or pools of information, about Clark Kent, one which is associated with the name "Superman" and another which is associated with the name "Clark Kent," and that these sets attribute different properties to Superman/Kent, and affect how one believes propositions containing Clark Kent. Braun and Saul write,

These two pools of information attribute different properties to Superman/Clark. For instance, the 'Superman' pool attributes to him the property of leaping tall buildings, while the 'Clark' pool does not. You also associate different images with the two names. When you *quickly* evaluated the proposition semantically expressed by [7], your pools of information and images appeared to you to support that proposition's truth. Therefore, you judged that [7] was true. You did not pause to consider the identity [Superman is Clark Kent] long enough to notice its logical consequences. Or, if you did, you erred in not considering this good reason to alter your original judgment.⁵

Braun and Saul imply here that an enlightened speaker, one who knows that Superman is Clark Kent, and one who is attentive, would see, and accept, that both (7) and (8) express falsehood. By the same token, an enlightened and attentive speaker should accept that both (3) and (4) express truth, given that she knows that Superman has prevented the destruction of her city. Our standard intuitions about sentences like (3) and (4), and (7) and (8), are incorrect, they claim.

Saul develops the Braun-Saul view further in her *Simple Sentences, Substitution, and Intuition*.⁶ There she argues that we have a reason to believe that a subject often stores information about an individual in different nodes or folders when learning about an individual that is presented in different ways. So, why is it then that we tend to have the intuition that (7) is true? Saul's answer is that we fail to reflect on the identity of Superman and Clark Kent and hence do not integrate the information about Superman and Clark Kent. Since we have good reasons for not reflecting on the identity and not making the relevant inferences, we don't do so. Hence, Saul concludes, we take (7) to be true.⁷

But what if we are aware of the relevant identity and do reflect upon it and still don't realize that (7) and (8) must have the same truth value? Saul's answer is that even if we reflect upon the double life of Clark Kent that does not mean that we make all the inferences that one can make

about him. Saul provides several possible explanations for this failure to make the relevant inferences. We might simply fail to make some of the inferences, just as students sometimes fail to make simple inferences. Or, it might be the case that we are more convinced of the truth and falsity of the relevant statements than the inference that shows that both have the same truth value. Or, it might be the case that one has greater confidence in one's truth-conditional judgments than one's inference making.⁸ So, the fundamental answer that Saul provides for our reluctance to freely substitute coreferential names in simple sentences remains the same as before; we fail to make the relevant inferences.

Saul further emphasizes her view that our intuitions against substitution depend on us failing to make the relevant inferences when she suggests an explanation as to why we speak the way we do. We do utter sentences such as

9. Shostakovich always signaled his connections to the classical traditions of St Petersburg, even if he was forced to live in Leningrad.

Someone who utters (9) may be aware of the identity in question, namely that St Petersburg is Leningrad, and so it is not lack of awareness of the relevant identity that explains this utterance. In fact, the speaker might take herself to be communicating something true and relevant. The speaker seems to be communicating a contrast between living in St Petersburg and Leningrad. However, if St Petersburg is Leningrad there cannot be such contrast. Presumably, the speaker would never utter (10).

10. Shostakovich always signaled his connections to the classical traditions of Leningrad, even if he was forced to live in Leningrad.

So how can we explain that a speaker who is aware of the identity of St Petersburg and Leningrad utters (9) while she would never utter (10)? Here is Saul's answer.

The answer is simple. Sometimes, as argued earlier in this chapter, enlightened speakers well aware of particular identities will fail to make all the inferences that they could from the relevant identity claims...The utterer of [9] knows that St Petersburg is Leningrad, and indeed reflects on this at the time that he utters [9]. Nonetheless, he may think that [9] is true, and not odd. This is because he simply doesn't use his knowledge of the identity in question to infer from [9] to [10], which might give him pause.⁹

Surely an unenlightened speaker could utter something like (9). However, Saul retains the primary explanation of the Braun-Saul view when she argues that an enlightened speaker would only utter something like (9) if he fails to make the relevant inferences.

Contrary to Braun and Saul, I will argue that even the enlightened speaker who makes all the relevant inferences could still have a good reason to utter something like (9), and still have a good reason not to freely substitute in simple sentences. That is, I will maintain that truth value is preserved with the substitution but that nevertheless our anti-substitution intuitions should be respected. I will therefore present a view that is strongly opposed to Saul when she writes "If...I decide that Naïve Millianism is right, I will have good reason to believe that substitution of coreferential names always succeeds—in any context."¹⁰

The Millian needs to tell a story that accounts for how (5) and (6), and (7) and (8), respectively, express the same propositions that, consequently, have the same truth value, while explaining how we rightly have strong anti-substitution intuitions, even when enlightened about the relevant identity. It is such a story that I will tell. The story will rely on a distinction between

- a. the proposition expressed by a sentence, and
- b. the information conveyed and/or elicited by a sentence expressing the proposition.

The proposition expressed by a sentence is a function of the semantic value of its components and so a study of (a) is a study in semantics. But often two sentences that express the same proposition nevertheless convey and/or elicit very different information. Sometimes the difference in

information between two sentences is due to pragmatic implicature,¹¹ but my concern is not with those, as we can provide an account of different information conveyed or elicited without resorting to pragmatic implicature. Instead we can rely on how different sentences convey different information to or elicit different information from the audience depending on how she organizes information. For example, (3) and (4) express the same propositions and so I cannot claim that (3) expresses a false proposition. However, if I tell someone that Clark Kent saved the city, using the sentence “Clark Kent saved the city,” then that certainly conveys the wrong and misleading information that he did so in his Clark Kent outfit and not in his Superman outfit.

I believe that the basic explanation that Braun and Saul provide of our semantic intuitions is a good one. However, I do not believe that their account adequately explains why we have anti-substitution intuitions. That is, I do not believe that the source of the problem and what lies behind our anti-substitution intuitions is mistaken evaluation. Even when I am informed about the relevant identity, that is, even when I am in the know, and consider the relevant statements and the relevant identity there is still something that gnaws at me. Even when I know that Superman is Clark Kent there is still this lurking feeling when I consider the relevant statements, a feeling that I share with Fregeans and many neo-Russellians, that we cannot freely substitute coreferential names. Even though I understand the semantic machinery that underlies names and that allows me to claim that truth value is preserved, my intuitions still have me resist substitution and so there is a reason to believe that there is *something else* at work, something beyond semantics, that makes me resist substitution.

Mental Files and Information Webs.

It is plausible to maintain that we file information away in our minds and that we normally organize information about individuals so that different pieces of information about what we take to be the same individual are kept together in a single file.¹² Several philosophers have used the analogy of a file, including my former self, suggesting that it provides a plausible explanatory account of how information is stored. However, there are reasons to believe that not everything in a file is equally salient. Some pieces of information are more salient than other pieces of information, and thus more readily recalled when one thinks about the relevant object. A web analogy, much like Quine's web of belief, captures the idea that some pieces of information are more central than other pieces of information, and that some pieces of information are therefore more stable and more readily evoked than other. For example, it is likely that central to most people's web for Superman is that he is very fast, wears a cape, can fly, and is very strong. It is less important to most that his given name is “Kal El,” and still less important that his father's name was “Jor-L.” Accordingly, that information is likely to be less central in the information web and hence less likely to be recalled when one thinks about Superman.

It is plausible to assume that in typical cases I form a file about what I take to be a new and unique object and then place information that I take to be of the object in the file. However, unlike many, I view the file metaphor as being just that, a metaphor. I am not granting a special metaphysical status to files. Instead, using the file metaphor is a convenient way of talking about us clustering information together. Also, unlike, e.g., Recanati, I am not assigning semantic values to files. For Recanati files refer and names derive their reference from the files they are associated with. On the account that I provide names are the primary referring items. However, we do use names as one of several ways to access the files they are associated with. And finally, unlike Recanati, I do not require subjects to stand in a special epistemic relationship, an ER relation, to an object the file is of. I see no difficulties in one having mental files of Santa Claus and Minnie Mouse without being able

to stand in an ER relation to them, as one can clearly have a wealth of information about these characters.

When my daughter was born, I already had a file in place for her, namely what we might call my Dagny file. When I had my son, I formed my Atli file. In the years since then, I have gathered all kinds of information in the two files. Some information is represented linguistically, either as predicates or propositions. Other information is not couched in linguistic form but rather in imagistic forms or non-linguistic form. I can recall vividly and see it in my mind's eye, much like if I was watching a film or a snapshot, how the kids ran around when younger, how they looked on a soccer field, and how they looked when sleeping in their beds. I can of course describe my mental episodes when I recall these moments, but I recall them not as propositions but rather as images.¹³ In some cases, I can vividly recall scent and mood from certain events in their lives and so the files have a varied and, sometimes a vast amount of information. While my children's elementary school teachers are likely to have files with information about my kids, the information that they have will be very different from what I have. One reason for that is more limited access to my kids than I had. Another reason is that the teachers are likely to focus on different aspects of my kids looks and behavior than do I, and they are likely to have a different view of what is important, and so they will gather different information and organize it in a different way than do I.

Generally, if we mistake a person's identity in such way that we take one person to be two, such as when I do not recognize a student in my class when I see him working the cash register at a local store, then I form two files for the person. In this case, one will contain information about the person as a student, and the other will contain information about him as a store clerk. When I later find out that they are the same person, I will combine the two files into one. But there are instances when we do not combine files when discovering the identity in question.

Suppose that I have one file about Clark Kent as Clark Kent and a second file about Clark Kent as Superman. If I now discover the identity of Clark Kent/Superman, then I will likely retain my two files. There are several reasons why I might do this. Suppose that I initially formed two files because I thought I was dealing with two individuals and then later find out that they are one. One reason I might keep the files separated after the discovery of the identity is that most people do not know about the identity of Superman and Clark Kent. If I combine the two files and all that is in them, then it is likely that my future communication with others about Superman/Clark Kent would be confusing and fraught with misunderstandings. If I, because I have combined the information in the two files, indicate in conversation that Clark Kent, using the name "Clark Kent," is strong and faster than a speeding train, then it is likely that my respondent will be lost, as she has no such information in her Clark Kent file and, in fact, what I am saying contradicts many of what is likely to be the central information in her Clark Kent file.

Second, I might know about the relevant identity of Superman/Clark Kent from the time I first encounter Clark Kent and nevertheless form two files about him and file information accordingly. The reason for me doing so might be the same as above: avoiding confusion in communication.

Third, if someone is developing or has developed a persona that is significantly different from the one he/she normally displays, then a new file might be useful and even necessary to avoid misunderstandings and miscommunications. For example, when attending a performance by David Bowie, I expect one type of show and music. If, on the other hand, I were to attend a performance by Ziggy Stardust, then I expect a different type of a show and for the most part a different style of music than would be performed at a Bowie concert. The personas are different enough so that one has very different expectations about a Ziggy concert than one has of a Bowie concert. Hence, it is often useful to keep files of different personas separated even though one knows that they are of the same object. Accordingly, if we have good reasons to keep files separate even if we know the

relevant identity, then we are hesitant to make inferences that result from combining the two files, as the results of such inferences are likely to be misleading. For example, if I know that I should expect a glam show at a Ziggy concert, it would be misleading to conclude that I should expect a glam show at a Bowie concert, even though I know the relevant identity. One of the reasons for keeping the Bowie file and the Ziggy file separated is the fact that the two put on very different shows.

Given that there are many reasons to keep more than one file about someone, even when the identity in question is known, we need to revisit why it is that we form a file about an object. Earlier I claimed that we file information of what we take to be the same individual or object in a single file. That now must be amended.

We file information about what we take to be the same individual or object in a single file *except* if that individual develops personas or displays aspects so different from his or its usual self that we expect significantly different looks and/or behavior from that persona/aspect than we do from his, hers, or its usual self.

Similarly, when we discover the identity of a person or object that we had taken to be two (or more) persons or object and thus had created two (or more) files, then we tend to combine the information in the files *unless* the personas or objects are so different that we expect significantly different looks and/or behavior from them.

Accessing Files.

I suggest that the file I have about an individual can be accessed by a marker. A marker directs one's thought to where information about the relevant individual or object is stored and so allows one to deposit or retrieve information about that individual or object. A marker can be a name, a description, scent, a picture, or a representation that we associate with a file in a given context. For example, the name "Superman" is a marker for a Superman file as is "The Man of Steel," and as is a picture or an image of the flying superhero. Similarly, the name "Clark Kent" is a marker for my Clark Kent file and so when I see or hear the name my thoughts are directed to the Clark Kent file. So, when I hear or see a sentence containing the name "Clark Kent," then that directs me to the appropriate file and information is added to or elicited from the Clark Kent file, and similarly, when I see or hear a sentence containing the name "Superman," then information is added to or elicited from the Superman file. Similarly, when I see a picture of Superman, then that directs me to my Superman file and not to my Clark Kent file. Typically, the information that is most pertinent in a file represents the most prominent features of the person or object, and it is the information that is most readily accessible and so is the information that is most easily elicited. Accordingly, when, for example, a conversation about Clark Kent prompts me to recall something about him, then some of his most prominent features are most likely to be elicited. The drab bespectacled reporter comes to mind as the drab bespectacled reporter, while his tie color or shoe size, the more peripheral information that is less salient, is less likely to be brought up.

The information elicited from a file determines how I represent the object in the proposition believed and thus determines to a large extent how I believe the proposition. It is therefore of great importance how a proposition is presented to me. Since the names "Clark Kent" and "Superman" are markers for different files, it makes a big difference whether a proposition containing Clark Kent as a constituent is presented to me with a sentence that contains "Clark Kent" or a sentence that contains "Superman." If the sentence contains "Clark Kent," then I am directed to my Clark Kent file, and if it contains "Superman," then I am directed to my Superman file. Accordingly, if someone tells me that Superman saved the city, then "Superman" elicits information from my Superman file, and I believe the proposition expressed accordingly. However, if someone tells me that Clark Kent

saved the city, then “Clark Kent” elicits information from my Clark Kent file. The result is likely to be confusion or disbelief on my behalf, since the information elicited from the Clark Kent file has me represent the hero as Clark Kent and not as Superman and nothing in my Clark Kent file indicates that he, in that guise, has superpowers.

We can expect similar result if different types of markers for files are used. Suppose that I am shown a picture of Superman charging his way to a meteor and then moving it to a safe distance from Metropolis. The image of Superman is a marker for my Superman file and so that is the file that is accessed for past and additional information about the superhero. Superman’s acts are perfectly consistent with the information that I already have about him and so this is just one more heroic act. Suppose on the other hand that I am shown a picture of Clark Kent in his drab suit charging towards the meteor and subsequently moving it. The picture of the journalist has me access my Clark Kent file. This time around the new information does not fit what I already know about Clark Kent, the journalist. The result is, again, likely to be confusion and disbelief.

Enlightened and Unenlightened Subjects.

The file that I have on Clark Kent figures heavily in how I believe a singular proposition that has him as a constituent when the proposition is presented to me with a sentence containing “Clark Kent.” The file I have about him as Clark Kent is what I draw from when I entertain the proposition so presented, and it thus largely determines the way in which I believe the proposition. Similarly, when I believe a singular proposition that has Clark Kent as a constituent and the proposition is presented to me with a sentence containing the name “Superman,” then the file that I have about Clark Kent as Superman is accessed and information is elicited from my Superman file. Accordingly, I believe the proposition in different ways when it is presented to me with the two different sentences.¹⁴

The picture presented provides a simple explanation of why I might assent to (4) and not to (3) while being fully rational. The names used to refer to the object in the proposition expressed by the sentences elicit information from different files and so I believe the proposition expressed in different ways. If I am not aware of the identity of Clark Kent and Superman, then I treat the names as referring to different individuals and my files are set up on the assumption that they contain information about different individuals. The two sentences express the same singular proposition and so, as we wanted, substitution of coreferential names preserves truth value. However, we preserve our anti-substitution intuitions as we can explain those with different names used in the sentences expressing the proposition eliciting information from different files. If we focus on truth values, substitution is fine. Once we turn our attention to information, we resist substitution.

Suppose that I am enlightened about the identity of Clark Kent and Superman. My anti-substitution intuitions are still there, and for a good reason. If someone who has witnessed Superman saving the city tells me that Clark Kent saved the city, using the name “Clark Kent,” then the use of the name “Clark Kent” prompts me to think of Clark Kent *as Clark Kent* saving the city and so indicates that he saved the city while in his Clark Kent persona; a highly unlikely scenario. The use of the name “Clark Kent” results in the Clark Kent file being accessed, and since the file does not contain information about Clark Kent *as Clark Kent* having super strength, the report of Clark Kent saving the city indicates that the superhero has revealed his identity to the public. Again, substitution of coreferential names, while preserving truth values, conveys and elicits very different information.

What if both I and my respondent are fully aware of Clark Kent being Superman? Even then we do not freely substitute the names when talking with each other. Since we are fully aware of the

dual lives that Clark Kent lives, we have two files for him; one for Clark Kent as Superman and one for Clark Kent as Clark Kent. If my respondent now reports that Clark Kent has saved the city, using the name “Clark Kent,” then that conveys that Clark Kent done so in his Clark Kent guise and so has revealed his identity. The use of “Clark Kent” would direct me to my Clark Kent file, prompting me to think of the drab journalist saving the city. The result would at first be disbelief, and if I accepted the new information, then doing so would require a rather massive overhaul of the information that I have on Clark Kent and Superman.

Similarly, if a fellow music aficionado had told me in 2010 that Ziggy Stardust was preparing a tour, then I would have made efforts to attend one of the shows, fully expecting Ziggy to appear as Ziggy and playing Ziggy’s music with the Spiders from Mars, his usual band. I would not have made the same effort had I been told that David Bowie was preparing a tour.

Even when I am in the know about the identity of Superman and Ziggy, I resist substitution in most situations because the information conveyed and/or elicited by a sentence containing the names “Clark Kent” or “Ziggy Stardust” is different than the information conveyed and/or elicited by a sentence containing the name “Superman” or the name “David Bowie.” Contrary to Braun and Saul, there is no mistaken evaluation here. I am in the know about the identity, I might make all the relevant inferences, and I still resist substitution, the reason being that even though the two sentences express the same proposition the information conveyed and elicited is very different. The information does not affect the truth value of the proposition expressed, but it does affect the cogency of free substitution.

Let us revisit Saul’s example of St Petersburg and Leningrad. The two problematic sentences were

9. Shostakovich always signaled his connections to the classical traditions of St Petersburg, even if he was forced to live in Leningrad.

and

10. Shostakovich always signaled his connections to the classical traditions of Leningrad even if he was forced to live in Leningrad.

Saul wrote that an enlightened speaker might utter (9) because she failed to make the relevant inferences. If she used her knowledge of the relevant identity and inferred (10) from (9), then that might give her a pause.

There is, I maintain, nothing odd about uttering (9), even after one has made all relevant inferences, because there is a systematic difference in the two uses of the name of the relevant city. It was known as Leningrad during most of the Soviet era, and St Petersburg both before the rise of the Soviet Union and after its fall. Because of this, it is very likely that people will speak of the city using the name “Leningrad” when they are talking about the city during the Soviet reign, and that they will use the name “St Petersburg” otherwise. For example, people talk about the siege of Leningrad and not the siege of St Petersburg during the second world war. So, while the enlightened speaker will know that St Petersburg and Leningrad are one city, the enlightened speaker will also know that the names refer to the city during different time periods. It makes sense, then, that the enlightened speaker will store information about the city from one time period in the St Petersburg file and information about the city from a different time period in the Leningrad file, and that includes information about a very different political climate during the time periods. Sentence (9) will therefore convey information to and/or elicit information from two different files. One file contains information about St Petersburg during the Soviet era, the other does not. Soviet officials appear to have been attentive to a similar distinction, as the following is true:

Shostakovich’s signaling his connections to St Petersburg instead of to Leningrad resulted in him having to make a public recantation for some of his works!

Even when the relevant information is in place, there is nothing odd about the previous sentence.

When Substitutions are Permitted.

It is evident that our choice of names is of great importance when making assertions as well as when reporting beliefs. If Lois sincerely assents to (4) and not to (3), then it certainly appears that (2) is true and (1) is not true. And if I utter (3) then I wrongly convey the information that Clark Kent *as Clark Kent* saved the city. The anti-substitution intuitions are alive and well. In spite of that the Millian holds that (3) and (4) as well as the embedded sentences in (1) and (2) express the same proposition and that the two simple sentences and the two belief reports have the same truth values.

If I see Superman deal with the meteor and report his feat to Lois by uttering (3), then that is likely to be met with disbelief and/or confusion, Lois not knowing about the identity. But it is likely to be met with the same reaction even if she knows about the identity, because my use of “Clark Kent” directs her to her Clark Kent file, and that file does not have information about Clark Kent doing these kinds of things in his Clark Kent guise.

Before finding out about Clark Kent’s true identity Lois would readily accept (5) while she would not accept (6). Once she is in the know about the identity we still cannot freely substitute “Clark Kent” and “Superman” due to the names directing us to different files and thus conveying and eliciting different information. Similarly, before finding out about the identity, Lois would readily accept (7) and not (8), even though the sentences express the same proposition. Once she knows about the identity, she is likely to still accept (7) due to how the proposition is expressed. The information elicited will include that Clark Kent leaps more tall buildings in his Superman guise than in his Clark Kent guise. In the know or not, (8) will still be puzzling, as the information elicited is not coherent.

The discussion of Clark Kent and Superman helps explain why we readily accept (11) and not (12):

11. David Bowie released his last album in 2016.
12. Ziggy Stardust released his last album in 2016.

Here we have different information associated with different personas and so we need to clearly separate information files on the two personas. Note also that the example can be presented without using the names “David Bowie” and “Ziggy Stardust.” We could, for example, use pictures of Bowie and Ziggy instead and we would readily accept the information with the Bowie picture and not with the Ziggy picture. The sentences as they stand express the same proposition, according to the Millian, but nevertheless we can only accept one of them as conveying or eliciting accurate information. In this case there is nothing odd about saying that Bowie’s last album bears little resemblance to Ziggy’s work.

The examples discussed so far are significantly different from the following. Suppose that I am talking with someone who knew Robert Zimmerman as a child and who never figured out that Zimmerman is Dylan. Suppose further that I know about my respondent’s lack of knowledge about the relevant identity. In that case I choose my words so that they match the information I know my respondent possesses. Accordingly, I utter (13) and not (14).

13. Zimmerman had a successful music career.
14. Dylan had a successful music career.

Alternatively, I might quickly inform my respondent that Zimmerman has taken the name “Dylan” and then I can proceed and utter (14), thus passing on the relevant information.

However, there are contexts in which we can freely substitute since neither the truth value nor the information conveyed and/or elicited is affected by the substitution. Generally, I can freely substitute names if my audience is in the know about the relevant identity and it is not the case that

the person named has developed a new identity or a persona. In that regard Zimmerman's name change is much like that of someone who takes the name of a long-time partner. His or her old friends can, in that case, talk about the person using either the old or the new name. Here I am, of course, assuming that there is not a significant character change that accompanies the name change.

What applies to simple sentences also applies to attitude reports. If I am reporting Stephen's belief to someone who is not in the know about the Clark Kent/Superman identity, then it at the very least misleading to say that if I told the person that Stephen believes that Clark Kent saved the city then I have informed her that Stephen believes that Superman saved the city. Sure, the embedded sentences in my belief reports express the same proposition and so the reports have the same truth values. But the two embedded sentences convey and elicit different information due to different names being employed in them. Assuming that most people have similar key information in their Clark Kent file, the former report conveys the information that Stephen believes that Clark Kent *as Clark Kent* saved the city. The information I thus convey about Stephen's beliefs is not accurate. On the other hand, I can accurately convey his beliefs with the latter report. And it would be outrageous to claim that if I report that Stephen believes that Kal El (Kent's given name on Krypton) saved the city, my respondent never having heard the name "Kal El," I thereby inform her that Superman saved the city. Even though the sentences "Kal El saved the city" and "Superman saved the city" express the same proposition, they convey and elicit very different information.¹⁵ While truth value is preserved with the substitution, the information conveyed and elicited prevents free substitution.

If I am reporting Stephen's belief to someone who is in the know about both of Clark Kent's names, then the same reasoning as above applies. My choice of names conveys information about how Stephen believes the proposition and so I need to choose my names accordingly.

Finally, in some cases we are not concerned with how a proposition is believed, but instead we are primarily concerned with retaining truth value. For example, when it comes to Zimmerman's childhood friend, then I can use either "Zimmerman" or "Dylan" when reporting her belief that the singer left his hometown at an early age. Assuming that I am reporting her belief to someone who only knows the singer by the name "Dylan," then I can substitute the name, thus adjusting to my audience, on the assumption that we are concerned with the proposition believed and not how it is believed.

As Kripke pointed out, simple sentences that contain different but codesignating names can raise different issues. He did not advocate or accept singular propositions as I do here, and he even speculated that we might have to give up the apparatus of propositions. My goal here has been twofold. First, to show that the Millian has resources that enable her to both accept singular propositions and explain our intuitions about the failure of substitutivity of coreferential names. Everyone, including the Millian, feels the intuitive pull against free substitutions when faced with sentences such as (3) and (4), and (5) and (6). Appealing to a three-place belief relation or to the hearer who is in the know about the relevant identity failing to draw all the inferences she should or could draw does not explain the extent of our intuitions. Simple sentences do not contain the belief relation that is present in attitude reports, and even when I am in the know about the relevant identity it seems evident that (5) and (6) can convey and elicit very different information. My second goal has been to provide an account that does not depend on a semantic solution to the problem of substitutivity. The solution I have provided does not depend on the meaning of names, as the same issues can be raised without the use of names, and the same explanations can be provided regardless of whether the issues are raised with names, picture, or scent.

Suggested readings

- Geirsson, Heimir. *Philosophy of Language and Webs of Information*, Routledge, 2013.
- Goodman, R., Genone, J., and Kroll, N. (eds). *Singular Thought & Mental Files*, Oxford University Press, 2020.
- Recanati, François. *Mental Files*, Oxford University Press, 2012.
- Salmon, Nathan. *Frege's Puzzle*, The MIT Press, 1986.
- Saul, Jennifer. *Simple Sentences, Substitution, and Intuitions*, Oxford University Press, 2007.

- Braun, David. "Understanding Belief Reports." *The Philosophical Review* 107 (1998): 555-95.
- Braun, David, and Jennifer Saul. "Simple Sentences, Substitutions and Mistaken Evaluations." *Philosophical Studies* 111 (2002): 1-41.
- Geirsson, Heimir. "Justification and Relative Apriority." *Ratio* 12 (1999): 148-61.
- . "Justification and Ways of Believing." *Disputatio* 12 (2002): 43-53.
- . *Philosophy of Language and Webs of Information*. New York: Routledge, 2013.
- Kripke, Saul. *Naming and Necessity*. Cambridge: Harvard University Press, 1980.
- Lawlor, Krista. *New Thoughts About Old Things: Cognitive Policies as the Ground of Singular Concepts*. Studies in Philosophy. Edited by Robert Nozick. New York: Garland Publishing, 2001.
- Perry, John. *The Problem of the Essential Indexical*. New York: Oxford University Press, 1993.
- Recanati, François. "Reference and Singular Thought." In *The Routledge Handbook of Linguistic Reference*, edited by Stephen Biggs and Heimir Geirsson: Routledge, 2020.
- Recanati, François. *Mental Files*. Oxford: Oxford University Press, 2012.
- Salmon, Nathan. *Frege's Puzzle*. Cambridge: The MIT Press, 1986.
- Saul, Jennifer. "The Pragmatics of Attitude Ascriptions." *Philosophical Studies* 92 (1998): 363-89.
- . *Simple Sentences, Substitution, and Intuitions*. Oxford: Oxford University Press, 2007.
- Soames, Scott. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. New York: Oxford University Press, 2002.

¹ Granted, some argue that pictures can have semantic properties. However, the case for that is far from settled and, at this time, I think we should be open to the idea that pictures do not have semantic properties. Additionally, one can construct examples involving Frege puzzle cases that are based on smell and not pictures and, arguably, smell does not have semantic properties.

² Note that I do not discuss knowledge context in this paper. I have argued elsewhere that while one can plausibly claim that if one believes that Clark Kent is Clark Kent, then one believes that Clark Kent is Superman, since the two sentences express the same proposition, it doesn't follow that if one knows the trivial identity statement one also knows that Clark Kent is Superman. Knowledge brings with it justification, and one can be justified in believing a proposition when apprehended in one way and not so justified when it is apprehended in a different way. While the Babylonian astronomers knew that Hesperus is Hesperus, it took them centuries to acquire the justification for believing that Hesperus is Phosphorus. For more on this, see Heimir Geirsson, "Justification and Ways of Believing," *Disputatio* 12 (2002); *Philosophy of Language and Webs of Information* (New York: Routledge, 2013).

³ Jennifer Saul, "The Pragmatics of Attitude Ascriptions," *Philosophical Studies* 92 (1998); David Braun, "Understanding Belief Reports," *The Philosophical Review* 107 (1998); David Braun and Jennifer Saul, "Simple Sentences, Substitutions and Mistaken Evaluations," *Philosophical Studies* 111 (2002).

⁴ See their "Simple Sentences, Substitutions and Mistaken Evaluations." Saul continues to develop this line in Jennifer Saul, *Simple Sentences, Substitution, and Intuitions* (Oxford: Oxford University Press, 2007).

⁵ Braun and Saul, "Simple Sentences, Substitutions and Mistaken Evaluations," 17.

⁶ Saul, *Simple Sentences, Substitution, and Intuitions*, 128-30.

⁷ *Ibid.*, 136.

⁸ Ibid., 137.

⁹ Ibid., 145.

¹⁰ Ibid., 143.

¹¹ The main proponent of the pragmatic implicature account is Nathan Salmon. See, e.g., his Nathan Salmon, *Frege's Puzzle* (Cambridge: The MIT Press, 1986). The account is further developed by Scott Soames in Scott Soames, *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity* (New York: Oxford University Press, 2002).

¹² The idea that we organize information in files or pools is a familiar one. It has been used in various forms in, to name a few, e.g., John Perry, *The Problem of the Essential Indexical* (New York: Oxford University Press, 1993); Krista Lawlor, *New Thoughts About Old Things: Cognitive Policies as the Ground of Singular Concepts*, ed. Robert Nozick, *Studies in Philosophy* (New York: Garland Publishing, 2001); Heimir Geirsson, "Justification and Relative Apriority," *Ratio* 12 (1999); "Justification and Ways of Believing.," *Philosophy of Language and Webs of Information*; Braun and Saul, "Simple Sentences, Substitutions and Mistaken Evaluations.," Saul, *Simple Sentences, Substitution, and Intuitions*; François Recanati, *Mental Files* (Oxford: Oxford University Press, 2012).

¹³ Compare this with Recanati's files. In "Reference and Singular Thought" he describes files as containing descriptions. François Recanati, "Reference and Singular Thought," in *The Routledge Handbook of Linguistic Reference*, ed. Stephen Biggs and Heimir Geirsson (Routledge, 2020). In *Mental Files* he writes "The role of a mental file based on a certain acquaintance relation is to store information acquired in virtue of that relation. The information in question need not be veridical; we can think of it in terms, simply, of a list of predicates which the subject takes the referent to satisfy." Recanati, *Mental Files*, 37-38. Italics are mine.

¹⁴ Compare this with what Kripke writes in the introduction to *Naming and Necessity*: "My view that the English sentence 'Hesperus is Phosphorus' could sometimes be used to raise an empirical issue while 'Hesperus is Hesperus' could not shows that I do not treat the sentences as completely interchangeable." Saul Kripke, *Naming and Necessity* (Cambridge: Harvard University Press, 1980), 20.

¹⁵ Had I informed her of Superman saving the city, then she would have been justified in believing that Superman saved the city. I have argued elsewhere that the justification for believing a singular proposition is crucially tied to *how* one believes it, and so given how the proposition is presented one cannot assume that when one is justified in believing a proposition is one way, one is thereby justified in believing the proposition *simpliciter*.