Arbitrary Reference

Wylie Breckenridge and Ofra Magidor

(Penultimate draft of paper to appear in *Philosophical Studies*
Please cite the final published version)

Two fundamental rules of reasoning are Universal Generalisation (UG) and Existential Instantiation (EI). The former is the rule that allows us, given that we stipulate that ABC is an arbitrary triangle and prove that ABC has a property p, to conclude that all triangles have p. The latter is the rule that allows us, given that there exists a continuous function, to stipulate that *f* is an (arbitrary) continuous function. Reasoning according to these rules (or, as Kit Fine calls it, 'instantial reasoning') is ubiquitous in both formal and informal contexts.

Yet applications of these rules involve stipulations (even if only implicitly) such as 'Let ABC be an arbitrary triangle', 'Let *n* be an arbitrary number', or 'Let John be an arbitrary Frenchman'.[1] And the semantics underlying such stipulations are far from clear. What, for example, does '*n*' refer to following the stipulation that *n* be an arbitrary number? In this paper, we argue that '*n*' refers to a number - an ordinary, particular number such as 58 or 2,345,043. Which one? We do not and cannot know, because the reference of 'n' is fixed *arbitrarily*.[2] Underlying this proposal is a more general thesis:

Arbitrary Reference (AR): It is possible to fix the reference of an expression arbitrarily. When we do so, the expression receives its ordinary kind of semantic-value, though we do not and cannot know which value in particular it receives.[3]

Our aim in this paper is defend AR. In §1, we consider and respond to the most obvious objections to AR, clarifying the thesis in the process. The remaining two sections provide our positive case in favour of AR. In a nutshell, our positive argument is an inference to the best explanation: we argue that AR provides the best explanation of a range of phenomena. In §2

[1] We do not need to always explicitly use the qualification 'arbitrary'. Given the right context, we may simply use stipulations such as 'Let *n* be a natural number' or 'Let Pierre be a Frenchman'. We insert the explicit qualification simply to ensure one focuses on the relevant kinds of readings.

[2] 'Knowledge which' claims are notoriously context sensitive. When we say that we do not know which number *n* is, we mean that we cannot describe the number in some informative mode of presentation, such as '*n* is 343'. Of course we do know that *n* is *n*, that '*n*' refers to *n*, that '*n*' refers to whatever number it refers to, and so forth. This qualification should be kept in mind throughout the paper.

[3] Except, of course, in special situations where we know there is only one value that the expression could receive (e.g. in the case of 'Let *n* be an arbitrary even prime number').

we discuss the application of AR to instantial reasoning. We argue that AR can be used to develop a semantics which accounts for instantial reasoning, and moreover that this account is better than the prominent alternatives. In §3, we point to other applications that AR may have. The applications in this final section are more suggestive and require further development. Nevertheless, we think that they suffice to show that AR is a highly promising thesis, and that adopting it is likely to lay the path to a range of new solutions to some difficult philosophical puzzles.

§1 Arbitrary Reference: objections and responses

In this section we consider the most obvious objections to AR. If nothing else, the discussion should at least serve to clarify the thesis.

Objection 1: AR cannot be correct because it violates some necessary conditions on reference. For an agent to refer to an object she must stand in some special kind of relation to the object: she must (in some sense) be 'acquainted' with the object, or stand in some kind of casual relation to the object or … (substitute your favourite condition on reference here). But often when one makes a stipulation such as 'Let Pierre be an arbitrary Frenchman', one does not stand in the appropriate relation to any particular Frenchman, and so one cannot succeed in referring to any particular Frenchman.

Response 1: First, we are highly sceptical of any such necessary conditions on reference. For example, we maintain that the stipulation "Let 'Julius' refer to the actual tallest spy" is entirely appropriate and, assuming that there is a unique tallest spy, we can successfully use it to fix the reference of 'Julius', even if we do not stand in any significant casual relation to that spy (or any other significant relation of acquaintance).[4]

Second, even if there are such acquaintance-like conditions on reference, this would not really threaten AR. All we need to consider are examples in which one *does* stand in the relevant acquaintance-like relation to the object referred to. To take an extreme case, suppose that one maintains that in order to refer to a concrete object one needs to have seen that object. Now suppose that there are two bottles on the table in front of Alice and she can see each of them. Alice says to herself 'Let Jake be an arbitrary bottle on the table'. According to

[4] For a detailed defence of the claim that reference does not require such special necessary conditions see Hawthorne and Manley (MS).

AR, 'Jake' (arbitrarily) refers to one of the bottles on the table. But Objection 1 does not apply to this case, because even the (artificially stringent) condition on reference envisaged is satisfied.

The crucial point is this. It does not follow from our view that *any* stipulation of the form 'Let *a* be an arbitrary F' successfully results in fixing a referent for '*a*'. One could plausibly require, for example, that if there are no Fs the stipulation fails and it is perfectly compatible with AR that it also fails when other, more complex conditions, do not obtain. All that AR requires is that at least some such stipulations succeed.

Objection 2: If a stipulation such as 'Let Pierre be an arbitrary Frenchman' succeeds in fixing the reference of 'Pierre' to some particular Frenchman, then something must determine which Frenchman is referred to. But it is difficult to see what that would be. There doesn't seem to be anything about our behaviour, the sounds we utter, our brain states, or even our external environment that determines which Frenchman is referred to.

Response 2: We agree that none of the factors mentioned determine which Frenchman is referred to. In fact, we propose that *nothing* determines which Frenchman is referred to – nothing, that is, other than the semantic fact that we have referred to the particular Frenchman in question.[5] We simply deny that for it to be a fact that some particular Frenchman is being referred to, some other facts need to determine this fact.

The mere claim that some facts are not determined by other facts is not in itself surprising. If the domain of facts is well-founded, then some facts are not such that they are true in virtue of other facts. Perhaps the worry is, though, that it there is something especially troubling about the claim that the particular fact in question (namely, who 'Pierre' refers to) is not determined by other facts. We turn to this point in the next objection.

[5] A tricky question is whether our intentions are sufficient to determine which Frenchman is being referred to. We certainly don't think that the reference is determined by any informative intentions of the sort 'I intend 'Pierre' to refer to Jacques Chirac'. But it may be that one has an intention that 'Pierre' refers to Pierre or that 'Pierre' refer to an arbitrary Frenchman, and that these in turn are sufficient to determine the reference of 'Pierre'. But in so far as our intentions determine the reference in the latter way, there is still an important semantic fact (one concerning the word 'Pierre', or the phrase 'an arbitrary Frenchman', or analogous phrases in one's language of thought), a semantic fact that is crucial in determining the reference of 'Pierre' and is not itself determined by non-semantic facts.

Objection 3: In response to Objection 2, you claim that when one stipulates that Pierre is an arbitrary Frenchman, nothing determines which Frenchman 'Pierre' refers to. Let us concede the general point that some facts are not determined by other facts. But facts about reference are *semantic* facts, and it is standardly accepted that semantic facts are not primitive: rather they fully determined by use facts (broadly construed). In short, AR seems to contradict the platitude that semantic facts supervene on use facts.[6]

Response 3: We accept that AR conflicts with the commonly-held view that semantic facts supervene on use facts. Indeed, one way to bring out this conflict is to consider possible worlds, w_1 and w_2, identical in all non-semantic respects, where in both Jill makes the stipulation 'Let Pierre be an arbitrary Frenchman'. According to AR as we would like to think of it, it is possible that following Jill's stipulation, 'Pierre' refers to one Frenchman (say Jacques Chirac) in w_1 but to another (say Nicolas Sarkozy) in w_2. Thus the two worlds are identical in their use facts (because they are identical in every non-semantic respect) but they differ in their semantic facts (because 'Pierre' has a different referent in each of the two worlds). So if AR is correct then semantic facts do not supervene on use facts.

We also recognise that Objection 3 is probably the most compelling objection to AR, and the main reason why it may be seen as radical. Nevertheless, we insist that the view that semantic facts supervene on use facts is simply incorrect. We rely here on the detailed defence of this claim that is provided in Kearns & Magidor (*forthcoming*). In particular, we note that this defence in no place relies on the truth of AR, and thus is not circular.[7]

---

[6] We allow that the objector interprets 'use facts' in a broad enough way so as to include facts about one's environment or facts about which properties are most natural. It is clear, however, that the objector does not intend count as use facts such semantic or intentional facts as the fact that one uses 'Pierre' to refer to, e.g., Jacques Chirac. For a more detailed discussion on how the claim that semantic facts supervene on use facts ought to be interpreted see Kearns & Magidor (*forthcoming*).

[7] It is worth noting that one could also consider a different interpretation of AR, one according to which facts about arbitrary reference *do* supervene on use facts. One constraint on developing the theory in this manner is that to ensure that it will allow that not all stipulations of the form 'Let *n* be an arbitrary number' result in *n* referring to the same number. (Otherwise one will have problems in applying AR to the case of instantial reasoning, in particular to stipulations such as 'Let *n* be an arbitrary number and let *k* be an arbitrary number). Two additional challenges for this interpretation of the view is to provide an explanation, on the one hand for how reference supervenes on use facts, and on the other hand for why – despite the supervenience claim - we do not and cannot know the referent. One avenue to explore in this context is a brute supervenience view: one according to which the semantic facts supervene on use facts in an entirely unexplanatory manner and are hence unknowable (cf. Williamson (1994) and Cameron (2010)). But we find this position dialectically inferior to our own: if one is going to allow bruteness concerning the semantic realm into one's theory, why insist on the supervenience claim, rather than simply postulating brute contingent facts concerning reference, as our own theory does? After all, contingent brute facts seem less offensive than necessary ones. (See also Kearns &

Objection 4: You claim that following the above stipulation, 'Pierre' refers to a particular Frenchman. How then, do you explain the fact that we do not and cannot know which Frenchman 'Pierre' refers to?[8]

Response 4: In a nutshell, our response is that we are ignorant of this fact precisely because nothing (non-semantic) determines which Frenchman 'Pierre' refers to.

Compare this with another case, one of determination over time. Suppose you are going to flip a coin (and assume that coin flips are genuinely non-deterministic). You do not and cannot know which side the coin will land. Why? The obvious answer is that this is because nothing in the current state of the world determines which side the coin will land. Of course, one could insist on a further explanation, namely an explanation of why you do not and cannot simply have direct knowledge of the primitive fact that coin will land on heads/tails.[9] Yet most of us seem content to simply shrug our shoulders at this stage: for some reason we simply don't have knowledge of such primitive facts about the future.

Return to the case of arbitrary reference. You do not and cannot know what 'Pierre' refers to. Why? The obvious answer is that this is because nothing in the state of the world determines what 'Pierre' refers to. Of course, one could insist on a further explanation, namely an explanation of why you do not and cannot simply have direct knowledge of the primitive semantic fact that 'Pierre' refers to so and so. Yet it seems that if we were happy to shrug our shoulders in the coin flipping case, we should be equally happy to shrug our shoulders in the case of arbitrary reference.

One could of course point out to a disanalogy: in the case of the coin flip one could come to know the result of the flip *after* the coin-flip had happened, while in the case of arbitrary reference one cannot come to know which Frenchman 'Pierre' refers to, even after the stipulation was made. This is correct, but misses the point of the original analogy. After the

---

Magidor (*forthcoming*), §2.3 for a similar argument). At any rate, at least for the purposes of this paper, we have chosen to develop a version AR which denies supervenience.

[8] It is worth point out in this context that our claim that one cannot know which Frenchman 'Pierre' refers to is not intended to exclude scenarios where an omniscient being knows who Pierre refers to, or where one comes to know this fact by testimony of an omniscient being. When we say that we cannot know who 'Pierre' refers to, we mean that we cannot know this in any reasonably ordinary scenario, one that does not involve omniscient beings. See Williamson (1997), p. 926 for a similar qualification regarding his epistemic view of vagueness.

[9] We are assuming here that even in non-deterministic settings there are such facts about the future.

coin was flipped, there are facts about the world that determine that the coin *was* flipped in a certain way (e.g. the facts that the coin wasn't moved after the flip, and that it is now lying with its heads-side up). But even after the stipulation concerning Pierre was made, there are no facts (again – no facts *other* than the fact that 'Pierre' refers to whoever it refers to) that determine which Frenchman 'Pierre' refers to, so it is no surprise that we cannot and do not know this fact.

Objection 5: Even if the stipulation succeeds in fixing the reference of 'Pierre', it does not fix it to be some particular Frenchman. After all, the stipulation would have achieved exactly the same effect if instead of saying 'Let Pierre be an arbitrary Frenchman' we would have said 'Let Pierre be an arbitrary Frenchman - *no particular one*'.

Response 5: We agree that the two stipulations would result in the reference of 'Pierre' being fixed in exactly the same way, but we still maintain that it would have been fixed (even in the latter case) to a particular Frenchman. 'Pierre' refers to a particular Frenchman. The qualification 'no particular one' is a qualification of the manner of fixing the reference, rather than of the Frenchman referred to. That is to say, it points out that we are fixing the reference *arbitrarily* rather than *particularly*. This is not an uncommon use of 'particular' in English. Consider for example the announcer of the results of a competition who says 'In no particular order, the winners are A, B, and C!'.[10] In one sense, of course, the winners are given in a particular order: first A, then B, then C. But in another sense they are not: the order was chosen arbitrarily and not particularly. Similarly, we suggest, Pierre is a particular Frenchman in the former sense (there is some specific guy who is Pierre), but not in the latter sense (Pierre was picked arbitrarily and particularly).

Objection 6: Even if the stipulation succeeds in fixing the reference of 'Pierre', it does not fix it to an ordinary particular Frenchman, but rather to a special kind of object: an *arbitrary Frenchman*.[11]

Response 6: We are happy to concede that Pierre is an arbitrary Frenchman, but we deny that this means that Pierre is a special kind of object, one that is not an ordinary, flesh and blood,

[10] We note that precisely this locution is used in the announcements of the results in the British television show 'The X factor'.
[11] Cf. the discussion of Fine's view in §2.1.3 below.

particular Frenchman. On our view an arbitrary Frenchman is just a Frenchman that is referred to arbitrarily.

Compare this to the following cases. Suppose you make an intentional mistake. It is perfectly true that the mistake you made is intentional. But this does not mean that the mistake you made is a special kind of mistake: it is simply a mistake made intentionally. Similarly, suppose someone gives you a pen as a gift. It is perfectly true that the pen in question is a gift. But this does not mean that the pen is a special kind of object ('gift-pen' as opposed to an ordinary pen). It is a gift simply because the pen was given in a certain way. Thus we maintain that Pierre is both an arbitrary Frenchman and a perfectly ordinary particular Frenchman.

§2 Arbitrary reference and instantial reasoning

The most obvious application of AR is to the interpretation of instantial reasoning. Much ordinary reasoning uses universal generalisation and existential instantiation. We would ideally like an account of such reasoning that satisfies two related constraints, a descriptive and a normative one.[12] On the descriptive side, we would like the account to correctly describe what it is that we are doing when making such inferences. More explicitly: when we make these inferences we use bits of language, and we would like an account of what those bits of language mean. On the normative side, we would like the account to explain why the inferences we make using UG and EI are justified. The two constraints are obviously linked: in so far as we believe that inferences made through instantial reasoning are justified, then an account would only be descriptively adequate if it also satisfies the second, normative constraint. In the discussion below it will be especially important to keep the descriptive constraint in mind: it is insufficient to suggest an account of instantial reasoning that provides a justification for some form of reasoning we could have been making but which in fact we do not make. Accounts which render instantial reasoning sound but are otherwise descriptively incorrect, are thus deficient.

Having clarified what we expect from an account of instantial reasoning we will proceed to argue for two claims: that AR underlies a good account of instantial reasoning, and that this account fares better than the alternative accounts. If we are right, then by inference to the best explanation one has good reason to accept AR. It should be noted, though, that the issue of

[12] Cf. Fine (1985b), p. 127, who suggests similar constraints.

how to interpret instantial reasoning is a complex one which has received a fair amount of discussion in the literature.[13] We cannot hope in the scope of this paper to discuss every alternative account in detail, and we do not take the section below to decisively refute every alternative account. But the discussion should at least give a sense of how difficult it is to provide a satisfactory account of instantial reasoning without appealing to AR.

§2.1 Problems with alternative accounts

Accounts of instantial reasoning fall into three kinds. First, there are *instrumentalist accounts,* according to which instantial terms[14] and the lines in the reasoning that include them are meaningless, and serve merely as convenient devices for reaching meaningful conclusions from meaningful premises. Second, there are *quantificational accounts*, according to which instantial terms are variables which are implicitly bound by quantifiers (or are themselves disguised quantifiers). Third, there are *referential accounts*, according to which instantial terms are names which refer to objects of some kind (arbitrary objects on Fine's account, ordinary objects on our own).

We begin by presenting some of the main difficulties with the instrumentalist and quantificational accounts, drawing heavily on Fine's seminal discussion of the topic.[15] We then proceed to argue that Fine's own view also faces serious difficulties.

It will be helpful throughout the discussion to have a particular argument involving instantial reasoning in mind. One paradigmatic argument which employs both UG and EI is the following argument, from the premise that there is someone who loves everyone to the conclusion that everyone is such that someone loves them. We present the argument, annotating each line with the rule of inference that seems, at least on the face of it, to be applied in each case:[16]

Argument 1:

[13] Literature on the topic includes at least Fine (1983), Fine (1985a), Fine (1985b), King (1991), Mackie (1985), Martino (2001), Price (1962), Rescher (1958), Shapiro (2004), and Tennant (1983).

[14] An instantial term is term *a*, such that in an application of UG we infer $\forall x\varphi(x)$ from $\varphi(a)$, or such that in an application of EI we infer $\varphi(a)$ from $\exists x\varphi(x)$.

[15] Fine (1985b), especially ch. 12.

[16] The semi-formal English used in this argument is intended to help keep track of the relative scopes of the quantifiers. We assume that a full syntactic parsing of ordinary English will provide us with similar formal properties.

(1) There is someone *x* such that for every person *y*, *x* loves *y* [Premise]
(2) Let John be such a person
(3) For every person *y*, John loves *y* [Existential Instantiation on 1]
(4) Let Jane be an arbitrary person
(5) John loves Jane [Universal Instantiation on 3]
(6) There is some person *x* such that *x* loves Jane [Existential Generalisation on 5]
(7) But since Jane was an arbitrary person, for every person *y* there is some person *x* such that *x* loves *y* [Universal Generalisation on 6]

§2.1.1 The instrumentalist account

According to the instrumentalist account, instantial terms and thus the lines in which they occur are simply meaningless. The reason we nevertheless include such lines in our arguments is that they are instrumentally valuable: there is a purely syntactic or proof-theoretic theorem which ensures that if we manipulate these meaningless symbols in a specified set of ways then we will only infer true meaningful conclusions from true meaningful premises.

The problem with this account is that it seems descriptively inadequate. According to the account, lines 2-6 in the above argument involve nothing but a meaningless manipulation of symbols. Yet it certainly seems to us that claims such as ‘For every person *y*, John loves *y*’ in line 3 are not only meaningful, but (assuming we believe the premise) true. Moreover, according to the instrumentalist account, the only way to see that inferences involving instantial terms are valid is not by acknowledging that each step of the inference is individually acceptable, but rather by knowing some highly theoretical result in proof theory which shows that the relevant kinds of manipulations of symbols will ultimately yield correct results. But it at least seems that agents can perfectly well recognise the soundness of the above reasoning without knowing this general proof-theoretic result. The instrumentalist view thus seems a descriptively inadequate explanation of our practices.

§2.1.2 Quantificational accounts

According to quantificational accounts, we should think of instantial terms as variables, ones that are implicitly bound by quantifiers.[17] The question, however, is how exactly are the instantial terms bound?

At first, suppose that the proof as a whole occurs within the scope of some quantifiers: universal quantifiers binding instantial terms connected with applications of UG, and existential quantifiers binding instantial terms connected with applications of EI[18] (call this 'the wide scope theory'). While the wide scope theory might sound initially tempting, trying to reconstruct an actual argument following its recommendation immediately runs in difficulties. Consider argument 1 above. We could try to construe it as a long conjunction of the following sort (with 'j_1'and 'j_2' replacing 'John' and 'Jane' respectively):

There is a person j_1 such that for every person *y*, j_1 loves *y*, and for all persons j_2:
There is someone *x* such that for every person *y*, *x* loves *y*
& for every person *y*, j_1 loves *y*
& j_1 loves j_2
& …
& for every person *y*, there is some person *x* such that *x* loves *y*.

But this completely distorts the structure of the argument. First and most importantly, it seems that we are simply asserting each of the lines (including the conclusion!), rather than inferring it from previous lines. Second, on this construal some of the lines stand for open formulas, so are not individually truth-valued. Third, we may later extend our reasoning so as to include more lines which mention John and Jane. But then the new tokens of the instantial terms will fall outside of the scope of our quantifiers (or else we will have to assume that the new argument is not really an extension of the original one). And it is hard to see how different variants of the wide scope proposal will avoid similar problems.

A much more promising proposal is to construe each step of the argument as individually bound by quantifiers (call this the 'narrow scope theory'). This has the advantage of allowing

[17] King (1991) offers a quantificational account that treats the instantial terms themselves as implicit quantifiers, rather than as variables. This point will make no difference to our argument.
[18] Many formal systems are phrased so that the same instantial term cannot act in both capacities. It is worth noting, though, that as long as we are careful to phrase the rules correctly, there should not be a principled difficulty in using a term for both purposes: an instantial term *a* used in an inference from ∃xFx to F*a*, can be treated as an arbitrary F, and generalised over to show that all Fs have a certain property.

each line of the argument to be truth-valued, and the argument as a whole to have a standard structure: it consists of a series of truth-valued statements, each of which follows from the previous steps. Still, construing instantial reasoning along the lines suggested by the narrow scope theory is far from smooth either. Consider for example the following argument, which from the premise that everything is F and everything is G, concludes that everything is F and G:

Argument 2:

(1) For all *x*, *x* is F [premise]
(2) For all *x*, *x* is G [premise]
(3) Let *a* be an arbitrary object
(4) *a* is F [UI on 1]
(5) *a* is G [UI on 2]
(6) *a* is F and *a* is G [Conjunction Introduction on 4 and 5]
(7) For all *x*, *x* is F and *x* is G [UG on 6]

According to the narrow scope theory, lines 4, 5, and 6 say the same as lines 1, 2, and 7 respectively, so it is hard to see why we need these steps in the argument. And it is worth noting that one cannot reply to this worry by claiming that these lines serve the purpose of making explicit content that was previously expressed implicitly: after all, according to the narrow scope construal, lines 4 and 5 do precisely the opposite, namely restate content that was previously expressed explicitly in an implicit way. Another problem is that on the current construal step 6 is not actually obtained from steps 4 and 5 via conjunction introduction, as our provisional annotation suggests. Rather, it consists of a fairly complex inference from the claims that for all *a, a* is F and that for all *a, a* is G, to the claim that for all *a*, *a* is F and *a* is G. While this inference is sound, it is far from trivial – indeed it is precisely the soundness of this inference that Argument 2 as a whole was supposed to establish!

As Fine illustrates, things become even trickier when we consider certain applications of EI. Consider for example the following argument, from the premises that there are some French men and that everyone is tall, to the conclusion that there is a tall French man.

Argument 3:

(1) There are some French men [premise]

(2) Let Jack be one of them
(3) Jack is a French man [EI on 1]
(4) Everyone is tall [premise]
(5) Jack is tall [UI on 4]
(6) Jack is tall and Jack is a French man [Conjunction Introduction on 3 and 5]
(7) So there is a tall French man [EG on 6]

This seems like a perfectly good argument. But now consider how the narrow scope theory treats step 6. According to the theory, step 3 says (in a disguised way) that there is a French man, step 5 says (in a disguised way) that there is someone who is tall, and step 6 says (in a disguised way) that there is someone who is tall and a French man. This means that not only does step 6 not follow from 3 and 5 by Conjunction Introduction (as our annotation suggests), but in fact it does not follow from 3 and 5 at all: from the claim that there exists an F and there exists a G, it does not follow that there exists something that is both F and G. Thus according to the narrow scope theory, Argument 3 is simply illegitimate.[19]

Finally, the narrow scope proposal has a problem accounting for the role of suppositions in arguments. Take an argument that contains the stipulation ‘Let *n* be an arbitrary number’. Further down the argument we might have the following line: ‘Suppose that *n* is even’, and we may then go on to show that certain things follow from this supposition. But, at least on its most straightforward version, the narrow scope theory construes the supposition as the claim that *all* numbers are even, which is clearly not what we intended.

Having objected to the narrow scope theory in general, it is worth saying a few words on what is by far the most sophisticated version of the theory - the account of instantial reasoning presented in King (1991). One key advantage of King’s account is that rather than suggesting the narrow scope theory in general terms, it offers a systematic way of determining (at least with respect to one particular formal system), precisely how each line in a proof should be interpreted. Another admirable feature of the account is that the interpretations are constructed in a careful way which avoids some of the technical worries mentioned above. Thus for example, argument 3 above is interpreted so that line 6 does

[19] That is, unless, the argument is construed so that step 6 follows directly from steps 1 and 4, and step 5 is taken to be completely redundant.

follow from line 3 and 5.[20] And suppositions are treated so that each line within the scope of the supposition is in effect interpreted as a conditional, which has the supposition as its antecedent.

Nevertheless, King's system does not manage to avoid the main problem facing the narrow scope theory: namely that it misconstrues the structure of arguments by instantial reasoning. Argument 2, for example, is construed in King's system just as we suggested above, where line 6 follows from lines 4 and 5 by a much more complex form of reasoning than Conjunction Introduction (indeed a form of reasoning the validity of which the whole argument was meant to establish).[21] Similarly, while the treatment of suppositions in King's system is technically adequate, it does not seem to correctly represent the structure of arguments involving suppositions. As Fine notes: "In making the supposition φ I am not asserting the trivial conditional $\varphi\to\varphi$ and in making an inference (say $\varphi\vee\psi$) from a supposition (say φ) I am not inferring one conditional ($\varphi\to(\varphi\vee\psi)$) from another ($\varphi\to\varphi$)."[22]

Given King's admirable work of construing a precise system which predicts the specific list of quantifiers and their relative scopes in each line of the proof, we can now see an additional problem that the narrow scope theory faces: the theory is forced to make some highly arbitrary choices. Consider Argument 1. Roughly speaking, on King's system the statement 'John loves Jane' in line 5 is interpreted with an existential quantifier within the scope of a universal quantifier ('There is a person *x* such that for every person *y*, *x* loves *y*'), and line 6 ('There is some person *x* such that *x* loves Jane') is interpreted with a universal quantifier within the scope of an existential one ('For every person *y* there is some person *x* such that *x* loves *y*').[23] This means that the crucial quantifier switch happens between lines 5 and 6 in the

---

[20] Very roughly this is achieved by interpreting line 5 to say that some French man is tall, rather than merely that someone is tall. But this is actually a gross oversimplification. What King's system in fact predicts is the following interpretations for line 3, 5, 6 (with 'F' standing for French man and 'T' for tall):
(3) $\exists y((\exists x\mathrm{F}x\to\mathrm{F}y)\wedge\mathrm{F}y)\wedge\forall y((\exists x\mathrm{F}x\to\mathrm{F}y)\to\mathrm{F}y)$
(5) $\exists y((\exists x\mathrm{F}x\to\mathrm{F}y)\wedge\mathrm{T}y)\wedge\forall y((\exists x\mathrm{F}x\to\mathrm{F}y)\to\mathrm{T}y)$
(6) $\exists y((\exists x\mathrm{F}x\to\mathrm{F}y)\wedge\mathrm{F}y\wedge\mathrm{T}y)\wedge\forall y((\exists x\mathrm{F}x\to\mathrm{F}y)\to\mathrm{F}y\wedge\mathrm{T}y)$
The complexity of these interpretations should already give us serious cause for concern.

[21] Things get even worse with other so called applications of Conjunction Introduction – as we see in the argument from 3 and 5 to 6 in Argument 3 (see footnote 20).

[22] Fine (1985b), p. 134.

[23] We say 'roughly' because again the system is more complex than that: line 5 will actually be interpreted as saying that $\exists a((\exists x\forall y\mathrm{L}xy\to\forall y\mathrm{L}ay)\wedge\forall b\mathrm{L}ab)\wedge\forall a((\exists x\forall y\mathrm{L}xy\to\forall y\mathrm{L}ay)\to\forall b\mathrm{L}ab))$, and line 6 as saying that $\exists a((\exists x\forall y\mathrm{L}xy\to\forall y\mathrm{L}ay)\wedge\forall b\exists x\mathrm{L}xb)\wedge\forall a((\exists x\forall y\mathrm{L}xy\to\forall y\mathrm{L}ay)\to\forall b\exists x\mathrm{L}xb))$. But what is crucial to the argument here are the claims that $\exists a\forall b\mathrm{L}ab$ and $\forall b\exists x\mathrm{L}xb$ which we get from the first conjuncts of the interpretations of

argument. But now suppose we move the stipulation that Jane be an arbitrary person (line 4), above line 2 (i.e. before applying EI to the premise). Now the argument will be construed on King's system so that line 5 ('John loves Jane') has a different interpretation, one where the existential quantifier is in the scope of a universal quantifier ('For every person y, there is some person x, such that x loves y'), which means that the crucial quantifier switch already happened by the time we reach line 5. The upshot is that on King's account, it makes a crucial difference both to the interpretation of each line and to the structure of the argument more generally whether the stipulation 'Let Jane be an arbitrary person' is made before or after the application of EI. But intuitively, there is no such difference between the two versions of the argument.

By reflecting on King's construal of Arguments 1 through 3 we observe four problematic features of his account, which are characteristic of narrow scope accounts more generally. First, the account construes certain lines in the proof as mere repetitions of the premises or conclusion, sometimes rendering implicit, material that was already presented explicitly. Second, the account construes certain lines in the proof that seem to have a very simple structure (e.g. 'John loves Jane') as making very complex claims ('$\exists a((\exists x\forall y \mathrm{L}xy \rightarrow \forall y \mathrm{L}ay) \wedge \forall b \mathrm{L}ab) \wedge \forall a((\exists x\forall y \mathrm{L}xy \rightarrow \forall y \mathrm{L}ay) \rightarrow \forall b \mathrm{L}ab)$').[24] Third, the account construes what seem to be very simple steps in reasoning, as actually employing quite complex forms of reasoning (e.g. the move from a '$\exists\forall$' scope structure to a '$\forall\exists$' scope structure) – complex forms of reasoning the validity of which is often supposed to be established by the argument as a whole, rendering these steps in the argument question begging. Fourth, the account entails that seemingly arbitrary and unimportant choices in the placing of stipulations in the argument make a substantial difference to how the argument is interpreted. We conclude that quantificational accounts do not provide a satisfactory account of instantial reasoning.

§2.1.3 Fine's referential account

Instantial terms exhibit, at least on the face of it, the syntactic behaviour of proper names. This fact, coupled with the difficulties facing the quantificational accounts, provides a strong

lines 5 and 6 respectively. We thus focus on these claims, and apply similar simplifications in the discussion below.

[24] We are fully aware that King is in no way committed to saying that the original claim has the same syntactic structure as that of the complex interpretation. But even so, the semantic complexity of the interpretations is in itself worrisome.

reason to think that instantial terms are proper names, and serve to refer to objects. The most prominent view of instantial reasoning which takes instantial terms to be referential is Kit Fine's view.[25] We cannot hope in this space to do full justice to Fine's detailed and intricate discussion, but we would like to point out several aspects of Fine's view which are considerably problematic.

According to Fine, instantial terms refer to special kinds of objects - 'arbitrary objects' – which are distinct from any particular object. Thus for example when we stipulate that *n* be an arbitrary number, we fix the reference of '*n*' to an arbitrary number, which is an entity distinct from any of the familiar particular numbers (2, 5467, and so forth). The rough picture is that each arbitrary object has a 'value-range' associated with it which determines which properties it has: an arbitrary object will have all and only the properties that are shared by all the (particular) objects in its value-range.

As Fine recognises, however, this rough picture requires substantial refinement. To begin with, we cannot accept without qualification that an arbitrary object has all and only the properties shared by all the objects in its value-range: the arbitrary number is an arbitrary object, even though not all of the objects in its value-range are arbitrary objects (in fact none of them are). Similarly, the arbitrary person is an abstract object, even though not all of the objects in its value-range are abstract objects (in fact none of them are). In order to address this problem, Fine distinguishes between what he calls 'generic' and 'classical' conditions. Generic conditions are properties such as *being a number* or *being even*, ones that hold of an arbitrary object if and only if they hold of every object in its value-range. Classical conditions, on the other hand, are properties such as *being arbitrary* – ones which do not satisfy this principle.

The distinction between generic and classical conditions looks suspiciously like an ad-hoc fix. But even if we grant Fine this distinction, there are further problems. Consider the following question: how many natural numbers are there between one and ten? Intuitively, the answer is 'ten'. But it seems that Fine's theory predicts otherwise: according to Fine, in addition to all the particular numbers between one and ten, there is also the arbitrary number

[25] See Fine (1983), Fine (1985a), and Fine (1985b)

between one and ten (call it 'Arb'). So there must be at least eleven numbers between one and ten!

Does the distinction between classical and generic conditions help with this problem? It is not clear that it does. After all, *being a number between one and ten* seems like a paradigmatic example of a generic condition; if so, then since each of the objects in the value-range of Arb satisfies this condition, Arb satisfies it as well – so Arb is a number between one and ten.

A further attempt to respond might appeal to Fine's suggestion that "there may be cases in which a predicate is ambiguous as between a generic and classical reading. The predicate 'is a number' is a good example. On a generic reading, it is inclusive of all arbitrary numbers; on a classical reading it is exclusive of them".[26] The proposal is that on a classical reading of 'number between one and ten', Arb is not a number between one and ten, while on a generic reading it is. So the statement 'There are exactly ten numbers between one and ten' is true on one reading and false on the other. We find this response unsatisfactory. For a start, it seems that the sentence 'There are exactly ten numbers between one and ten' has *only* a true reading, and no false reading. Moreover, we see no reason to believe that 'number between one and ten' is ambiguous in the manner Fine suggests.[27] Consider for example the following speech: 'There are ten numbers between one and ten. Let Arb be (an arbitrary) one of them'. This seems like a perfectly standard way of introducing an instantial term, but Fine's ambiguity proposal predicts otherwise. Since the term 'them' is anaphoric upon the term 'numbers between one and ten' in the preceding sentence, both phrases must receive the same interpretation. So either 'number between one and ten' is interpreted classically, in which case Arb cannot be an arbitrary number which has this property, or 'number between one and ten' is interpreted generically, in which case the initial claim that there are ten numbers between one and ten is false.

Next we turn to a second refinement which Fine's theory requires, one which concerns the relationship between distinct instantial terms appearing in the same argument. Consider the stipulation 'Let m be an arbitrary real number and let k be an arbitrary real number greater than m'. According to Fine's theory, m and k will both be arbitrary real numbers, and will

[26] Fine (1985b), p. 14.
[27] Note also that since the counting problem seems to generalize, Fine would need to argue that pretty much every predicate is ambiguous in this manner.

have as their value-range the full set of real numbers. But this cannot be the whole story: we want somehow to capture the idea that *k* and *m* are related, and in particular that *k* must be greater than *m*. To that end, Fine proposes two additional pieces of machinery. The first is the idea that a collection of arbitrary objects may have 'joint' rather than 'individual' value-ranges. So, for example, instead of saying that *m* and *k* each have the full range of real numbers as their value-range, we can say that *m*, *k* jointly have as their value-ranges pairs of real numbers (pairs where the second member is larger than the first member). The second piece of machinery involves the idea that arbitrary objects come in two kinds: dependent and independent. While *m* is an independent arbitrary number, *k* is a dependent arbitrary number: in a sense "its value" depends on "the value" of *m*.

Other than adding a further layer of complication to the theory, this complex hierarchy of arbitrary objects runs into trouble when we consider the relation of identity between them. Consider the stipulation 'Let *m* be an arbitrary number, and let *k* be an arbitrary number'. One natural suggestion is to assume that *m* and *k* are both independent arbitrary objects, which have all numbers as their value-range. The problem is that Fine insists that there is only one independent arbitrary object associated with each value-range.[28] But this in turn means that we would be able to infer that $m=k$, which we clearly ought not to.

One could try responding to this problem by appealing to the ambiguity between classical and generic readings, claiming that '=' is ambiguous in this manner. But this will not do: first, it is not clear how we get a 'generic' reading of '=' here: after all, as applied to *independent* arbitrary object the generic reading is supposed to be interpreted by considering the range of values for each of the objects independently. Second, in so far as we can make sense of the generic reading for '=', it will be one according to which it is not true that $m=k$ (because it is not the case that all potential values for *m* and *k* will be identical to each other). But just as we ought not to infer from the above stipulation that $m=k$, equally we ought not to infer that $m\neq k$! Third, as Fine defined the 'classical' reading of a predicate, it is one where as applied to an arbitrary object it would be automatically false. So it turns out that *both* the classical reading and the generic one should yield the result that $m\neq k$, and thus we cannot appeal to the alleged ambiguity of '=' to justify our indecision as to whether $m=k$ or $m\neq k$. Finally, we note that the ambiguity theory is particularly unappealing in the case of identity: perhaps more

[28] See Fine (1985b), p. 18. Fine later relaxes this condition (ibid. p. 35), but claims he does this only because "for certain technical purposes, a smoother theory is obtained".

than any other predicate, we have a clear sense that identity has a complete and determinate definition, one that already applies in the standard way to any kind of object (even to arbitrary ones, if there are such things).

A different response to the identity problem would be to suggest that there can be more than one independent arbitrary object with the same value-range. So perhaps following the above stipulation, it is possible that '*m*' and '*k*' refer to the same arbitrary number but also possible that they refer to two qualitatively identical but numerically distinct arbitrary numbers, thus accounting for our indecision. The problem is that if this proposal is adopted, one would need to explain how it is that the reference of '*m'* is fixed to one independent arbitrary number rather than another. And given that the different independent arbitrary numbers are qualitatively identical, it is hard to see how could give such story, without already appealing to something like AR.[29] But as we will go on to show, if one already accepts AR, there is a much simpler theory of instantial reasoning, and thus there is no need for Fine's theory of arbitrary objects in the first place.

A final solution to the identity problem is to suggest that both *m* and *k* are *dependent* rather than *independent* arbitrary objects: perhaps they both depend upon the 'arbitrary pair of numbers', *p*. The value range of *p* is the set of pairs of numbers; *m* depends upon *p* in this manner: the value of *m* must be the same as the first element of the value of *p*; *k* depends upon *p* in this different manner: the value of *k* must be the same as the second element of the value of *p*.[30] But this solution will not do either. For a start the solution requires that the referent of '*m*' is fixed to one object (the independent arbitrary number) if the stipulation 'Let *m* be an arbitrary number' appears on its own in an argument, but fixed to an entirely different object (one that is dependent upon the arbitrary pair of numbers in the manner described above) if the stipulation 'Let *n* be an arbitrary number' appears later in the argument. But this seems entirely wrong: the two stipulations are completely unrelated, and it is odd that we cannot fix the reference of '*m*' until we see what other instantial terms appear further down the argument. As Fine himself notes (in another context): "It is a natural requirement on a derivation containing A-names that we know what those names denote as soon as they are introduced; their interpretation should not depend upon what subsequently

[29] Cf. §3.2 below.
[30] This seems to be the solution favoured by Fine (see Fine (1985b), p. 19).

happens in the derivation".[31] An even graver problem with the current proposed solution is that it too would lead to the undesirable conclusion that $m \neq k$: two dependent arbitrary objects are identical, according to the theory, if and only if they depend on the same arbitrary objects in exactly the same manner. But m depends on the arbitrary ordered pair in a different manner than k does, so we should be able to infer that the two are not identical, i.e. that $m \neq k$.

Leaving the identity problem, we go on to note a third layer of complication for Fine's theory. This has to do with the observation that Fine's theory seems to face the same objection concerning the role of suppositions that the narrow-scope quantificational view faced. Consider a proof containing the stipulation 'Let n be an arbitrary number', where later we add a supposition such as 'Suppose n is even'. Since on Fine's view an arbitrary number n is even if and only if every number is even, this suggests that the content of the supposition is that every number is even, which is clearly not what we want to suppose. Fine's response to this problem involves introducing the notion of a 'vacancy value':[32] Arbitrary objects are split into 'vacant' and 'occupied' objects, where an 'occupied' object is one that is treated (in some sense) *as if* it has been assigned one particular value, while a 'vacant' arbitrary object is treated in full generality. The logic is then amended so that vacancy values are marked, and inferences involving suppositions as above can only be made where the supposition is interpreted as involving an occupied arbitrary object. We will not discuss the details of this proposal, but we merely note that it adds yet another complication to Fine's theory, one that involves a technical apparatus which is hard to motivate on independent grounds, or to interpret in any natural way.

Our final major worry concerning Fine's theory is that it involves a rejection of classical logic. One way to see this is to note that according to Fine, "the basic principle is that a sentence concerning A-objects is true (false) just in case it is true (false) for all of their values".[33] This in turn entails that it is neither true nor false that the arbitrary number is even, and the principle of bivalence must be rejected. A related issue concerns the semantics of the connectives: it is true that the arbitrary number is either odd or even, even though it is not true that it is odd and it is not true that it is even. Fine is happy to endorse the rejection of classical logic (he notes that the phenomenon of vagueness should anyhow motivate us

[31] Fine (1985b), p. 101.
[32] Fine (1985b), p. 75-80.
[33] Fine (1985b), p. 41.

towards a similar rejection).[34] But is far from clear that classical logic should be rejected due to other phenomena[35], and at any rate, it certainly seems doubtful that the problem of accounting for instantial reasoning provides a convincing enough case for such a rejection. The rejection of classical logic seems like a high price for Fine's theory to pay.

We would have had no objection to postulating the existence of arbitrary objects, if this postulation was necessary (or even sufficient) for a good account of instantial reasoning. But as we have seen, it is far from easy to provide an account of instantial reasoning by appealing to arbitrary objects. Fine's theory is both highly complicated (requiring the distinctions between classical and generic conditions, dependent and independent arbitrary objects, vacant and occupied objects), and even given these distinctions the theory faces serious objections.

§2.2 The AR account of instantial reasoning

Having shown some of the difficulties with the prominent alternative accounts, we next turn to show that we can give a good account of instantial reasoning by appealing to AR.[36]

Consider Argument 1 again. According the account we propose, both instantial terms ('John' and 'Jane'), function as proper names that refer to particular persons. Following the stipulation in line 2, 'John' (arbitrarily) refers to a person who loves everyone. Following the stipulation in line 4, 'Jane' (arbitrarily) refers to a person.

At a first pass it may seem that this account cannot explain why Universal Generalization is a valid rule of inference. After all, if Jane is some arbitrarily chosen particular person, then Jane will have many properties not shared by all people. Suppose for example that Jane, our arbitrary person, happens to be French. It does not follow that all persons are French.

---

[34] Fine (1985b), p.11.

[35] As is well familiar there are various solutions to the problem of vagueness that do not involve such a rejection (see e.g. Williamson (1994) or Fara (2000)) and we note in particular that AR may well provide such a theory (see §3.3 below).

[36] Our account is in some ways very close to that involved in systems of Hilbert's Epsilon Calculus though as far as we know Hilbert was not particularly concerned with the metaphysical underpinnings of the epsilon operator. The idea that instantial terms associated with EI refer to particular objects appears in Mackie (1958), pp. 30-31, though it is not clear that Mackie is sensitive the problem of having more than one object potentially satisfying the existential quantifier, and to the metaphysical implications of such cases. Finally, Fine (1985b), pp. 136-138 discusses (and criticises) the view that instantial terms refer to ordinary particular objects, though the account he has in mind seems to be one where the agent is aware of which particular object is being referred to, rather than an account based on AR.

In responding to this worry, it is helpful to draw upon Prawitz's distinction between 'proper' and 'improper' inferences, and Fine's illuminating discussion of it.[37] Proper inferences are ones that are valid in the straightforward manner: if the premises are true then the conclusion must also be true. An improper inference is one that is valid in a slightly more roundabout manner. Two paradigmatic cases of rules of inference that are improper are the rule of Conditional Introduction in systems of natural deduction, and the rule of Necessitation in normal modal logics. In the case of Conditional Introduction, one supposes φ, proves ψ, and then infers $\varphi\rightarrow\psi$ (while discharging the supposition). It is easy to see why this is a semantically valid rule of inference, but this is not quite a straightforward case of inferring $\varphi\rightarrow\psi$ from the premises. Rather, we seem to infer $\varphi\rightarrow\psi$ from the fact that we have shown that ψ follows from φ. Similarly, the rule of Necessitation is one that allows us to infer $\Box\varphi$, given that we have proved φ. Again, this rule has a good semantic justification, but this justification is not quite that the conclusion follows from contents of the previous line in the proof: after all, it is not in general true that if φ is true then $\Box\varphi$ must also be true. The idea is that given that φ was *proved*, then φ must be a logical truth, and therefore a necessary truth as well. As Fine suggests, one helpful way to think of improper rules of inference is to assume that every line in an argument or proof serves two roles simultaneously: "On the one hand, the statement made, call it φ, will have been *inferred* from the suppositions used, call them Δ. On the other hand, it will have been *demonstrated* that φ can be inferred from the supposition Δ. Although the two goals are distinct, it should be noted that in achieving one I have, in effect, achieved the other. If I succeed in inferring φ from Δ, then that very inference can serve as a basis for demonstrating that φ is inferable from Δ".[38] Improper rules of inference rely on this second role: if φ was inferred from ψ, then we have demonstrated that φ is inferable from ψ, and hence that $\varphi\rightarrow\psi$ is true. And if φ was inferred from no premises, then we have thereby demonstrated that φ is inferable from no premises, so φ is a logical truth, and thus $\Box\varphi$ is true.

It should now be clear what we would like to say about Universal Generalization: UG is an improper rule of inference. Let Jane be an arbitrary person. It is not the case that if Jane has some property then it must be that all persons have this property. However, what is true is that if we can *demonstrate* that Jane has a certain property, then all persons have this

[37] Prawitz (1965) and Fine (1985b), pp. 69-74.
[38] Fine (1985b), p. 71.

property.[39] The reason for this is as follows: according to AR, we cannot know which person 'Jane' refers to. All we know is that Jane is some person or other. Thus in demonstrating that Jane has a certain property, the only properties of Jane that we can appeal to are ones that Jane shares with all other persons.

It is worth noting three advantages of this account of UG. First, by treating UG as an improper rule of inference, we have a simple explanation of why UG cannot be applied unrestrictedly in the scope of suppositions (an issue that, as we have seen, raises a problem for alternative accounts). Recall our proof which started by letting *n* be an arbitrary number, and proceeded to suppose that that *n* is even. We cannot apply UG here and infer that all numbers are even. But if UG were a proper rule of inference there would be no reason why it cannot be applied in this case (from the supposition that *n* is even we can, for example, infer that either *n* is even or *q*, or that *not not n* is even). But improper rules of inference should always be applied with care within the scope of suppositions. Consider again the rule of Necessitation. If we suppose that p, we can infer (under the supposition) that p. But we should not apply the rule of Necessitation to infer (under the supposition) that $\Box$p. (Otherwise we would have been able to discharge the supposition and infer $p \rightarrow \Box p$). The same is true for UG: if the claim that *n* is even is only inferred via a supposition (that *n* is even), then we haven't properly demonstrated that *n* is even, and hence should not infer that all numbers are even. In discussing the role of suppositions in proofs, Fine himself acknowledges that "in supposing that an arbitrary number *n* is even, I do not seem intuitively, to be restricting its values to those numbers that are even. Rather I seem to be following through the fate of particular arbitrary number, even one that 'might' be odd, and supposing that it is 'even'"[40]. Interestingly, this is precisely how our own account interprets suppositions involving instantial terms.

A second advantage of our account of UG is that it can explain one seemingly peculiar aspect of standard uses of UG in informal reasoning. Consider line 7 in argument 1. In the previous line, we have shown that there is some person *x* such that *x* loves Jane. Yet we seem reluctant to straightforwardly assert the universal conclusion without reminding the reader that 'Jane was an arbitrary person'. In fact, our use of 'Since…was arbitrary' in this context seems very

---

[39] One class of exceptions is that we can, in a sense, demonstrate that Jane has been named, referred to, referred to by us, and so forth: all properties that Jane may not share with all persons. But this does not seem to us a grave problem: there is indeed no temptation to apply UG to these special cases.
[40] Fine (1985b), p. 75.

typical of applications of UG.[41] Our suggestion is that this peculiarity can be explained precisely because UG is an improper rule of inference. The move from the claim that Jane has a certain property to the claim that everyone has this property is not in general valid, and hence we are reluctant to make it. The reader has to be reminded that this is a special case where this seemingly invalid inference is warranted.

Here is a third advantage of our account of UG. It seems that reasoning according to UG is highly reminiscent of another form of reasoning which we will call 'without loss of generality' arguments. Suppose one wants to prove a certain theorem in geometry. One will often proceed by drawing a diagram of a *particular triangle* ABC, demonstrating the theorem (without loss of generality) with respect to ABC, and then inferring that since the proof did not appeal to any features of ABC that it does not share with other triangles, then the theorem is true of all triangles.

While this kind of argument has some differences from cases where we apply AR (in the former case but not in the latter, we have a particular triangle in mind), it nevertheless seems that 'without loss of generality' arguments have a very similar structure to proofs by UG. And indeed, this is exactly what our account predicts: according to our account, an argument by UG with respect to an arbitrary triangle ABC, involve reasoning pertaining to particular triangle, except that this time we do not know *which* triangle is involved in the proof. And just as in the case of without loss of generality arguments, the conclusion is justified because we have not appealed to any property of ABC that it does not share with other triangles. But without loss of generality proofs are notoriously prone to error: when one is working with a particular object, it is often hard to ensure that one ignores the properties it does not share with other members of its kind. The advantage of referring to an object arbitrarily is that our ignorance ensures that we will not make this type of mistake.

It is natural to think of ignorance as a kind of epistemic deficiency. But reflection on the case of UG reveals that sometimes ignorance can be of epistemic advantage. The fact that when we let be n be an arbitrary number we do not know which number n is, helps us ignore those properties of n which in this context could be distracting. There are other cases where ignorance provides us with a similar kind of epistemic advantage. Suppose for example that

[41] Fine also observes this peculiarity and thinks it needs to be explained, but offers a different explanation than ours – one involving his distinction between occupied and vacant objects. (see Fine (1985b), p. 80).

you are serving on a hiring committee, and want to ensure that you assess the candidates only on the basis of their academic merit, rather than on the basis of their race or gender. One way to try and achieve this goal is to make a conscious decision to try and focus only on the applicants' qualification and ignore available information about their race and gender. However, as psychological research shows, even with good intentions this aim is hard to achieve: subconsciously, knowing certain facts about an applicant's background is likely to influence one's judgment. A much more reliable way to achieve your aim is to ensure that the application files you consider simply do not contain any information about the applicants' race or gender. If you are completely ignorant of these facts, then you can be absolutely sure that you are ignoring them in your decision making process. It is often acknowledged that there are sometimes practical advantages in ignorance (e.g. ignorance of certain facts might make you happier), but ignorance can often have epistemic advantages as well.

Having considered UG, we turn to the case of Existential Instantiation. Here our account is even more straightforward. Suppose that there exists an F. Let '*a*' refer to one of the Fs. Which one? Well, if there is more than one, then let '*a*' arbitrarily refer to one of them. Since '*a*' refers to an F, we can validly infer that F*a*. Again, this account has several advantages in predicting our reasoning practices. For one thing, it seems to us that in typical applications of EI, reasoners often seamlessly think as if they are referring to some particular F. (Consider for example a typical application of the Axiom of Choice. The axiom merely established that there exists a choice function. But applications of the axiom typically proceed by letting *f* be one such function, and going on to talk as if one can make use of this function and apply it in building various constructions). For another thing, consider the applications of EI where it is clear that there can be *at most one* F. Suppose for example that we prove that a certain set of real numbers S has a (strictly) minimal member. Now let *k* be the minimal number in S. It seems almost obvious that '*k*' here refers to a particular object (namely the minimal member of S). But the alternative accounts of instantial reasoning face an unhappy dilemma: they must either deny '*k*' refers to a particular number, or they must claim that there are two fundamentally different EI-rules: one that applies in cases where the predicate is satisfied by at most one object (or when it is known to be satisfied by at most one object?), and one for all other cases. Our own account of EI, on the other hand, avoids this dilemma: we argue that the two cases of applying EI are on a par, since both involve reference to a particular object.

There is, however, one difficulty with our account of EI: how should EI be interpreted in cases where the existential premise is false? This problem is especially acute in the case of proofs by contradiction, where we may well suppose an existential premise which we believe from the outset to be false. A promising technical solution to this problem is to assume that if the existential premise is $\exists x Fx$, then EI allows us to infer that Fa, where '*a*' denotes an arbitrary F if there is one, and an arbitrary object which is a non-F, if there isn't one. (This will ensure, for example, that if $\neg Fa$, then $\neg\exists x Fx$, which seems desirable). But whether or not this technical solution is ultimately accepted, we note that this problem is not really peculiar to the case of AR. Consider again the applications of EI which involve conditions which can be applied to at most one object. For example, a proof by contradiction might start with the assumption that there is a greatest prime number, and proceed to say 'Let *n* be the greatest prime'. But if there were a greatest prime, it would have been unique, so one would not need to appeal to AR in order to refer to it. We thus suggest that the problem of applying EI when the existential premise is false is part of the more general problem of empty names, which would need to be addressed independently. And we see no reason to think that whatever general solution is provided to that problem would not be applicable specifically to cases where AR is involved.

We conclude that AR underlies a simple and compelling account of instantial reasoning. Given the difficulties facing the alternative accounts, this provides a strong reason to prefer our account, and thus AR more generally.

§3 Further applications of Arbitrary Reference

In the previous section we argued that AR underlies the best account of instantial reasoning. In this section, we briefly note three further fundamental problems that AR might help to (though working the details of these solutions will require further development).

§3.1 Benacerraf's problem

A famous problem raised by Paul Benacerraf is the following.[42] It is widely accepted that a set theoretic reduction for the notion of an ordered pair is adequate if and only if it satisfies the constraint that $\langle a,b\rangle=\langle c,d\rangle$ if and only if a=c and b=d. The problem is that there is more than one reduction which satisfies this constraint. For example, one could propose that the

[42] Benacerraf (1965). (Benacerraf uses a different example then the one we present above, but the problem is basically the same).

ordered pair <a,b> is identical to the set {a, {b}}; or one could propose that it is identical to the set {{a,b}, a}. Each reduction seems individually perfectly good. But are they correct? They cannot both be correct, because that would entail that {a, {b}}=<a,b>={{a,b}, a}. One could say that neither is correct, but that too seems odd. Consider for example the first reduction. It seems perfectly adequate on its own. How could it be that the existence of yet another good reduction diminishes the adequacy of the original? Finally, Benacerraf argues, it is also seems difficult to claim that one, but not the other reduction, is correct: after all they seem equally adequate, so it would completely arbitrary to prefer one over the other.

Our suggestion is that AR provides a new way to respond to Benacerraf's problem: of all the possible functions that exhibit the adequacy constraint on ordered pairs, the function that the term 'ordered pair' picks out is fixed arbitrarily. (Thus we could introduce the notion of an order pair by a stipulation like: 'Let 'ordered pair' denote an arbitrary adequate function'). This entails that there is a determinate fact of the matter as to whether or not <a,b>={a,{b}}, but we do not and cannot know this fact. According to this proposal, Benacerraf is correct in thinking that it is arbitrary which adequate reduction order pairs receive, but is incorrect in thinking that this prevents one such reduction from being the correct one.

This proposal has the advantage of capturing the structuralist intuition that the only feature relevant to being an ordered pair is the adequacy property. This is precisely why we use only those structural features in order to fix the reference of the term arbitrarily, and although ordered pairs in fact have some properties other than their structural ones, we do not and cannot know what these properties are, and hence they are irrelevant for our purposes. On the other hand, our proposal manages to avoid some of difficulties problems that traditional structuralism faces, namely questions concerning the nature of structures.[43]

§3.2 Referring to indiscernibles

Suppose we are presented with two objects, which are (qualitatively) identical in all non-haecceitistic respects. Can we refer to these objects? Consider a world containing two

[43] One interesting question about our account is whether it entails that we all refer to the same object by '<a,b>'. It is fairly easy to work out our account so that at least different speakers of the same community use the term with the same reference: after all, it may be that the reference is fixed arbitrarily once and for all for the community, and then used by everyone parasitically on the initial reference fixing, as is standard in the framework of the casual theory of reference. It may be harder to show that a different community with similar 'order pair' practices refers by '<a,b>' to exactly the same object as we do, but we do not see this as terribly worrying.

qualitatively identical (co-located) spheres.[44] Call one of them 'A' and the other 'B'. At least on the face of it, this stipulation names one of the spheres 'A'. But which one? No description we can give will tell the two apart, so it is hard to see how we can successfully fix the reference of 'A'.

Given AR we can easily explain how we successfully fix the reference of 'A': 'A' arbitrary refers to one of the spheres. One might worry, though, that this is not a particularly impressive achievement on the part of AR. Why not simply accept that 'A' does not successfully refer in this case? Other examples of reference to indiscernibles reveal that the 'unsuccessful reference' hypothesis is not very appealing. There are two numbers x such that $x^2 + 1 = 0$ - we call them 'i' and '$-i$'. But how did we succeed in naming one of them 'i'? There are no non-haecceitistic properties that one number has but the other lacks, so we could not have used any such properties to fix the reference of 'i'. And haecceitistic properties are of no use either, because even if there is a haecceitistic property that i has but $–i$ lacks (e.g. being identical to i) we could not have referred to this property until after we had already succeeded in referred to i, so the property could have been no help to us in fixing the reference of 'i'. It is not especially tempting in this case to claim that we failed to fix the referent of 'i': we seem to make many true mathematical statements containing the name.[45] AR provides a simple solution to the problem: 'i' had its reference fixed arbitrarily to one of these two numbers.

§3.3 Vagueness

A third promising application of AR is in providing a new account of vagueness. As Williamson has persuasively argued, there are strong reasons to favour an epistemicist view of vagueness, one according to which vague expressions such as 'tall' must pick a determinate, sharp property, though we do not know which one.[46]

However, Williamson's epistemicist view seems insufficient to us. The main reason for this is that we are not convinced by Williamson's explanation for why we do not and cannot

[44] It is worth noting that the actual world might well contain examples of this kind. For examples, according to some inteprertations of quatum field theory, there are in fact cases of distinct yet completely qualitively identical particles.
[45] This is not to say that one cannot try to provide an account of how such statements can be truthful even if 'i' fails to refer. The point here is that such an account will be needed, and that one cannot simply be content by saying that this is a case of reference failure.
[46] See especially Williamson (1994).

know the sharp boundaries of vague expressions.[47] An alternative proposal is to assume that the reference of vague terms is determined arbitrarily. Consider the predicate 'tall' and the range of admissible precisfications for it. The proposal is that 'tall' had its reference fixed arbitrarily to one of these properties. This means, that 'tall' refers to a particular sharp property (so it does not make for any failures of classical logic). But the fact that we do not and cannot know which property 'tall' denotes has nothing in particular to do with vagueness: it has to do with the ignorance associated with applications of AR more generally (see §1 of this paper). Thus we obtain a different epistemicist view from Williamson's, one that offers an alternative explanation of our ignorance.

The AR view of vagueness has an additional advantage. We suspect that part of the uneasiness many philosophers have with the epistemicist position stems precisely from the intuition that if 'tall' were to refer to a particular property it would be purely arbitrary which one. How could it be that our linguistic practices are sufficient to determine that 'tall' refers to the property of being over 1.80001cm rather than over 1.80002cm? The AR proposal can accommodate this intuition: it is indeed the case that nothing in our practices is sufficient to decide between each of the options, and the choice between them is indeed arbitrary. Still, this is consistent with the epistemicist claim that 'tall' determinately refers to a particular sharp property.[48]

§4 Conclusion

AR may strike many as a highly radical claim. We hope to have shown, however, that the claim is not only conceptually coherent, but also carries considerable explanatory power: it can be used to provide a good account of instantial reasoning, as well as pave the way to solutions to some of the most fundamental problems in philosophy of language such as vagueness, reference to indiscernibles, and Benacerraf's problem.[49] Many of these applications require further development, but if they are ultimately successful, this should give us serious reason to believe that it is, after all, possible to refer arbitrarily.[50]

---

[47] See Kearns & Magidor (2008) for a detailed argument.

[48] One difficult question, however, is how the AR proposal would address the problem of higher-order vagueness. We leave this issue for further work.

[49] Further potential applications of AR that we have not discussed here include the problem of partially defined predicates, the problem of inscrutability of reference, the semantics of variables, the semantics of pronouns, and the semantics of conditionals.

[50] We are grateful to audiences at Cornell University, CSMN Oslo, Macquarie University, and the University of Oxford, as well as to Ross Cameron, John Hawthorne, Stephen Kearns, Jeff King, Vann McGee, Moritz Schulz, Stewart Shapiro, Nick Smith, and Robbie Williams for helpful discussions.

References

Benacerraf, P. (1965): 'What numbers could not be', *Philosophical Review* 74, pp. 47-73.

Cameron, R. (2010): 'Vagueness and naturalness', *Erkenntnis* 72, pp. 281-193.

Fara, D. (2000): 'Shifting sands: An interest-relative theory of vagueness', *Philosophical Topics* 28, pp. 45-81. Originally published under 'Delia Graff'.

Fine, K. (1983): 'A defence of arbitrary objects', *Proceedings of the Aristotelian Society*, Supplementary Volume 57, pp. 55-77.

Fine, K. (1985a): 'Natural deduction and arbitrary objects', *Journal of Philosophical Logic* 14, pp. 57-107.

Fine, K. (1985b): *Reasoning with Arbitrary Objects*, Basil Blackwell

Hawthorne, J. and Manley, D. (MS): *The Reference Book*, unpublished manuscript.

Kearns, S., and Magidor, O. (2008): 'Epistemicism about vagueness and meta-linguistic safety', *Philosophical Perspectives* 22, pp. 277-304

Kearns, S., and Magidor, O. (*forthcoming*): 'Semantic Sovereignty', forthcoming in *Philosophy and Phenomenological Research*

King, J. (1991): 'Instantial terms, anaphora and arbitrary objects', *Philosophical Studies* 61, pp. 239-65.

Mackie, J. (1959): 'The rules of natural deduction', *Analysis* 19, pp. 27-35.

Martino, E. (2001): 'Arbitrary reference in mathematical reasoning', *Topoi* 20, pp. 65-77.

Prawitz, D. (1965): *Natural Deduction: A Proof-theoretical Study*, Almqvist & Wiksell.

Price, R. (1962): ‘Arbitrary individuals and natural deduction’, *Analysis* 22, pp. 94-96.
Rescher, N. (1958): ‘Can there be random individuals’, *Analysis* 18, pp. 114-117.

Shapiro, S. C. (2004): ‘A logic of arbitrary and indefinite objects’, in D. Dubois, C. Welty, & M. Williams (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004)*, AAAI Press, pp. 565-575.

Tennant, N. (1983): ‘A defence of arbitrary objects’, *Proceedings of the Aristotelian Society*, Supplementary Volume 57, pp. 79-89.

Williamson, T. (1994): *Vagueness*, Routledge

Williamson, T. (1997): ‘Précis of vagueness’, *Philosophy and Phenomenological Research* 57, pp. 921-928.