

HOW PROPER NAMES REFER

Abstract – This paper develops a new account of reference-fixing for proper names. The account is built around an intuitive claim about reference-fixing: the claim that I am a participant in a practice of using α to refer to ϱ only if my uses of α are constrained by the representationally relevant ways it is possible for ϱ to behave. §1 raises examples that suggest that a right account of how proper names refer should incorporate this claim. §2 provides such an account.

Author's contact details: Philosophy Department, University of Toronto, 170 St. George St., Toronto M5R 2M8, Ontario, Canada; imogen.dickie@utoronto.ca

Acknowledgements:

Thanks to audiences at York University, Toronto; the 2008 meeting of the Society for Exact Philosophy; the University of California at Berkeley; the University of Southern California; my 2009 graduate class at New York University; the Jowett Society at Oxford; Kevan Edwards's 2010 graduate class at Syracuse University; and the Aristotelian Society. For conversations and correspondence that led to major improvements in the paper, particular thanks to Troy Cross; Gary Ebbs; Henry Jackman; Kirk Ludwig; Michael Nelson; Lewis Powell; Jim Pryor; Stephen Schiffer; James Shaw; Scott Soames; David Velleman; and Daniel Velleman. Much of my work towards this paper was supported by a Bersoff Faculty Research Fellowship at New York University

How Proper Names Refer

Imogen Dickie

This paper proposes an account of reference-fixing for proper names. The account is intended as a successor to the view that Gareth Evans proposed in ‘The Causal Theory of Names’ and revised in *The Varieties of Reference*¹. Evans argued that the bearer of a proper name, as used by a community of speakers, is the dominant causal source of the information that members of the community associate with the name. I think Evans’s suggestion was a major advance. But I am going to argue that it fails, and that reflection on why it fails opens the way for a new and better account.

Let me start by making the question I am concerned with more precise, and distinguishing it from a different but currently more widely discussed question about proper names. The question I am concerned with is

The reference-fixing question – How is it determined which object (if any) is the bearer of a proper name as used by a community of speakers?

This question needs clarification at the following two points. What is a proper name? And what is it for an object to be a name’s ‘bearer’?

I propose answering these questions in terms of one another. I take ‘proper name’ to be a semantic natural kind term. Exemplars of the kind include ordinary proper names of people (‘Bertrand Russell’), places (‘Manhattan’), and things (‘HMS Victory’). The kind includes these exemplars, and expressions relevantly similar to them. (Compare “‘Water’ is stuff relevantly similar to *these* samples.’) I take it that proper names have a standard or default semantic role – a role they can be used to perform without any extra stage-setting. A proper name’s default semantic role is to introduce its bearer as the object the speaker intends to be communicating about. And a proper name’s ‘bearer’ is the object it introduces when it performs its default semantic role. ‘Relevant similarities’ for the purposes of delineating the class of proper names are similarities with respect to default semantic role, and with respect to how the name’s bearer is determined. So α is a proper name if and only if its default semantic role is to introduce its bearer as the object the speaker intends to communicate about, and its bearer is determined in the same way as bearers are determined for ordinary proper names of people, places, and things.

Note that some expressions that look syntactically like proper names do not belong to the semantic natural kind I have described. For example, ‘John Bull’, used as an abbreviation for ‘English men of a traditional red-faced variety’ is not a proper name.² Similarly, some expressions that look syntactically like typical members of some other category are in fact proper names. For example, ‘The North Island’, and ‘The Black Country’ are proper names that look syntactically like descriptions.

¹ For recent work endorsing something like Evans’s view see (Campbell 2002, pp. 38-40, 95-96, 100-101); (Recanati 1993, ch. 10); (Reimer 2010, §2.3). Field (1994, p. 118) says that an Evans-type view (with its appeal to ‘dominant causal sources’) is ‘the positive view of reference’ which emerges from ‘the writings of Kripke, Putnam, etc.’.

² Similarly, given that ordinary proper names do not have their references fixed by initial descriptive stipulations, ‘descriptive names’ (Evans 1982, p. 31) like ‘Jack the Ripper’ and ‘Deep Throat’ are not proper names.

Note also that the claim that a proper name has a default semantic role (to introduce its bearer as the object the speaker intends to communicate about) allows for the fact that proper names can be used in other ways given suitable stage-setting. In particular, it allows that we can use proper names to communicate about objects other than their bearers. (X, who looks like Russell, is growing irate on the other side of the room. I say to you ‘Don’t look now, but Russell is getting testy’.) But I shall not be concerned here with exactly how we manage to use proper names to perform these other roles.³

So I am supposing that a proper name’s default use is to introduce a particular – its bearer – as the object the speaker intends to communicate about. And I am supposing that there will be a uniform account of how it is determined which particular this is. The reference-fixing question concerns how this uniform account should go. An exhaustive treatment of the question would defend the suppositions required to set it up.⁴ But to keep the paper to a reasonable length I shall follow other participants in the debate about reference-fixing for proper names in taking them for granted.

The nearby question from which I want to distinguish the reference-fixing question is

The formal semantic question – How should a proper name be treated by a semantic theory for our language (a theory stating how a sentence’s truth conditions depend on contributions made by its parts)?

We would have to go far afield to explore the relation between the formal semantic question and the reference-fixing question. However, what is important for my purposes here is just that the questions are distinct. The quickest way to bring out their distinctness is to consider some answers to the formal semantic question. Here are three examples:

The Metalinguistic Description Theory of the Semantic Value of a Proper Name – A semantic theory should treat ‘ α is F’ the way it treats the ‘The bearer of “ α ” is F’.⁵

Millianism – If \underline{o} is α ’s bearer, \underline{o} is also α ’s semantic value, (so a semantic theory’s theorems for sentences of form ‘ α is F’ should take form ‘‘ α is F’ is true if and only if \underline{o} satisfies F’).⁶

³ There is a question about whether in this kind of case I am using ‘Russell’ to express a proposition about Russell and am merely implicating a proposition about X, or using ‘Russell’ to express a proposition about X. Compare (Kripke 1977) especially pp. 262-264. Nothing I shall say in this paper demands taking sides on this question. In Kripke’s terms, my concern here is with what determines the ‘semantic referent’ of a proper name as used by a speaker or group of speakers. A token name’s ‘speaker’s referent’ is the object the speaker intends to be communicating about in using the name. The name’s ‘semantic referent’ is the object that would be introduced as speaker’s referent in the absence of any special stage-setting. In the example in the text, speaker’s referent diverges from semantic referent.

⁴ For a summary of the case against them see (Bach 2006) especially pp. 536-543.

⁵ For views of this general kind see (Katz 1994); (Bach 1981 and 2002); (Burge 1973).

⁶ Braun (2005 and 2006, §§1-2); Soames (2002, ch. 3).

A treatment of proper names in Generalised Quantifier Theory – If \underline{o} is α 's bearer, α 's semantic value is the set of sets containing \underline{o} (so a semantic theory will lay down that $\ulcorner \alpha \text{ is } F \urcorner$ is true if and only if the set of F things is an element of the set of sets containing \underline{o} '⁷).

Each of these answers to the formal semantic question supposes that there is some relation between α and \underline{o} in virtue of which \underline{o} is the bearer of α . So each supposes that there is an answer to the reference-fixing question. But none takes a stand on what this answer is.⁸

I do not think it is possible to address both the formal semantic question and the reference-fixing question in a single paper. So I shall have no further truck with the formal semantic question here.

The paper has two parts. §1 states Evans's answer to the reference-fixing question and says why I think it fails. §2 develops my alternative proposal.

§1 Evans's View and Why it Fails

Evans sets up his initial account of reference-fixing for proper names (in 'The Causal Theory of Names') as a response to Kripke's discussion in *Naming and Necessity*. Kripke had argued against the Sophisticated Description Theory and explored the possibility of a Causal Inheritance view:

The Sophisticated Description Theory of Reference Fixing – \underline{o} is the bearer of α in S's linguistic community if and only if \underline{o} is the unique satisfier of the most important descriptions most members of S's linguistic community associate with α ⁹.

The Causal Inheritance Theory of Reference Fixing – \underline{o} is the bearer of α in S's linguistic community at time t if and only if at some time at or before t, α was assigned as a name for \underline{o} , and the uses of α made by members of S's linguistic community at t descend from this initial 'baptism' by a causal chain whose links consist in a speaker's taking up α with the intention of using it to refer to the object it is being used to refer to already¹⁰.

Evans endorses Kripke's anti-descriptivist arguments. But he argues that the Causal Inheritance Theory will not do either. He claims that his own view captures what is right, while discarding what is wrong, about each kind of view. I shall summarise Evans's

⁷ So, for example, 'Harry sneezed' is true if and only if one of the sets of which Harry is a member is the set of sneezers. See (Barwise and Cooper 1981, §1.6).

⁸ Compare Braun's comments on the distinction between the reference-fixing question and the formal semantic question at (2005, pp. 12-13).

⁹ See (Kripke 1977, p. 31) for a characterisation. See (Strawson 1959, ch. 1) and (Searle 1958) for Sophisticated Descriptivist proposals. Note that Searle's proposal does not require a precise account of 'weighting' or 'importance': he says that a proper name's bearer must satisfy a sufficient number of the descriptive conditions associated with it, but there is no non-arbitrary account of which (1958, p. 171).

¹⁰ (Kripke 1977) – see pp. 91-92 and 96-97 for summaries, though note Kripke's claim on p. 97 that he is trying to present only a 'picture' of how reference fixing works, and not a complete 'theory' (set of necessary and sufficient conditions).

discussion in just enough detail to bring out how his account is motivated and how it is supposed to work.

Kripke had discussed examples like¹¹

Case 1: 'Gödel'

The only contender to be a description associated with the name 'Gödel' by most people who understand it is something like 'the prover of certain incompleteness theorems'. But now suppose that Gödel in fact did not prove these theorems. The proofs were discovered by Gödel's rival, Schmidt. Gödel found Schmidt's proofs and published them as his own.

If the Sophisticated Description Theory were right, we would conclude that the referent of 'Gödel' is the actual prover of the theorems (Schmidt). But this result runs counter to ordinary speaker intuition. If we discovered that Schmidt proved the theorems we would not say "'Gödel'" refers to a man known throughout his lifetime as "Schmidt". We would say 'Gödel did not prove those theorems after all'. Evans (following Kripke) concludes that the Sophisticated Description Theory is wrong: ordinary speaker intuitions about which objects proper names refer to are data that a theory of reference-fixing is supposed to explain; the Sophisticated Description Theory should be rejected because it generates predictions contrary to these data.¹²

Evans's argument against the Causal Inheritance Theory is built around examples involving reference change. Consider

Case 2: 'Madagascar'

'Madagascar' is a corruption of a name initially used to refer to a part of continental Africa. Marco Polo intended to use the name with the same referent as it was used by the Arab sailors from whom he heard it. But he applied it instead to the great African island (writing 'Madagascar is F' and 'Madagascar is G' in his reports of what the island is like).

According to the Causal Inheritance Theory, if S is a speaker removed from the initial baptism, S's use of α refers to \underline{a} if and only if S inherited the practice of using α from S*; S*'s uses of α refer to \underline{a} ; and S's intention in picking up the use of α was to use it the way it was already being used by S*. The speakers with whose use of 'Madagascar' Marco Polo intended to keep faith used it to refer to a part of continental Africa. So, assuming the Causal Inheritance Theory, Marco Polo's uses of the name also referred to this part of the continental Africa. We have inherited the use of 'Madagascar' from Marco Polo through a chain of transactions in which speakers taking up the name intend to go on using it the way it is already being used. Therefore, according to the Causal Inheritance Theory, our uses of 'Madagascar' refer to what Marco Polo's did: they refer to a part of continental Africa. Therefore, for example, our utterances of 'Madagascar is an island' are false. But this result runs counter to ordinary speaker intuition. Intuitively, our uses of 'Madagascar' refer to the island. So there has been an unintended change in which thing the name refers to. The Causal

¹¹ For the purposes of this paper I follow Evans in discussing only Kripke's arguments against description theories of reference-fixing. Kripke's 'modal' arguments are directed against the claim that proper names are synonymous with descriptions, and so against a specific kind of description theory of semantic value.

¹² (Evans 1973, p. 3); compare (Kripke 1977, pp. 83-85).

Inheritance Theory cannot explain how this kind of change is possible. Evans concludes that it must be rejected.¹³

Evans sets up his own account of reference-fixing for proper names against the background of an intuitive, and now widespread, view of how a system of beliefs about particular things is structured, and the role names play in such a system. This is the view that a system of beliefs is structured like a filing system, with names acting as labels on files. Each file is a cluster of information that the speaker takes to be about a single thing. A name plays a label-like role in determining how information entering the filing system is processed. For example, all the information I receive in the form of sentences containing the name 'George W. Bush' is sent to my 'George W. Bush' file. A name's bearer (if it has one) is the object (if there is one) that the file of beliefs the name labels is about.¹⁴

If this background view is assumed, the reference-fixing question becomes the question of how the files of beliefs speakers associate with α (speakers' ' α -files') must relate to \underline{o} if the beliefs in the files are to be about \underline{o} . Evans argues that we know that the required relation is not satisfactorial¹⁵: the anti-descriptivist arguments show that there is no descriptive condition contained in, or culled from, speakers' α -files whose role is to single \underline{o} out as the bearer of α . Yet, Evans argues, the relation must have something to do with the information in speakers' files. For the problem with the claim that we are all using 'Madagascar' to refer to a part of continental Africa must trace to the fact that the information in our 'Madagascar' files derives from what the island, not the part of the mainland, is like.¹⁶

Evans proposes that the required relation is a relation between an object and the information in the file, but it is a causal rather than a satisfactorial relation: a name's bearer (if it has one) is the dominant causal source of the information the associated files contain. So the suggestion is that the beliefs speakers would express using α do play a role in determining which object (if any) is the bearer of α in their linguistic community, but this role does not depend on whether these beliefs are true. It depends on how they are caused. The bearer of α in our linguistic community is the causal source of most of the most important beliefs that our α -files contain.¹⁷

To see Evans's view to best advantage, consider how it copes with the 'Madagascar' example. On Evans's view, the bearer of 'Madagascar' in our linguistic community is the dominant causal source of the beliefs contained in the community's 'Madagascar' files. The dominant causal source of contemporary 'Madagascar' beliefs is the island. Therefore, 'Madagascar' as used by us refers to the island, even though there is an unbroken causal chain, each new speaker intending to use the name as it is already being used, running from our uses of 'Madagascar' back to a time when the name was used to refer to something else.

¹³ (1973, p. 11).

¹⁴ (1973, pp. 14-16). Evans cites (Grice 1969) as his source for the notion of a 'file' or 'dossier'. The notion also played a central role in (Strawson 1974). For more recent uses of the 'files' view of how systems of beliefs about particulars work see (Recanati 1993); (Campbell 2002); (Jeshion 2001).

¹⁵ This term was introduced by Bach: (1981, p. 12).

¹⁶ (1973, pp. 12-13).

¹⁷ (1973, p.14), and compare (1973, p. 17): 'in general a speaker intends to refer to the item that is the dominant source of his associated body of information'.

So Evans's view can deliver what his motivating example demands: an account of how a name's referent might change despite speakers' intentions to go on using it the same way.

A final qualification allows for deferential cases. S is a 'deferential participant' in a proper-name-using practice if and only if S uses α with the 'overriding intention' of talking about whatever is the bearer of α as used by some group of other speakers. In this kind of case, Evans suggests, S counts as using α with the same bearer as it has in the mouths of the speakers to whom S intends to defer.¹⁸

Evans's main problem in precisifying this view was to explain the notion of 'causal source'. The view needs an account of 'causal source' according to which \underline{Q} may be the causal source of the belief I would express by saying ' α is F' even though \underline{Q} is not F – otherwise the view is just descriptivist. In 'The Causal Theory of Names', Evans proposed a reliabilist account. He said that \underline{Q} is the 'causal source' of the belief S would express by saying ' α is F' if and only if the belief is caused by S's standing in a relation to \underline{Q} of a type such that, in general, standing in this relation to a thing puts you in a position to know that the thing is F¹⁹. For example, suppose I form the belief I would express by saying 'Fred is in my fourth year seminar' as a result of seeing Fred sitting among the students during the seminar. Seeing someone sitting among the students during a seminar is, in general, a way to gain knowledge that this person is in the seminar. So, according to Evans's proposed criterion, Fred is the causal source of the belief. But a relation that generally puts you in a position to know that the thing to which you stand in the relation is F might fail to put you in this position in a specific instance. For example, it may be that Fred was just too lazy to leave the room at the end of the previous class.

However, this reliabilist account of the notion of 'causal source' cannot perform the role Evans needs it to. For consider

Case 3: The astrologer

It is a time of faith in astrology. An astrologer makes a series of predictions about a small child: 'She will free us from tyranny'; 'She will build a great city'; and so on. These predictions are widely disseminated, and are accepted by whoever hears them. Other details about the child ('She is two feet tall'; 'She is the child of X and Y') are left behind as irrelevant to the prophecy. In this way, many speakers in the community end up with files labelled by the child's name and containing only information derived from the astrologer's predictions.

Ordinary speaker intuition declares that speakers in this community are using the child's name to stand for the child. But Evans's 'The Causal Theory of Names' view does not generate this result. Believing the sayings of an astrologer is not, in general, a way to acquire knowledge. So the files that, for most speakers, are labelled by the child's name do not contain any information that Evans's initial account of 'causal source' treats as information causally 'of' or 'from' the child.

Case 3 would evaporate as a counterexample if we found a reason to say that the beliefs speakers would express using the child's name are, after all, formed by a reliable method. But it is hard to see what this method could be. Believing whatever you hear is not a

¹⁸ (1973, p. 21).

¹⁹ See (1973, p. 15): 'X is the [causal] source of the belief S expresses by uttering "Fa" if there was an episode which caused S's belief in which X and S were causally related in a type of situation apt for producing knowledge that something F's ... a type of situation in which the belief that something F's would be caused by something's F-ing.' The context indicates that Evans means 'if and only if'.

reliable method. Believing what you hear through generally reliable channels is a reliable method. But a channel originating in an astrologer is not generally reliable.²⁰

In *The Varieties of Reference*, Evans presents a revised version of his ‘dominant causal source’ view from which the initial reliabilist element has been excised. The new view centres on a distinction between roles that participants in a name-using practice might play: the distinction between what Evans calls ‘producers’ and ‘consumers’ in the practice of using α to refer to \underline{Q} .

A ‘producer’ in the practice of using α to refer to \underline{Q} is somebody who, as Evans says, ‘knows \underline{Q} as α ’. Evans intends a special sense of ‘know as’, which I shall mark by italicising ‘as’. S ‘knows \underline{Q} as α ’ if and only if S has a specific kind of rapport with \underline{Q} , where the use of α forms part of this rapport: S has the capacity to identify \underline{Q} demonstratively and reidentify \underline{Q} after breaks in observation; S exercises this capacity from time to time; S uses α in storing information gained by interactions with \underline{Q} , keeping this information in a file on which α is the label; if \underline{Q} is a person or animal, S may use α in addressing \underline{Q} ; and S is very likely to use α in transmitting information about \underline{Q} to others (expecting that they have α -files too).²¹

A ‘consumer’ with respect to the practice of using α to refer to \underline{Q} is a participant in the practice who does not know \underline{Q} as α .²² Though Evans himself does not use these terms, his account of a name-using practice incorporates a distinction between what I shall call ‘participating’ and ‘parasitic’ consumers.

A ‘participating’ consumer in the practice of using α to refer to \underline{Q} is somebody who does not know \underline{Q} as α , but does have an α -file that stands in an appropriate (reference-fixing) relation to \underline{Q} . This notion of an appropriate relation is, as in Evans’s earlier view, cashed out in terms of the causal origins of the information in the file. But now the suggestion is that the α -file of somebody who does not know \underline{Q} as α is about \underline{Q} if and only if the information in the file is dominantly derived from the contents of the α -files of people who do (or did) know \underline{Q} as α . Evans puts this new suggestion in terms of the intention with which a participating consumer uses a name:

... it is reasonable to attribute to a speaker the intention to participate, by his use of a name, in the same practice as was being participated in by those speakers from whose use of the name the information he has associated with the name derives. (1982, p. 387)²³

A ‘parasitic’ consumer in the practice of using α to refer to \underline{Q} counts as using α to refer to \underline{Q} solely in virtue of intending to use α the way it is used by some group of speakers who are producers or participating consumers in the practice.²⁴ (So when you first pick up

²⁰ The next step in discussing how Evans’s (1973) view might be defended against apparent counterexamples like *Case 3* would be to consider special cases of objections to reliabilist views of knowledge or justification which focus on the notion of a ‘generally reliable method’. But Evans moved on from his early view, so I shall not dwell on this issue here.

²¹ Evans (1982, p. 376).

²² (1982, p. 377).

²³ Contrast this with (Evans 1973, p. 17): ‘in general a speaker intends to refer to the item that is the dominant source of his associated body of information’.

²⁴ This requires that the speaker have a way of identifying the group. But there are no strong constraints on which ways will do: I need not be able to pick the group out other than as the group of name-users to which people from whom I first heard the name belong. The

the use of a name for a thing with which you lack producers' rapport, and you associate little or no information with the name, you are a parasitic consumer. You may become a participating consumer or a producer over time.)

According to Evans's *The Varieties of Reference* view, \underline{a} is the bearer of my uses of α if and only if I am either a producer, a participating consumer, or a parasitic consumer in a practice of using α to refer to \underline{a} .²⁵

This revised view is a major improvement over the proposal it replaces.

The revised view still generates the result ordinary speaker intuition wants in the 'Madagascar' case: I am a participating consumer in the practice of using 'Madagascar' whose producers include Marco Polo and the people since him who have known the island \underline{as} 'Madagascar'. That is why my uses of 'Madagascar' refer to the island.

The revision has excised the reliabilist notion of a relation which 'generally results in knowledge'. So Evans's revised view is not subject to counterexamples like *Case 3: The Astrologer*.

Finally, the revised view recognises the different roles speakers play in establishing and perpetuating a name-using practice. In 'The Causal Theory of Names', Evans had recognised the possibility of deferential reference-fixing. But he had recognised only one kind of non-deferential case. He had treated the speakers his later view distinguishes as 'producers' and 'participating consumers' in a name-using practice as standing in the same kind of relation to the name's bearer (the relation of having a file whose contents are dominantly derived from the bearer). The revised view recognises what does seem to be a genuine asymmetry between the roles that producers and participating consumers play.

But it is not hard to find counterexamples which suggest that this revised view is still not getting things right. Consider

Case 4: 'Geoffrey Chaucer'

Chaucer lived from about 1343 to 1400. He was well known in his lifetime. But in the centuries after his death, for reasons to do with the invention of the printing press and Henry VIII's desire to create an English national literature, the pool of claims made using Chaucer's name became flooded with invented attributions of literary works to him, and fabrications about his life, ancestry, place of birth, and so on. As a result of this flood of invention, there was a period of several hundred years (ending with the 'purging of the apocrypha' in the nineteenth century) during which even Chaucer experts had 'Chaucer' files most of the information in which was derived from fabrications made long after Chaucer's death.²⁶

During the centuries before the purging of the apocrypha, even the most careful and knowledgeable users of 'Chaucer' had 'Chaucer' files most of the information in which was not derived from beliefs held by speakers who knew Chaucer \underline{as} 'Chaucer'. So according to Evans's *The Varieties of Reference* view, during these years 'Chaucer' was not being used as a name for Chaucer. But ordinary speaker intuition cries out against this result. In describing the years before the purging of the apocrypha we do not say 'For several hundred years

distinction between participating and parasitic consumers is based on Evans's distinction between 'understanding' and merely 'using' a name in (1982, ch. 11 §5).

²⁵ The framework allows for the possibility of mixed cases, in which the speaker's intention does not fit neatly under a single classification.

²⁶ For historical details see (Gillespie 2006). Alexandra Gillespie *Print Culture and the Medieval Author: Chaucer, Lydgate, and Their Books, 1473-1557* (Oxford: Oxford UP, 2006).

“Chaucer” was not used to refer to Chaucer.’ We say ‘For several hundred years most things that even the experts believed about Chaucer were pure invention.’

Case 4 suggests that the conditions generated by Evans’s revised view are in fact not necessary for reference. I might be a participating consumer in the practice of using α to refer to α even though the majority of the information in my α file is not derived from information that producers in the practice associate with the name.

Another family of counterexamples suggest that Evans’s conditions are not sufficient for reference either. Consider

Case 5: ‘Rio Ferdinand’

During a conversation about football teams and players, somebody tells me that Rio Ferdinand is based in Leeds and plays in white strip. This is the first time I have heard the name ‘Rio Ferdinand’. For no very good reason, I assume that ‘Rio Ferdinand’ is a name for a team (rather than a player). I follow news stories expressed using ‘Rio Ferdinand’, and after some years I have quite a rich ‘Rio Ferdinand’ file formed by incorporating what I pick up from various news sources in a way consistent with my belief that Rio Ferdinand is a football team. So, for example, my file contains the beliefs I would express by saying ‘Rio Ferdinand is moving to Manchester’; ‘Rio Ferdinand scored two goals on Saturday’; ‘Some of Rio Ferdinand were in London at noon on Friday and some were in Glasgow’; ‘Two thirds of Rio Ferdinand are vegetarian’; ‘Rio Ferdinand might merge with Crystal Palace’.

In this case the information in my ‘Rio Ferdinand’ file is derived from the beliefs of speakers who know a particular person as ‘Rio Ferdinand’. But intuition cries out that I am not using ‘Rio Ferdinand’ to refer to this person. My mistake is not a matter of thinking that a person (Rio Ferdinand) has a series of properties which can only be properties of a collection. It is a matter of thinking that a person’s name (‘Rio Ferdinand’) is a name for a collection that has such-and-such properties.

If you doubt this intuition, consider how I will discuss my mistake after its discovery. Note first that (1) would be perfectly ordinary in this case, while (2) (said without any special emphasis) would be very strange (or ‘marked’):

- (1) I used to think that ‘Rio Ferdinand’ was a name for a football team.
- (2) I used to think that Rio Ferdinand was going to merge with Crystal Palace.

The strangeness can be removed from (2) by adding emphasis:

- (3) I used to think that RIO FERDINAND was going to merge with Crystal Palace.

One function of added emphasis is to indicate that the stressed expression is occurring metalinguistically. This is not the only function. But, given the availability of the metalinguistic report, (1), and the markedness of (2), it at least seems fair to say that linguistic intuitions about how it would be appropriate to report my mistake support the referential intuition that before I discover the mistake I am not using ‘Rio Ferdinand’ to refer to Rio Ferdinand²⁷.

²⁷ There are further markedness intuitions that suggest the same conclusion. For example, consider

- (1) Rio Ferdinand was not sometimes the best team on the day. Rio Ferdinand is a player.
- (2) Rio Ferdinand won’t merge with Crystal Palace. Rio Ferdinand is a player.

Evans does not discuss examples like *Case 5* explicitly. But ‘The Causal Theory of Names’ does contain a suggestion as to how his view might be tinkered with to address them. The suggestion is that the account of reference-fixing be supplemented with a ‘degree of fit’ requirement, so that for the case of participating consumers we have

The bearer of α as used by S is the object known as α by speakers from whose uses of α the information in S’s α -file is dominantly derived, so long as this object fits (or approximately fits) whatever kind-specific information the file contains; if the fit requirement is not met S’s uses of α fail to refer.²⁸

But this amended version of Evans’s view is subject to a different kind of counterexample. For my α -file might be about \underline{o} even though the file contains kind information that \underline{o} does not even approximately fit. Consider

Case 6: ‘The Oracle’

The sayings of the Oracle at Delphi are decided upon by a committee of priestesses who have the practice among themselves of referring to the corporate decision-making body as ‘The Oracle’, and spread oracular sayings among ordinary people using sentences of form ‘The Oracle says that p’. A typical ordinary speaker’s file contains the beliefs she would express by saying ‘The Oracle is to be consulted at Delphi’; ‘The Oracle may not be consulted on a Tuesday’; ‘The Oracle predicted last year’s earthquake’; ‘The Oracle appreciates cash donations’. It also contains the belief she would express by saying ‘The Oracle is a god’. The fact that The Oracle is actually a committee is a secret unknown outside the committee itself.

In this case, most speakers’ ‘Oracle’ files contain kind information (‘is a god’) that is not even approximately correct. But, intuitively, these speakers are still using ‘The Oracle’ to refer to the actual Oracle (the committee of priestesses). Again, if you doubt the intuition, consider how the mistake would be reported. A speaker reporting the mistake might say

(4) I used to think that the Oracle was a god.

Unlike (2), (4) may be asserted without special emphasis. In the ‘Rio Ferdinand’ case, it seems that after the discovery of the mistake the beliefs the subject would have used ‘Rio

In (1), a positive polarity item (‘sometimes’) occurs within the scope of a negation operator. In (2), the second sentence contradicts a presupposition of the first (the presupposition that Rio Ferdinand is the kind of thing for which ‘merging’ is defined). These are characteristics of ‘metalinguistic’ negation contexts – contexts where the negated sentence is treated as mentioned rather than used, and the negation operator expresses dissatisfaction with this mentioned sentence. (So the first sentence in (1) is really an abbreviation for ‘The sentence “Rio Ferdinand was sometimes the best team on the day” is not true’.) See (Horn 2001); (Carston 1996). It would take us too far afield to discuss these data about ‘Rio Ferdinand’ in detail here.

²⁸ See (1973, p. 13): ‘...I think it is likely that there is [a degree of fit] requirement for referring. We learn, for example, from E. K. Chambers *Arthur of Britain* that Arthur had a son Anir “whom legend has perhaps confused with his burial place”. If ...[a] notion of reference fixing is such that those who said Anir was a burial place of Arthur might be denoting a person it seems it has little to commend it....But the existence or nature of this “degree of fit” requirement will not be something I shall be concerned with here.’

Ferdinand' to express before this discovery can be reported only metalinguistically. In the 'Oracle' case things seem to be different: the mistake did not generate reference failure; it was a matter of assigning a wrong kind to the thing referred to.

So the situation is this. Suppose that we leave aside the phenomenon of parasitic consumerism. (I take it that Evans is right that we sometimes function as parasitic consumers. But since a parasitic consumer's referential intention is just to use a name they way *they* do, it is always going to be possible to graft the claim that there are parasitic consumers onto whatever account of non-parasitic reference-fixing we come up with.) Once parasitic consumerism is set aside, Evans's revised account of reference-fixing for proper names divides into two parts. Firstly, there is an account of how a proper-name- using practice is established: the practice of using α to refer to \underline{o} is established by a core group of speakers who know \underline{o} *as* α . Secondly, there is an account of what is required for participation in a name-using practice by speakers not in this core group: I count as participating in the practice, established by some core group of speakers, of using α to refer to \underline{o} if and only if

- (i) I have an α -file most of the information in which is derived from beliefs that members of the core group would express, or would have expressed, by saying $\lceil \alpha \text{ is } F \rceil$, and
- (ii) \underline{o} fits (or approximately fits) any kind information that my α -file contains.

But in fact neither (i) nor (ii) seems to be necessary for participation in a name-using practice by a speaker outside the core group. There are cases (like *Case 2: 'Madagascar'*) where contamination by information not derived from producers' beliefs does interrupt transmission of a name-using practice. But there are also cases (like *Case 4: 'Chaucer'*) where name-using practices survive this kind of contamination. Similarly, there are cases (like *Case 5: 'Rio Ferdinand'*) where a mistake about kind blocks a speaker's participation in a name-using practice. But there are also cases (like *Case 6: 'The Oracle'*) where speakers are participating consumers in the practice of using α to refer to \underline{o} even though their α -files contain kind information that \underline{o} does not even approximately fit.

These results suggest the following diagnosis with respect to Evans's proposal. The notions of information 'of' or 'from' an object and 'degree of fit' do have something to do with the conditions for participation in the practice of using a proper name to refer to an object. Otherwise participation in such a practice could not be undermined by contamination with extraneous information, or blocked by radical classificatory error. But Evans is wrong to formulate conditions for participation in a name-using practice directly in terms of these notions. For contamination by extraneous information and radical classificatory error only sometimes interfere with participation in name-using practices. So where contamination or classificatory error does have this result, it must do so by undermining some other, more basic, as yet unidentified condition on participation in a name-using practice – a condition which is left undisturbed in the cases where contamination or classificatory error leaves participation in the practice intact.

The rest of the paper is about what this as yet unidentified condition might be.

§2 The Governance View of Reference Fixing for Proper Names

Consider how representation works in the following toy case.

You are using tokens that you move around on a map to represent cyclists as you describe the events of the day's race. You say '*This* is the Spanish rider; *this* is the Italian' and so on. Then you move the tokens along the route on the map, mirroring the ways you think the riders moved ('Now *this* rider breaks away; now *these two* give chase...') Your use of the tokens to represent the riders does not depend on whether you are getting the details of the race right. (Your friend might correct you: 'No – at that point the Italian was in the lead.')

It does depend on your deploying the tokens in ways that match the central ways it is possible for cyclists to behave. If you say 'Then *this* one went *this* way' (boring through a mountain-side) or '*This* one disappeared *here* and instantaneously reappeared *over here*', you are no longer using the tokens in ways that keep faith with the initial stipulations, so the anchor for the claim that you are using the tokens to stand for the cyclists falls away.

In this toy case representation works by what I shall call 'governance': use of the tokens to represent the cyclists requires deploying the tokens in ways that are governed by the ways it is possible for cyclists to behave.

I am going to argue that the notion of governance provides the condition on participation in name-using practices that we are looking for. My uses of α refer to \underline{o} only if they are governed by \underline{o} 's possible behaviour.²⁹

In a little more detail, what I want to suggest is this. Let us take over from Evans the notion of a 'producer' in a proper-name-using practice:

Definition: A 'producer' in a practice of using α to refer to \underline{o} is a speaker who knows \underline{o} as α .

Let us add the notion of an 'information channel' leading from one speaker's uses of a name back to another's:

Definition: There is an 'information channel' connecting S_1 's and S_n 's uses of α if and only if there is a chain of speakers leading from S_1 to S_n such that each speaker inherits information expressed using α from his or her predecessor, and passes information causally derived from this information to his or her successor.

And let us define 'transmission' of a name-using practice as follows:

Definition: An information channel with S_1 at its source and S_n at its outlet 'transmits' a proper-name-using practice if and only if, if S_1 is a producer in this practice, S_n is a participating consumer in the practice.

²⁹ The main historical antecedent for this kind of condition on reference is the account of the relation between a name and its bearer entailed by Wittgenstein's 'picture theory'. Wittgenstein claims that α is a name for \underline{o} if and only if for every possible sentence containing α there is a possible state of affairs containing \underline{o} whose structure matches the structure of the sentence, and for every possible state of affairs containing \underline{o} there is a possible sentence containing α whose structure matches the structure of the state of affairs (1961, §§ 2.1 – 2.225). The notion of governance I am going to develop is much weaker than Wittgenstein's notion. A full discussion of points of similarity and difference is not possible here.

Then I suggest that proper-name-using practices are established by producers' rapport, and are transmitted by some, but not all, information channels and networks of channels: an information channel, or network of channels, transmits a proper-name-using practice if and only if it transmits governance.

So I am going to suggest that my uses of α refer to \underline{a} if and only if either

- (i) I am a producer in a practice of using α to refer to \underline{a} ; or
- (ii) there is an information channel (or network of channels) leading from my α -file back to the α -files of speakers who know as α , and this channel (or network of channels) transmits governance; or
- (iii) I am a parasitic consumer in a practice of using α to refer to \underline{a} .

I shall call this view the 'Governance View' of reference-fixing for proper names. My aim in this section is to get the Governance View on the table.

I shall show how the Governance View copes with the examples from §1 as details emerge. But it helps with intuitive motivation to approach the Governance View with some idea of how the account of the examples will go. Very roughly, I am going to suggest that classificatory error or contamination by extraneous information blocks transmission of a proper-name-using practice if and only if it destroys governance. That is why contamination by extraneous information undermines reference in *Case 2: 'Madagascar'* but not *Case 4: 'Chaucer'*. In the 'Madagascar' case Marco Polo's mistake results in a switch in governance. In the 'Chaucer' case no such switch occurs. Similarly, that is why classificatory error sometimes does, and sometimes does not, generate reference failure. In *Case 5: 'Rio Ferdinand'*, my mistake about kind is central to the way I maintain the file of beliefs I would express using the name. So I end up going radically wrong in the way I maintain the file: I maintain it in a way appropriate to representing a collection (a football team) rather than an individual (a person). In the 'Oracle' case, in contrast, ordinary people's beliefs that the Oracle is a god are peripheral to the way they maintain their 'Oracle' files. The central beliefs and hypotheses that they would use 'The Oracle' to express all concern the accuracy of the Oracle's predictions, how the Oracle is to be consulted, and so on – beliefs and hypotheses that could be true of either a god or a committee. So their uses of the name keep faith with ways it is possible for the Oracle (a committee) to behave.

There are three questions that must be answered to get the basic details of the Governance View in place. Firstly, how is the relevant notion of 'possible behaviour' to be understood? Secondly, what does it take for a subject to use a name in a way that 'keeps faith with' an object's possible behaviour? Thirdly, how is governance (keeping faith with possible behaviour) established by producer's rapport and transmitted down information channels? I shall address each of these questions in turn.

This first question is 'How is the relevant notion of "possible behaviour" to be understood?' According to the Governance View, I am a producer or participating consumer in a practice of using α to refer to \underline{a} only if I use α in ways that keep faith with the ways it is possible for \underline{a} to behave. But there are different senses of 'way it is possible for a thing to behave'. For example, quite a small range of behaviours are 'possible' for my dog given his actual size and abilities (in this sense of 'possible', it is possible for him to swim 100m; it is not possible for him to swim the English Channel). A wider range of behaviours are 'possible' for him in that he is a dog and there is no reason why a dog should not do such a thing. There is a still wider range of behaviours that are 'possible' for him in that it would be consistent with the most fundamental laws of nature for him to behave in these ways. And

there is a range of behaviours wider again that are ‘possible’ for him in that they would not violate laws of logic. Which kind of ‘way it is possible for \underline{Q} to behave’ does the Governance View demand that our uses of names keep faith with?

As a first step towards an answer to this question, let us return to the toy example I used to introduce the notion of governance, and have a closer look at exactly how Representation as Governance works.

The toy example has the following structure.

(i) The representations are being used in a model-building task (you are using the little tokens in your attempt to construct a (dynamic) model of a day’s cycling race).

(ii) The uses made of each representation depend on (a) rules governing how the representation may be deployed (move the tokens in *these* ways but not *those*); and (b) specific input about the properties of represented entities (the French rider was *here* at t and *here* at $t + n$). The rules determine how the information associated with each representation will be marshalled (for example, they determine that a rider cannot be represented as *here* at t and *there* (ten miles away) a second later).

(iii) The rules for deploying tokens guarantee that if you fail at the task it will be the property information you feed into your model that is to blame, not the way of marshalling this information determined by your rules for deploying object representations. (If you end up with a wrong model of the day’s race, this will be because you fed in wrong information about who was where when, and how fast the riders went, not because you got this property information right but marshalled it inappropriately.)

I suggest that all cases of Representation as Governance share this structure. To use X as a governed representation of Y is to deploy X as a proxy for Y in a model-building task. It is to deploy X in such a way that the roles X can play in the model match the roles Y can play in the modelled state of affairs.

This characterisation of Representation as Governance entails that the relevant notion of ‘possible behaviour’ is task-relative. For it entails that which possibilities you need to match depends on the model-building task at hand. For example, in the initial cycling-race case we are supposing that you are using your model to describe the day’s actual events (‘The Italian was *here* at t .’) and engage in simple counterfactual reasoning (‘If the Spanish rider had opened up a bigger lead at *this* point the Italian would not have won’). In this case, to get the ‘object’ part of the model right you must deploy the tokens in ways that match the ways the corresponding objects behave across the relevant range of situations (the actual situation and the situations that are relevant to evaluating the counterfactuals.) But there are other model-building tasks in which you might be engaged. You might start out saying ‘Here are some cyclists’, and go on to tell a children’s story in which your cycling protagonists develop super-powers. In this case, you are envisaging a much wider range of possible behaviours for the objects you are talking about. And the constraints on how the little tokens representing cyclists may be deployed are (consequently) much looser than they are in the initial prosaic scenario.

Now let us return to the case of proper names. Given (i)-(iii), the Governance View entails that when we use proper names we are engaged in model-building tasks. I suggest that when we use proper names in assertoric sentences we are engaged in tasks of narrative construction. I construe ‘narrative’ broadly to include an account of a course of events up to a

time; counterfactual claims to the effect that the course of events would have turned out differently at t if things had been different at $t - n$; and claims about how the course of events can be expected to develop in the future. We construct different kinds of narrative. And differences between the kinds of narrative we construct generate differences in the range of a thing's possible behaviours with which you must be keeping faith to count as referring to it.

In many contexts, we are engaged in constructing what I shall call 'mundane narratives'. A 'mundane narrative' incorporates (a) an account of observable events involving ordinary middle-sized objects; (b) basic counterfactual reasoning about how the course of observable events might have gone differently; and (c) basic speculation about how the course of observable events might develop in the future. The task of constructing a mundane narrative determines a range of permissible uses for representations deployed in attempting it. If you are constructing a mundane narrative, you count as using X to represent Y only if the ways you deploy X match what I shall call Y 's 'mundane behavioural possibilities': observable behaviours possible for Y across the situations you are trying to model when you construct a mundane narrative (the actual situation; basic counterfactual situations; likely future situations).³⁰

Note that mundane behavioural possibilities are possible paths in evolution of a thing's observable properties. So I might use α as a name for \underline{o} in constructing my mundane narrative even though my α -file contains beliefs about non-observable properties that are wildly at odds with \underline{o} 's possible behaviour in mundane situations. For example, suppose that I believe in life after death. Then my α -file may contain the belief I would express by saying ' α (my dead enemy) is rejoicing in my current suffering'. If there is no life after death, this belief is at odds with the possible behaviours for a person in the actual world and nearby possible worlds. But my false belief does not infect my deployment of α in constructing an account of observable events. So it does not undermine reference.

Similarly, I might have completely false beliefs about why a thing I am referring to has the observable properties it does. I might believe that Dorothy (in fact a cleverly designed robot) is a living thing (a cat), and might deploy the name 'Dorothy' in constructing my mundane narrative in accordance with this belief. My mistake will not undermine my use of 'Dorothy' to refer to Dorothy as long as my use of 'Dorothy' continues to keep faith with Dorothy's observable behaviour across mundane situations. But a cat and a cleverly designed cat-robot do share most of their observable behaviours across mundane situations: these are behaviours to do with retaining a stable colour, size, and shape unless interfered with; moving the way a living thing moves; interacting with various objects in various ordinary observable ways; and so on. In most cases, a mistaken belief about a thing's hidden nature does not interfere with reference because it does not generate widespread mistakes about

³⁰ Suppose Lewis's account of counterfactual conditionals (according to which ' A counterfactually entails C ' is true if and only if C is true in all of the nearest worlds in which A is true – see (Lewis 1973, ch. 1). And let N be the set of sentences in my mundane narrative at t . Then the suggestion is that I am using X as a proxy for Y in my mundane narrative if and only if, whenever my rules for deploying X generate observable consequent C from N together with observable antecedent A , if C is not true in the nearest A -worlds, these worlds are not N worlds. (So that if I end up wrongly committed to ' A counterfactually entails C ', my background beliefs about the pattern of observable properties are to blame, not my rules for deploying X .)

ordinary observable behaviour. (The ‘most’ qualification will be cashed out in the course of explaining the notion of ‘keeping faith’ below.)

If this account of representationally relevant possibility is along the right lines, we should find examples showing that the range of possibilities with which use of a name for an object must keep faith depends on the narrative task in hand. And there are examples that at least lean in this direction. Consider

Case 7: ‘Balto’

In 1925 a dog-team relay delivered vaccine to the Alaskan town of Nome, averting an epidemic. Balto, the lead dog of one of the teams became the focus of celebration, resulting in bitter controversy about whether Balto should get all the attention. A careless historian stumbles on records of the controversy, and decides to write a paper defending Balto’s right to special recognition. The careless historian has not realised that Balto was a dog. The resulting paper contains the claims ‘Balto covered more miles in the relay than anyone else’; ‘Balto motivated everybody to carry on’; ‘Balto gave up his own food for his dogs so that they would have the strength to complete the trip’.

Case 8: ‘Diamond’

Newton had a dog called ‘Diamond’. The canon of sayings attributed to Newton includes both claims that Diamond could prove mathematical theorems (Newton liked to joke about this), and lamentations on the damage Diamond had done to Newton’s work (Diamond knocked over a lamp, causing a fire that destroyed a large body of experimental results). A bad scholar comes across some of these sayings, and writes a paper on Diamond’s unacknowledged importance to Newton’s scientific career. The bad scholar has not realised that Diamond was a dog. The resulting paper contains the claims ‘Diamond developed the mathematical methods Newton used in his early work’; ‘Diamond and Newton corresponded about inverse square laws’; ‘Diamond’s wrong suggestions about the nature of light held back Newton’s later work in optics’.

Now consider the following two claims (made by someone reporting the bad scholar’s and careless historian’s mistakes):

- (1) There was no Diamond – ‘Diamond’ was the name of Newton’s dog.
- (2) There was no Balto – ‘Balto’ was the name of the lead sled-dog.

I do not suggest that there is a clear-cut intuitive difference with respect to the appropriateness of (1) and (2). But I think there is at least a difference in degree: (1) is more appropriate than (2) is; we are more inclined to say that the ‘Diamond’ mistake generates reference failure than that the ‘Balto’ mistake does. And developing the examples in accordance with the mistakes increases this difference. Adding more details about observable physical behaviour in the ‘Balto’ case (‘Balto met Calvin Coolidge’; ‘Balto died in 1933’) does nothing to weaken the claim that the bad scholar’s uses of ‘Balto’ refer to Balto. But the more details about alleged intellectual interactions with Newton that we add to the ‘Diamond’ case, the better (1) looks as a report of the careless historian’s mistake.

Assuming the account of representationally relevant possibilities that I have proposed, this difference between Case 7 and Case 8 can be explained in terms of a difference between the tasks involved.

In Case 7, the careless historian is attempting a contribution to the mundane narrative about the sled relay. Given this narrative aim, use of ‘Balto’ to refer to Balto requires keeping faith with a dog’s mundane behaviour across the actual situation, nearby counterfactual situations, and likely future situations. But the laws governing mundane interactions for a dog are the laws governing possible observable movement and change for an ordinary middle-sized thing capable of self-initiated movement. And these are the same laws as govern

mundane interactions for a person. So from the point of view of modelling mundane history, deploying a dog's name in a way appropriate to a name of a person is not a mistake that destroys reference.

Now consider *Case 8: 'Diamond'*. The bad scholar in *Case 8* is not just constructing a mundane narrative. Mundane narratives include only observable physical events. The bad scholar is telling a story that includes other kinds of events (development of mathematical methods; intellectual collaboration). And dog's potential for participation in events of this kind is radically limited compared with a person's. So the bad scholar ends up deploying 'Diamond' in a way that does not keep faith with Diamond's representationally relevant possibilities given the task at hand. That is why there is at least some temptation to say that the bad scholar's uses of 'Diamond' do not refer to Diamond.

So here is what I suggest about the notion of 'possible behaviour' that is relevant to the Governance View. We use proper names in constructing narratives. Narrative construction is a model-building task. And a model building task determines a range of Y's possible behaviours with which your use of X must keep faith if you are to count as using X as a proxy for Y in building your model. In many cases, the model-building task in which we are engaged when we use ordinary proper names is the task of constructing a mundane narrative. This task determines an associated range of representationally relevant possible behaviours for an object: the object's mundane behavioural possibilities. In cases where the task is to construct a mundane narrative, use of α as a proper name for \underline{a} requires that you deploy α in ways that keep faith with this range of possible behaviours. But we also use proper names in constructing narratives of other kinds (as in the 'Diamond' example). And these other kinds of narrative determine their own ranges of representationally relevant possibility.

The second element of the Governance View that needs explanation is the notion of using a name in a way that 'keeps faith with' an object's representationally relevant behavioural possibilities. In the 'files' shorthand that I am taking over from Evans, the central use of a name is its use as a label on a file of information. So we are looking for an account of what it takes for the way a file of information develops over time to 'keep faith' with the representationally relevant ways an object might behave.

Note first that 'keeping faith' cannot be just a matter of luck. In terms of the cycling race example, it is not enough, to count as using a token to represent a rider, just that the ways you move the token happen to mirror ways a cyclist could move: this result could be achieved by tossing a coin to decide where to move the token next.

But nor can the relevant notion of 'keeping faith' demand an absolute match between ways a subject would be prepared to let a file develop and the object's representationally relevant possibilities. For we have seen that a file might be 'about' an object even though it contains kind information that the object does not even approximately match (recall *Case 6: 'The Oracle'*). And in this kind of case the ways the subject would be prepared to let the file develop do diverge from the representationally relevant possibilities for the thing the file is about.

To motivate the notion of 'keeping faith' that the Governance View requires, consider the following schematic account of how files actually develop. Suppose that you start supplying me with information expressed using α (a name I am hearing for the first time). I open an α -file. My default expectation is that you are a reliable source of information. So if you start by saying ' α is F', ' α is G', and ' α is H' then, provided that there is no overt contradiction between these three statements, the inaugural contents of my α -file will be 'is F', 'is G' and 'is H'. As you supply more information I build up a picture of

what the thing is like. I weigh prospective new additions against this picture. If an incoming piece of information does not cohere with the existing picture, I either revise the picture, or reject the incoming information. I also engage in internally driven updating, revising items of information to keep pace with changes that I expect the object the file is about to have undergone. And, given the picture I am building up, I have expectations about the file's future development (I anticipate that tomorrow I will be adding either 'is J' or 'is K' to my file; I do not anticipate that I will be adding 'is L').

This story about file development generates a distinction between what I shall call the 'developmental core' of a file and its 'developmental periphery':

Definition: (i) the 'developmental core' of a file at a time is the part of the file through which potential additions to the file are filtered; which is subject to internally-driven updating; and which informs speculation about likely future developments .

(ii) the 'developmental periphery' of a file at a time is the part of the file outside its developmental core.

The most straightforward examples of developmentally peripheral beliefs are beliefs we would express using sentences of form ' α is said by X to be F', where we have no opinion as to X's reliability. If my α -file contains 'is said by X to be F', where I have no opinion as to X's reliability, I should not allow 'is not said by X to be F' into the file. But I could allow 'is not F', or even 'is said by X not to be F'. Unless I have a belief about X's reliability, I could, in general, add 'is said by X to be F' to my α -file, or take it away, without incurring the obligation to make any widespread changes. In contrast, items of information like 'is square', 'is heavy', or 'is in Winnipeg' are not usually developmentally peripheral in this way. Adding this kind of information, or taking it away, will usually have widespread implications concerning how I maintain my file: implications for which other beliefs I will be prepared to add to the file, how I will update what the file already contains, and my expectations about future development.

Let us say

Definition: A file's 'core developmental path' between times t and t^* is the path traced by its developmental core over this interval.

Definition: the 'maintenance policies' which generate a file's core developmental path are the generalisations, laws, and rules of thumb that the subject employs in getting from a state of the developmental core at one time to the state of the developmental core at the next.

Then I suggest that the notion of 'keeping faith' that the Governance View requires should be understood like this:

Definition: S's α -file 'keeps faith with' \underline{o} 's representationally relevant possible behaviours across a time interval if and only if

(i) the path traced out by the file's developmental core across the interval matches \underline{o} 's representationally relevant possible behaviours, and

(ii) the paths which would be generated using the same maintenance policies as generate the core developmental path, and given the same kind of informational input, also match \underline{o} 's representationally relevant possible behaviours. (So, for example, if the core developmental path results from combining colour, shape, size, solidity, and degree of motion information

with maintenance policies M, this condition requires that keeping M constant and varying colour, shape, size, solidity and degree of motion information within a reasonable range of difference would not result in a path deviating from ϱ 's representationally relevant possibilities.)

Given (i), keeping faith depends on a file's developmental core, not its developmental periphery. Given (ii), keeping faith requires more than just a mere match between core developmental path and what is possible for an object. It requires a non-lucky match: a match that is 'safe'³¹ in that it is generated using policies that would not easily generate a mismatch.

Combining this account of 'keeping faith' with the account of representationally relevant possible behaviours from earlier in the section, we can now make the following diagnosis about the circumstances under which a classificatory error blocks participation in a name-using practice. A classificatory error blocks participation in a name-using practice if and only if it is developmentally central and generates (or could easily generate) a mismatch between the file's core developmental path and the developmental paths that are representationally relevant for the object given the model-constructing task at hand.

Here is how this diagnosis applies to the examples involving classificatory error that I have raised.

First consider *Case 4: 'Rio Ferdinand'*. Here, the belief I would express by saying 'Rio Ferdinand is a football team' determines how I interpret the information expressed using 'Rio Ferdinand' that I receive from my newspaper sources. It determines which internally-driven updatings of my 'Rio Ferdinand' file I go in for. And it determines which updatings of the file I expect to be making in the future. As a result, the core developmental path of my 'Rio Ferdinand' file does not match what is possible for an individual. It matches what is possible for a collection. So the mistake about kind is developmentally central. And it generates a mismatch between my file's core developmental path and Rio Ferdinand's representationally relevant possible behaviours. That is why reference fails.

Now consider *Case 6: 'The Oracle'*. In this case, ordinary people's 'Oracle' files contain things like 'is to be consulted by asking the priestesses at Delphi'; 'predicted x'; 'failed to predict y'. They also contain 'is a god'. But the 'god' information, though it is part of people's picture of what the Oracle is like, is not a part which plays a role in determining how incoming information is incorporated, how the picture is updated over time, or how people expect to be updating their files tomorrow. People maintain their 'Oracle' files in ways that are geared to the kinds of dealings with the Oracle that they have always engaged in: adding new predictions; adding new instructions for consultation; updating predictions and instructions for consistency; developing expectations according to past patterns of reliability. So ordinary people are employing file-maintaining policies that make no use of their belief that the Oracle is a god. The mistake about kind does not interfere with reference because it is developmentally peripheral.

Finally consider *Case 7: 'Balto'* and *Case 8: 'Diamond'*. In each case, the mistake about kind is developmentally central. But the cases differ with respect to whether the

³¹ This is the notion of 'safety' adapted to the case of knowledge-which. Knowledge-that epistemologists say that a subject's true belief that p is 'safe' if and only if the subject could not easily be wrong about p in a similar case (compare Williamson *Knowledge and its Limits* 123-125; 147). Similarly, a match between a file's core developmental path and an object's representationally relevant possible behaviours is 'safe' if and only if the policies that generate the match could not easily generate a mismatch.

resulting maintenance policies generate departures from representationally relevant possibility, given the narrative-constructing task at hand. In *Case 7*, where the task is to contribute to the mundane narrative, there is no such departure, so ‘Balto’ refers to Balto. In *Case 8*, the narrative task is different. So the same mistake results in reference failure.

The third element of the Governance View that requires explanation is how governance is established by producers’ rapport, and transmitted to non-producers by some but not all information channels.

Let us start with the question of how governance is established by producers’ rapport.

A full characterisation of producers’ rapport would be an account of the kinds of relations to an object that a group of speakers must have if they are to be in a position to introduce and establish among themselves a practice of using α as a name for \underline{o} . Exactly how this account should go is an empirical question (the question of under what circumstances we actually introduce proper names). It is not possible to address this question here. But I take it that any viable answer will involve the claim that producers in a practice of using α to refer to \underline{o} are able to identify \underline{o} demonstratively (and identify it *as* α); reidentify \underline{o} after breaks in observation (and reidentify it *as* α encountered again); and use α to communicate about \underline{o} among themselves. For these conditions are all met in the situations in which ordinary proper names of people, places, and things are introduced. (Recall that we are treating ‘proper name’ as a semantic natural kind term whose central exemplars are these ordinary cases.)

Each of these components of producers’ rapport – demonstrative identification; reidentification; communication – is in its own right a subject of discussion among philosophers working on reference to particulars. It is not possible to enter properly into these discussions here. Nor is it possible to consider the finer points of how producers’ rapport is built up from these, and perhaps other³², components. So it is not possible here to show exactly *how* producers’ rapport establishes governance.

But I think it is possible to show *that* it does. For consider what is involved in reidentifying a thing encountered at $t + n$ as the same one you encountered at t . In ordinary cases, we reidentify things on the basis of (i) their apparent observable properties on the occasions when we encounter them, and (ii) a calculation about how a thing with *these* observable properties at t is likely to appear at $t + n$. But you will be reliable about getting this calculation right for \underline{o} only if you are right about the kinds of changes in observable properties that \underline{o} is likely to undergo, which is to say, only if you are right about \underline{o} ’s possible behaviours over mundane situations. So, if I have the capacity to reidentify \underline{o} *as* α , I am already maintaining my α -file in a way that keeps faith with the possibilities for \underline{o} that are relevant to mundane narrative.

Similarly, consider what is required if speakers S_1 and S_2 are to use α to communicate about \underline{o} . Communication requires that each speaker be in a position to understand why the other takes various observations as evidence for claims of form ‘ α is F ’, and why the other acts as he or she does having accepted a claim of form ‘ α is F ’ as true.³³ But, in general, the

³² Jeshion (2009) argues that we can introduce α as a proper name for \underline{o} only where \underline{o} is an object that has ‘significance’ for us. Note that Jeshion does not think that introduction of a proper name requires demonstrative contact with the named object or the capacity to reidentify it after a break in observation.

³³ Here I am in step with the view of communication as involving rational engagement defended in, for example, (Dummett 1995, ch. 4); (Evans 1982, ch. 10); (Heck 1995); (Dickie and Rattan 2010).

conclusions you draw from evidence expressed using α will make sense to me only if my procedures for incorporating information into my α -file parallel yours. I will understand why you move from ' α is F' and ' α is G' to ' α is H' only if that move makes sense to me as well. And whether the move makes sense will depend on the maintenance policies I employ in maintaining the file. So if we are to use α to communicate about α , we will at least need maintenance policies that generate matching patterns of possible development.³⁴

So I take it that producers' rapport with an object does establish governance. There is a further, harder, question about exactly how it does. But, again, my aim in this paper is just to get the Governance View on the table. I shall not try to address this harder question here.

What about the transmission of governance from one speaker to another down an information channel, and from producers to participating consumers through a network of such channels? I suggest that transmission of governance be explained in terms of buffered counterfactual dependence. Roughly, governance is transmitted across the transaction by which S^* acquires the use of α from S if and only if S and S^* 's α -files both keep faith with α 's representationally relevant possible behaviours; the fact that S^* 's α -file meets this condition depends counterfactually on the fact that S 's does; and this counterfactual holds in virtue of S^* 's intention to go on using α to refer to the thing S uses it to refer to. More precisely, let us say

Definition: an information channel with the α -files of S_1, \dots, S_i (producers in a name using practice) at its source and S_n 's α -file at its outlet 'transmits governance' if and only if

- (i) S_1, \dots, S_i 's α -files and S_n 's α -file keep faith with the same range of representationally relevant possibilities, and
- (ii.a) the fact that S_n 's α -file keeps faith with this range of possibilities depends counterfactually on the fact S_1, \dots, S_i 's α -files do (so that both of 'If producers' files kept faith with range R, S_n 's file would keep faith with range R' and 'If producers' files did not keep faith with range R, S_n 's file would not keep faith with range R' are true),
- (ii.b) the counterfactual dependence at (ii.a) holds in virtue of the intention of each of the speakers whose uses of α contribute to the information channel to use α with the reference it has in the mouths of the speakers from whom he or she acquired it.

It is now possible to see why some cases of contamination by extraneous information undermine reference while others do not.

First consider *Case 4: 'Chaucer'*. Fabrications find their way into the 'Chaucer' stream of information when speakers accept the contents of fabricated claims as legitimate additions to their 'Chaucer' files. But potential additions to speakers' 'Chaucer' files are filtered through these files' developmental cores. And these developmental cores contain the information that Chaucer is a person. Because incoming information is held up against the claim that Chaucer is a person, the incoming fabrication must either fill some gap in the skeletal life of Chaucer that the recipient's file depicts, or displace some piece of information from the existing story. In either case, the fabricated content is added to the recipient's 'Chaucer' file without interrupting counterfactual dependence of the ways the recipient's 'Chaucer' file could develop on the developmental possibilities for producers' files. So the injection of fabrications does not undermine the transmission of governance. That is why,

³⁴ I show how perceptual demonstrative identification secures governance in (Dickie 2010) and (Dickie forthcoming).

before the purging of the apocrypha, people were still using ‘Chaucer’ to refer to Chaucer. Speakers with contaminated ‘Chaucer’ files were still participating consumers in the initial name-using practice.

Now consider *Case 2: ‘Madagascar’*. There is an information channel leading from our uses of ‘Madagascar’ back to producers in the practice of using ‘Madagascar’ to refer to a part of continental Africa. But this channel does not transmit governance. Rather, our ‘Madagascar’ files inherit their possibilities for deployment from producers for whom the thing known *as* ‘Madagascar’ is the island, not the part of continental Africa. So Marco’s Polo’s mistake has brought about reference change by switching the source of governance.

Note that the point here is not that reference has switched because we maintain our ‘Madagascar’ files as ‘island’ files while the initial producers’ ‘Madagascar’-files were ‘part-of-a-larger-land-mass’-files. Rather, it is that reference has switched because buffered counterfactual dependence has been broken. To consolidate this point, let us consider a final example:

Case 9: ‘Lollius’

Lollius was a Roman boy to whom Horace wrote advising him to read Homer. The start of the relevant passage should be translated as something like ‘Lollius, while you are practising rhetoric in Rome, I am re-reading the greatest of writers of the Trojan Wars’. A widespread medieval mistranslation read ‘While you are practising rhetoric at Rome, I am re-reading Lollius, the greatest writer of the Trojan Wars’. This mistranslation caused many authors to use ‘Lollius’ in naming alleged sources for Trojan War narratives. For example, Chaucer credited Lollius as the source for his ‘Troilus and Cressida’.

Intuition declares that Chaucer and other writers who shared his mistake were not using ‘Lollius’ to refer to the Roman boy. And the Governance View explains why we make this intuitive judgement. Chaucer and those who shared his mistake were using ‘Lollius’ in a way that matched the developmental possibilities for the ‘Lollius’ files of the initial producers (the people who know the Roman boy, Lollius, *as* ‘Lollius’). So the first condition for the transmission of governance ((i) in the definition on the previous page) is met. But the second condition (buffered counterfactual dependence) is not. Let S_{trans} be the initial perpetrator of the mistranslation. S_{trans} acquired the use of the name ‘Lollius’ from Horace. But consider the counterfactuals

(Horace’s ‘Lollius’-file is a ‘person’-file) $\square \rightarrow$ (S_{trans} ’s ‘Lollius’-file is a ‘person’-file.)

and

(Horace’s ‘Lollius’-file is not a ‘person’-file) $\square \rightarrow$ (S_{trans} ’s ‘Lollius’-file is not a ‘person’-file.)

Whether these counterfactuals are true depends on the exact nature of the mistake involved in the mistranslation. One possibility is that S_{trans} was so incompetent that it was a mere matter of chance that he translated ‘Lollius’ as a name for a person (rather than for example, a name for a battle). In this case the counterfactuals are false. Alternatively, it may be that the textual clues (capitalisation; sentence position) were such that S_{trans} was not in danger of translating ‘Lollius’ as anything except a name for a person. In this case, the counterfactuals are true. But they are not true for the right reason. They are true in virtue of S_{trans} ’s intention to translate ‘Lollius’ in the way required by category markers in the text, not in virtue of S_{trans} ’s

intention to go on using ‘Lollius’ to refer to the referent it has for Horace. In either case, the mistranslation interrupts reference because it interrupts the transmission of governance, even though it does not introduce a mistake about kind.

§3 Conclusion – comparison with Evans

I shall close by summarising the similarities and differences between Evans’s view and the view I have proposed.

I take it that Evans is right that our proper-name using practices have an essentially asymmetric structure – a structure which assigns a privileged role to some core group of speakers. A proper-name-using practice is set up when a group of speakers who have a characteristic kind of rapport with an object begin to use the name to collect together the information about the object that they have in virtue of this rapport. These speakers are the producers in the practice of using the name to refer to the object. The practice of using α to refer to \underline{q} spreads beyond the immediate community of producers if and only if non-producing speakers (speakers who do not have rapport with the object) come to have α -files which are appropriately related to producers’ α -files.

But I think Evans is wrong about why producers’ rapport is important in setting up a name-using practice, and wrong about what it takes to be a consumer in such a practice. Evans thinks that the role of producers’ rapport is to establish a bank of information derived from encounters with \underline{q} and expressed using α . And he thinks that I am a participating consumer in the practice of using α to refer to \underline{q} if and only if the contents of my α -file are dominantly derived from this initial bank of information. In contrast, I have suggested that the role of producers’ rapport is to secure governance: the nature of producers’ rapport (‘knowing *as*’) is such that if you know \underline{q} *as* α , you will maintain your α -file in a way that keeps faith with \underline{q} ’s representationally relevant possible behaviours. And I have suggested that the ‘appropriate relation’ between a consumer’s α -file and producers’ α -files does not involve inheritance of information. It involves inheritance of governance. I am a participating consumer in the practice of using α to refer to \underline{q} if and only if my α -file inherits its possibilities for deployment from the α -files of producers in the practice.

References

- Bach, K. 1981: ‘What’s in a Name?’ *Australasian Journal of Philosophy*, 59, pp. 371-386.
- Bach, K. 1994: *Thought and Reference*. Oxford: Oxford University Press.
- Bach, K. 2002: ‘Giorgione was So-Called Because of his Name’. *Philosophical Perspectives*, 16, pp. 73-103.
- Bach, K. 2006: ‘What Does It Take To Refer?’ In E. Lepore and B. Smith (eds.) *The Oxford Handbook of Philosophy of Language*. Oxford: Oxford University Press, pp. 516-554
- Barwise, J. and R. Cooper. 1981: ‘Generalised Quantifiers and Natural Language’. *Linguistics and Philosophy*, 4, pp. 159-219.
- Braun, D. 2005: ‘Empty Names, Mythical Names, Fictional names’. *Nous*, 39, 596-631.

- Braun, D. 2006: 'Names and Natural Kind Terms'. In E. Lepore and B. Smith (eds.) *Handbook of Philosophy of Language* Oxford: Oxford University Press, pp. 490-515.
- Burge, T. 1973: 'Reference and Proper Names'. *Journal of Philosophy*, 70, 425-439.
- Campbell, J. 2002: *Reference and Consciousness*. Oxford: Oxford University Press.
- Carston, R. 1996: 'Metalinguistic Negation and Echoic Use'. *Journal of Pragmatics*, 25, pp. 309-330.
- Dickie, I. 2010a: 'We Are Acquainted With Ordinary Things'. In Robin Jeshion (ed.) *New Essays on Singular Thought*. Oxford: Oxford University Press, pp. 213-245.
- Dickie, I. (forthcoming): 'Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck'. To appear in C. Mole, D. Smithies, and W. Wu (eds.), *Attention: Philosophical and Psychological Essays*, forthcoming from Oxford University Press.
- Dickie, I. and G. Rattan. 2010: 'Sense, Communication, and Rational Engagement'. *Dialectica*, 64, pp. 131-151.
- Dummett, M. 1995: *The Logical Basis of Metaphysics*. London: Duckworth.
- Evans, G. 1973: 'The Causal Theory of Names'. Reprinted in Gareth Evans *Collected Papers*. Oxford: Oxford University Press, 1985, pp. 1-24.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.
- Field, H. 1994: 'Deflationism about Meaning and Content'. Reprinted in Hartry Field *Truth and the Absence of Fact*. Oxford: Oxford University Press, 2001, pp. 104-140.
- Gillespie, A. 2006: *Print Culture and the Medieval Author: Chaucer, Lydgate, and Their Books, 1473-1557*. Oxford: Oxford University Press.
- Grice, H. P. 1969: 'Vacuous Names'. In Donald Davidson and Jaakko Hintikka (eds.) *Word and Objections*. Dordrecht: D. Reidel, 1969, pp. 118-145.
- Heck, R. 1995: 'The Sense of Communication'. *Mind*, 104, pp. 79-106.
- Horn, L. 2001: *A Natural History of Negation*. Reissue Edition. Stanford. CSLI Publications, 2001.
- Jeshion, R. 2001: 'Acquaintanceless *De Re* Belief'. In J. Campbell, M. O'Rourke and D. Shier (eds.) *Topics in Contemporary Philosophy: Truth and Meaning*. New York: Seven Bridges Press, 2001, pp. 53-78.
- Jeshion, R. 2009: 'The Significance of Names'. *Mind and Language*, 24, 370-403.
- Katz, J. 1994: 'Names without Bearers'. *Philosophical Review*, 103, pp. 1-39.
- Kripke, S. 1972: *Naming and Necessity*. Oxford: Blackwell.
- Kripke, S. 1977: 'Speaker's Reference and Semantic Reference'. *Midwest Studies in Philosophy II*, pp. 255-276.
- Lewis, D. 1973: *Counterfactuals*. Oxford: Blackwell.
- Recanati, F. 1993: *Direct Reference*. Oxford: Blackwell.
- Reimer, M. 'Reference'. *Stanford Encyclopaedia of Philosophy (Spring 2010 Edition)*.
- Russell, B. 1918. 'The Philosophy of Logical Atomism'. Reprinted in Russell 1956, pp. 177-281.

- Russell, B. 1905. 'On Denoting'. Reprinted in Russell 1956, pp. 39-56.
- Russell, B. 1956: *Logic and Knowledge*. Ed. R. C. Marsh. London: MacMillan.
- Searle, J. 1958: 'Proper Names'. *Mind*, 67, pp. 166-173.
- Soames, S. 2002: *Beyond Rigidity*. Oxford: Oxford University Press.
- Strawson. P. F. 1959: *Individuals*. London: Routledge.
- Strawson. P. F. 1974: *Subject and Predicate in Logic and Grammar*. London: Methuen..
- Williamson, T. 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. 1961: *Tractatus Logico-Philosophicus*. Trans. David Pears and Brian McGuinness. London: Routledge. 1961.